

Alibaba Cloud

Apsara Stack Enterprise

Apsara Stack Enterprise
User Guide - Cloud Essentials
and Security

Product Version: 2105, Internal: V3.14.0

Document Version: 20210820

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.
3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).
5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.
6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

| Style | Description | Example |
|--|---|---|
|  Danger | A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results. |  Danger: Resetting will result in the loss of user configuration data. |
|  Warning | A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results. |  Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance. |
|  Notice | A caution notice indicates warning information, supplementary instructions, and other content that the user must understand. |  Notice: If the weight is set to 0, the server no longer receives new requests. |
|  Note | A note indicates supplemental instructions, best practices, tips, and other content. |  Note: You can use Ctrl + A to select all files. |
| > | Closing angle brackets are used to indicate a multi-level menu cascade. | Click Settings > Network > Set network type . |
| Bold | Bold formatting is used for buttons, menus, page names, and other UI elements. | Click OK . |
| <code>Courier font</code> | Courier font is used for commands | Run the <code>cd /d C:/window</code> command to enter the Windows system folder. |
| <i>Italic</i> | Italic formatting is used for parameters and variables. | <code>bae log list --instanceid</code> <i>Instance_ID</i> |
| [] or [a b] | This format is used for an optional value, where only one item can be selected. | <code>ipconfig [-all -t]</code> |
| { } or {a b} | This format is used for a required value, where only one item can be selected. | <code>switch {active stand}</code> |

Table of Contents

| | |
|--|----|
| 1. Apsara Uni-manager Management Console | 75 |
| 1.1. What is the Apsara Uni-manager Management Console? | 75 |
| 1.2. User roles and permissions | 75 |
| 1.3. Log on to the Apsara Uni-manager Management Console | 77 |
| 1.4. Web page introduction | 78 |
| 1.5. Initial configuration | 79 |
| 1.5.1. Configuration description | 79 |
| 1.5.2. Configuration process | 80 |
| 1.6. Monitoring | 81 |
| 1.6.1. View the workbench | 81 |
| 1.6.2. CloudMonitor | 82 |
| 1.6.2.1. Cloud Monitor overview | 82 |
| 1.6.2.2. Metrics | 82 |
| 1.6.2.3. View monitoring charts | 95 |
| 1.6.3. Alerts | 95 |
| 1.6.3.1. View alert overview | 95 |
| 1.6.3.2. Enable or disable alert notification | 95 |
| 1.6.3.3. View alert logs | 95 |
| 1.6.3.4. Alarm rules | 96 |
| 1.6.3.4.1. View alert rules | 96 |
| 1.6.3.4.2. Create an alert rule | 96 |
| 1.6.3.4.3. Disable an alert rule | 98 |
| 1.6.3.4.4. Enable an alert rule | 98 |
| 1.6.3.4.5. Delete an alert rule | 98 |
| 1.7. VMware Cloud on Alibaba Cloud | 99 |
| 1.7.1. Log on to the VMware Cloud on Alibaba Cloud console | 99 |

| | |
|--|-----|
| 1.7.2. Bind a VMware Cloud on Alibaba Cloud region | 99 |
| 1.7.3. Instructions | 99 |
| 1.7.3.1. Limits | 100 |
| 1.7.3.2. Suggestions | 101 |
| 1.7.4. Instances | 101 |
| 1.7.4.1. Create a VMware Cloud on Alibaba Cloud instance | 101 |
| 1.7.4.2. View instance information | 104 |
| 1.7.4.3. Modify an instance | 105 |
| 1.7.4.4. Remotely connect to an instance | 105 |
| 1.7.4.5. Stop an instance | 106 |
| 1.7.4.6. Start an instance | 106 |
| 1.7.4.7. Restart an instance | 107 |
| 1.7.4.8. Delete an instance | 107 |
| 1.7.4.9. Change the instance type of an instance | 108 |
| 1.7.5. Images | 108 |
| 1.7.5.1. Create a custom image | 108 |
| 1.7.5.2. View images | 109 |
| 1.7.6. Snapshots | 109 |
| 1.7.6.1. Create a snapshot | 109 |
| 1.7.6.2. Delete a snapshot | 110 |
| 1.7.6.3. View snapshots | 110 |
| 1.7.7. Disks | 111 |
| 1.7.7.1. Create a disk | 111 |
| 1.7.7.2. View disks | 112 |
| 1.7.7.3. Detach a data disk | 113 |
| 1.7.8. ENIs | 113 |
| 1.7.8.1. Create an ENI | 113 |
| 1.7.8.2. View ENIs | 114 |

| | |
|---|-----|
| 1.7.8.3. Delete an ENI | 115 |
| 1.8. Enterprise | 115 |
| 1.8.1. Organizations | 115 |
| 1.8.1.1. Create an organization | 115 |
| 1.8.1.2. Query an organization | 116 |
| 1.8.1.3. View organization information | 116 |
| 1.8.1.4. Modify the name of an organization | 116 |
| 1.8.1.5. Change organization ownership | 117 |
| 1.8.1.6. Obtain the AccessKey pair of an organization | 117 |
| 1.8.1.7. Delete an organization | 117 |
| 1.8.2. Resource sets | 118 |
| 1.8.2.1. Create a resource set | 118 |
| 1.8.2.2. View the details of a resource set | 118 |
| 1.8.2.3. Modify the name of a resource set | 118 |
| 1.8.2.4. Add a member to a resource set | 119 |
| 1.8.2.5. Add or remove a user group of a resource set | 119 |
| 1.8.2.6. Delete a resource set | 120 |
| 1.8.3. Roles | 120 |
| 1.8.3.1. Create a custom role | 120 |
| 1.8.3.2. View the details of a role | 121 |
| 1.8.3.3. Modify custom role information | 122 |
| 1.8.3.4. Copy a role | 122 |
| 1.8.3.5. Disable a role | 123 |
| 1.8.3.6. Enable a role | 123 |
| 1.8.3.7. Delete a custom role | 123 |
| 1.8.4. Users | 124 |
| 1.8.4.1. System users | 124 |
| 1.8.4.1.1. Create a user | 124 |

| | |
|---|-----|
| 1.8.4.1.2. Query a user | 125 |
| 1.8.4.1.3. Modify user information | 125 |
| 1.8.4.1.4. Change user roles | 126 |
| 1.8.4.1.5. Modify the information of a user group | 126 |
| 1.8.4.1.6. Modify a user logon policy | 127 |
| 1.8.4.1.7. View the initial password of a user | 127 |
| 1.8.4.1.8. Reset the password of a user | 128 |
| 1.8.4.1.9. Disable or enable a user account | 128 |
| 1.8.4.1.10. Delete a user | 128 |
| 1.8.4.2. Historical users | 129 |
| 1.8.4.2.1. Query historical users | 129 |
| 1.8.4.2.2. Restore historical users | 129 |
| 1.8.5. Logon policies | 129 |
| 1.8.5.1. Create a logon policy | 130 |
| 1.8.5.2. Query a logon policy | 131 |
| 1.8.5.3. Modify a logon policy | 132 |
| 1.8.5.4. Disable a logon policy | 132 |
| 1.8.5.5. Enable a logon policy | 132 |
| 1.8.5.6. Delete a logon policy | 132 |
| 1.8.6. User groups | 133 |
| 1.8.6.1. Create a user group | 133 |
| 1.8.6.2. Add users to a user group | 134 |
| 1.8.6.3. Delete users from a user group | 134 |
| 1.8.6.4. Add a role | 135 |
| 1.8.6.5. Delete a role | 135 |
| 1.8.6.6. Modify the name of a user group | 135 |
| 1.8.6.7. Delete a user group | 136 |
| 1.8.7. Resource pools | 136 |

| | |
|--|-----|
| 1.8.7.1. Update associations | 136 |
| 1.8.8. Change the ownership of an instance | 136 |
| 1.8.9. Cloud instances | 137 |
| 1.8.9.1. Manage Apsara Stack cloud instances | 137 |
| 1.8.9.1.1. Export data of the current cloud | 137 |
| 1.8.9.1.2. Add a secondary Apsara Stack node | 137 |
| 1.8.9.1.3. View managed cloud instances | 139 |
| 1.8.9.1.4. Modify a cloud instance | 139 |
| 1.8.9.1.5. Manage cloud instances | 140 |
| 1.8.9.2. Manage VMware nodes | 140 |
| 1.8.9.2.1. Add a VMware node | 140 |
| 1.8.9.2.2. Modify a VMware node | 141 |
| 1.8.9.2.3. Test VMware node connectivity | 142 |
| 1.8.10. Data permissions | 142 |
| 1.8.10.1. Overview | 142 |
| 1.8.10.2. Set the data permissions of resource instances | 142 |
| 1.8.10.3. Edit user permissions | 143 |
| 1.8.10.4. View the permissions of a user | 143 |
| 1.9. Configurations | 144 |
| 1.9.1. Security policies | 144 |
| 1.9.1.1. Configure password policies | 144 |
| 1.9.1.2. Configure logon control | 144 |
| 1.9.2. Menus | 144 |
| 1.9.2.1. Create a menu | 145 |
| 1.9.2.2. Modify a menu | 146 |
| 1.9.2.3. Delete a menu | 147 |
| 1.9.2.4. Show or hide menus | 147 |
| 1.9.3. Specifications | 148 |

| | |
|---|-----|
| 1.9.3.1. Specification parameters | 148 |
| 1.9.3.2. Create specifications | 151 |
| 1.9.3.3. View specifications | 151 |
| 1.9.3.4. Disable specifications | 152 |
| 1.9.3.5. Export specifications | 152 |
| 1.9.3.6. View specifications of each resource type in previous... | 152 |
| 1.9.4. Message center | 152 |
| 1.9.4.1. View internal messages | 152 |
| 1.9.4.2. Mark messages as read | 153 |
| 1.9.4.3. Delete a message | 153 |
| 1.9.5. Resource pool management | 153 |
| 1.9.6. Custom configurations | 154 |
| 1.9.6.1. Configure brands | 154 |
| 1.10. Operations | 155 |
| 1.10.1. Quotas | 155 |
| 1.10.1.1. Quota parameters | 155 |
| 1.10.1.2. Set quotas for a cloud service | 158 |
| 1.10.1.3. Modify quotas | 159 |
| 1.10.1.4. Reset quotas | 159 |
| 1.10.2. Usage statistics | 159 |
| 1.10.2.1. View the usage statistics of cloud resources | 159 |
| 1.10.3. Statistical analysis | 160 |
| 1.10.3.1. View reports of current data | 160 |
| 1.10.3.2. Export reports of current data | 160 |
| 1.10.3.3. Download reports of historical data | 161 |
| 1.11. Security | 163 |
| 1.11.1. View operation logs | 163 |
| 1.12. RAM | 164 |

| | |
|---|-----|
| 1.12.1. RAM introduction | 164 |
| 1.12.2. Permission policy structure and syntax | 164 |
| 1.12.3. RAM roles | 166 |
| 1.12.3.1. View basic information about a RAM role | 166 |
| 1.12.3.2. Create a RAM role | 166 |
| 1.12.3.3. Create a policy | 167 |
| 1.12.3.4. Modify the content of a RAM policy | 168 |
| 1.12.3.5. Modify the name of a RAM policy | 168 |
| 1.12.3.6. Add a RAM role to a user group | 168 |
| 1.12.3.7. Grant permissions to a RAM role | 169 |
| 1.12.3.8. Remove permissions from a RAM role | 169 |
| 1.12.3.9. Modify a RAM role name | 169 |
| 1.12.3.10. Delete a RAM role | 169 |
| 1.12.4. RAM authorization policies | 170 |
| 1.12.4.1. Create a service-linked role | 170 |
| 1.12.4.2. View the details of a service-linked role | 170 |
| 1.12.4.3. View RAM authorization policies | 170 |
| 1.13. Personal information management | 171 |
| 1.13.1. Modify personal information | 171 |
| 1.13.2. Change the logon password | 171 |
| 1.13.3. Switch the current role | 172 |
| 1.13.4. View the AccessKey pair of your Apsara Stack tenant ... | 173 |
| 1.13.5. Create an AccessKey pair | 173 |
| 1.13.6. Delete an AccessKey pair | 174 |
| 1.13.7. Disable an AccessKey pair | 174 |
| 1.13.8. Enable an AccessKey pair | 174 |
| 1.13.9. MFA | 175 |
| 1.13.9.1. Overview | 175 |

| | |
|--|-----|
| 1.13.9.2. Bind a virtual MFA device to enable MFA | 175 |
| 1.13.9.3. Unbind a virtual MFA device to disable MFA | 176 |
| 1.13.9.4. Forcibly enable MFA | 176 |
| 1.13.9.5. Reset MFA | 176 |
| 2.Elastic Compute Service (ECS) | 177 |
| 2.1. What is ECS? | 177 |
| 2.1.1. Overview | 177 |
| 2.1.2. Instance lifecycle | 177 |
| 2.2. Instructions | 179 |
| 2.2.1. Restrictions | 179 |
| 2.2.2. Suggestions | 179 |
| 2.2.3. Limits | 179 |
| 2.2.4. Notice for Windows users | 180 |
| 2.2.5. Notice for Linux users | 180 |
| 2.2.6. Notice on defense against DDoS attacks | 180 |
| 2.3. Quick start | 180 |
| 2.3.1. Overview | 180 |
| 2.3.2. Log on to the ECS console | 181 |
| 2.3.3. Create a security group | 181 |
| 2.3.4. Create an instance | 183 |
| 2.3.5. Connect to an instance | 187 |
| 2.3.5.1. Instance connecting overview | 187 |
| 2.3.5.2. Connect to a Linux instance by using SSH command | 187 |
| 2.3.5.3. Connect to a Linux-based instance by using remote | 188 |
| 2.3.5.4. Connect to a Windows instance by using RDP | 188 |
| 2.3.5.5. Connect to an ECS instance by using the VNC | 189 |
| 2.4. Instances | 190 |
| 2.4.1. Create an instance | 190 |

| | |
|---|-----|
| 2.4.2. Connect to an instance | 195 |
| 2.4.2.1. Instance connecting overview | 195 |
| 2.4.2.2. Connect to a Linux-based instance by using SSH c... | 196 |
| 2.4.2.3. Connect to a Linux-based instance by using remote... | 196 |
| 2.4.2.4. Connect to a Windows instance by using RDC | 196 |
| 2.4.2.5. Install the certificate for VNC in Windows | 197 |
| 2.4.2.6. Connect to an ECS instance by using the VNC | 199 |
| 2.4.3. View instances | 200 |
| 2.4.4. Modify an instance | 201 |
| 2.4.5. Stop an instance | 201 |
| 2.4.6. Start an instance | 201 |
| 2.4.7. Restart an instance | 202 |
| 2.4.8. Delete an instance | 202 |
| 2.4.9. View the monitoring information of an instance | 203 |
| 2.4.10. Change the instance type of an instance | 203 |
| 2.4.11. Change an instance logon password | 204 |
| 2.4.12. Change the VNC password | 204 |
| 2.4.13. Add an ECS instance to a security group | 204 |
| 2.4.14. Customize instance data | 205 |
| 2.4.15. Modify a private IP address | 207 |
| 2.4.16. Install the CUDA and GPU drivers for a Linux instanc... | 208 |
| 2.4.17. Install the CUDA and GPU drivers for a Windows inst... | 210 |
| 2.5. Disks | 211 |
| 2.5.1. Create a disk | 211 |
| 2.5.2. Attach a disk | 213 |
| 2.5.3. Partition and format disks | 214 |
| 2.5.3.1. Format a data disk for a Linux instance | 214 |
| 2.5.3.2. Format a data disk of a Windows instance | 217 |

| | |
|---|-----|
| 2.5.4. View disks | 217 |
| 2.5.5. Roll back a disk | 218 |
| 2.5.6. Modify the disk properties | 219 |
| 2.5.7. Modify the disk description | 219 |
| 2.5.8. Expand a disk | 220 |
| 2.5.9. Encrypt a disk | 221 |
| 2.5.9.1. Encrypt a system disk | 221 |
| 2.5.9.2. Encrypt a data disk | 222 |
| 2.5.10. Reinitialize a disk | 223 |
| 2.5.11. Detach a data disk | 224 |
| 2.5.12. Release a data disk | 224 |
| 2.6. Images | 225 |
| 2.6.1. Create a custom image | 225 |
| 2.6.2. View images | 225 |
| 2.6.3. View instances related to an image | 226 |
| 2.6.4. Modify the description of a custom image | 226 |
| 2.6.5. Share custom images | 227 |
| 2.6.6. Encrypt a custom image | 227 |
| 2.6.7. Import custom images | 228 |
| 2.6.7.1. Limits on importing custom images | 228 |
| 2.6.7.2. Convert the image file format | 232 |
| 2.6.7.3. Import a custom image | 233 |
| 2.6.8. Export a custom image | 234 |
| 2.6.9. Delete a custom image | 235 |
| 2.7. Snapshots | 235 |
| 2.7.1. Create a snapshot | 235 |
| 2.7.2. View snapshots | 236 |
| 2.7.3. Delete a snapshot | 237 |

| | |
|---|-----|
| 2.8. Automatic snapshot policies | 237 |
| 2.8.1. Create an automatic snapshot policy | 237 |
| 2.8.2. View automatic snapshot policies | 238 |
| 2.8.3. Modify an automatic snapshot policy | 239 |
| 2.8.4. Configure an automatic snapshot policy | 239 |
| 2.8.5. Configure an automatic snapshot policy for multiple di... .. | 240 |
| 2.8.6. Delete an automatic snapshot policy | 240 |
| 2.9. Security groups | 240 |
| 2.9.1. Create a security group | 240 |
| 2.9.2. View security groups | 242 |
| 2.9.3. Modify a security group | 242 |
| 2.9.4. Add a security group rule | 243 |
| 2.9.5. Clone a security group rule | 245 |
| 2.9.6. Modify a security group rule | 245 |
| 2.9.7. Export security group rules | 245 |
| 2.9.8. Import security group rules | 246 |
| 2.9.9. Add an instance | 246 |
| 2.9.10. Remove instances from a security group | 246 |
| 2.9.11. Delete a security group | 247 |
| 2.10. Elastic Network Interfaces | 247 |
| 2.10.1. Create an ENI | 247 |
| 2.10.2. View ENIs | 249 |
| 2.10.3. Modify a secondary ENI | 250 |
| 2.10.4. Bind a secondary ENI to an instance | 250 |
| 2.10.5. Unbind a secondary ENI from an instance | 251 |
| 2.10.6. Delete a secondary ENI | 251 |
| 2.11. Deployment sets | 252 |
| 2.11.1. Create a deployment set | 252 |

| | |
|---|-----|
| 2.11.2. View deployment sets | 253 |
| 2.11.3. Modify a deployment set | 254 |
| 2.11.4. Delete a deployment set | 254 |
| 2.12. Install FTP software | 254 |
| 2.12.1. Overview | 254 |
| 2.12.2. Install and configure vsftpd in CentOS | 254 |
| 2.12.3. Install vsftpd in Ubuntu or Debian | 255 |
| 2.12.4. Build an FTP site in Windows Server 2008 | 256 |
| 2.12.5. Build an FTP site in Windows Server 2012 | 257 |
| 3.Container Service | 258 |
| 3.1. Container Service support for Kubernetes 1.18 | 258 |
| 3.2. What is Container Service? | 258 |
| 3.3. ACK@Edge overview | 259 |
| 3.4. Planning and preparation | 260 |
| 3.5. Quick start | 260 |
| 3.5.1. Procedure | 260 |
| 3.5.2. Log on to the Container Service console | 261 |
| 3.5.3. Create a Kubernetes cluster | 261 |
| 3.5.4. Create an application from an orchestration template | 266 |
| 3.6. Kubernetes clusters | 268 |
| 3.6.1. Authorizations | 268 |
| 3.6.1.1. Assign RBAC roles to a RAM user | 268 |
| 3.6.2. Clusters | 270 |
| 3.6.2.1. Create a Kubernetes cluster | 270 |
| 3.6.2.2. Create an edge Kubernetes cluster | 274 |
| 3.6.2.3. View log files of a cluster | 279 |
| 3.6.2.4. Connect to a cluster through kubectl | 280 |
| 3.6.2.5. Connect to a master node by using SSH | 281 |

| | |
|---|-----|
| 3.6.2.6. Expand a cluster | 282 |
| 3.6.2.7. Renew a certificate | 283 |
| 3.6.2.8. Delete a Kubernetes cluster | 283 |
| 3.6.2.9. View cluster overview | 284 |
| 3.6.3. Nodes | 285 |
| 3.6.3.1. Add existing nodes to a Kubernetes cluster | 285 |
| 3.6.3.2. Add nodes to an edge Kubernetes cluster | 286 |
| 3.6.3.3. View nodes | 289 |
| 3.6.3.4. Manage node labels | 289 |
| 3.6.3.5. Set node schedulability | 290 |
| 3.6.3.6. Remove a node | 291 |
| 3.6.3.7. View node resource usage | 292 |
| 3.6.3.8. Upgrade the NVIDIA driver on a GPU node | 293 |
| 3.6.3.9. GPU scheduling for Kubernetes clusters with GPU-a... | 296 |
| 3.6.3.10. Use labels to schedule pods to GPU-accelerated n... | 301 |
| 3.6.3.11. Manually upgrade the kernel of a GPU node in a ... | 303 |
| 3.6.3.12. Node pools | 305 |
| 3.6.3.12.1. Create a node pool | 305 |
| 3.6.3.12.2. Scale out a node pool | 306 |
| 3.6.3.12.3. Schedule an application pod to a specific node.. | 307 |
| 3.6.4. Storage | 310 |
| 3.6.4.1. Overview | 310 |
| 3.6.4.2. Mount disk volumes | 310 |
| 3.6.4.3. Mount NAS volumes | 315 |
| 3.6.4.4. Mount OSS volumes | 322 |
| 3.6.4.5. Create a PVC | 326 |
| 3.6.4.6. Use PVCs | 327 |
| 3.6.5. Network management | 328 |

| | |
|--|-----|
| 3.6.5.1. Set access control for pods | 328 |
| 3.6.5.2. Set bandwidth limits for pods | 330 |
| 3.6.5.3. Work with Terway | 331 |
| 3.6.6. Namespaces | 336 |
| 3.6.6.1. Create a namespace | 336 |
| 3.6.6.2. Set resource quotas and limits | 337 |
| 3.6.6.3. Modify a namespace | 339 |
| 3.6.6.4. Delete a namespace | 339 |
| 3.6.7. Applications | 340 |
| 3.6.7.1. Create an application from an image | 340 |
| 3.6.7.2. Create an application from an orchestration templa... | 349 |
| 3.6.7.3. Use commands to manage applications | 351 |
| 3.6.7.4. Create a Service | 351 |
| 3.6.7.5. View a Service | 353 |
| 3.6.7.6. Update a Service | 353 |
| 3.6.7.7. Delete a Service | 354 |
| 3.6.7.8. Use a trigger to redeploy an application | 354 |
| 3.6.7.9. View pods | 355 |
| 3.6.7.10. Manage pods | 356 |
| 3.6.7.11. Schedule pods to specific nodes | 358 |
| 3.6.7.12. Simplify application deployment by using Helm | 359 |
| 3.6.8. SLB and Ingress | 362 |
| 3.6.8.1. Overview | 362 |
| 3.6.8.2. Use SLB to access Services | 362 |
| 3.6.8.3. Configure Ingress monitoring | 365 |
| 3.6.8.4. Ingress support | 366 |
| 3.6.8.5. Ingress configurations | 370 |
| 3.6.8.6. Create an Ingress in the console | 372 |

| | |
|---|-----|
| 3.6.8.7. Update an Ingress | 378 |
| 3.6.8.8. Delete an Ingress | 378 |
| 3.6.9. Config maps and secrets | 379 |
| 3.6.9.1. Create a ConfigMap | 379 |
| 3.6.9.2. Use a ConfigMap in a pod | 380 |
| 3.6.9.3. Update a ConfigMap | 385 |
| 3.6.9.4. Delete a ConfigMap | 385 |
| 3.6.9.5. Create a Secret | 385 |
| 3.6.9.6. Modify a Secret | 386 |
| 3.6.9.7. Delete a Secret | 387 |
| 3.6.10. Templates | 387 |
| 3.6.10.1. Create an orchestration template | 387 |
| 3.6.10.2. Update an orchestration template | 389 |
| 3.6.10.3. Save an orchestration template as a new one | 390 |
| 3.6.10.4. Download an orchestration template | 390 |
| 3.6.10.5. Delete an orchestration template | 391 |
| 3.6.11. Auto scaling | 391 |
| 3.6.11.1. Auto scaling of nodes | 391 |
| 3.6.11.2. Horizontal pod autoscaling | 396 |
| 3.6.12. Sandboxed-containers | 399 |
| 3.6.12.1. Overview | 399 |
| 3.6.12.2. Create a Kubernetes cluster that runs sandboxed c...----- | 400 |
| 3.6.12.3. Expand a Container Service cluster that runs sand...----- | 404 |
| 3.6.12.4. Create an application that runs in sandboxed cont...----- | 406 |
| 3.6.12.5. Configure a Kubernetes cluster that runs both san...----- | 415 |
| 3.6.12.6. How do I select between Docker and Sandboxed-C...----- | 417 |
| 3.6.12.7. Benefits of Sandboxed-Container | 420 |
| 3.6.12.8. Differences between runC and runV | 424 |

| | |
|---|-----|
| 3.6.12.9. Compatibility notes | 427 |
| 3.6.13. Use the Kubernetes event center | 429 |
| 3.6.14. Use Log Service to collect log data from containers | 430 |
| 4.Auto Scaling (ESS) | 441 |
| 4.1. What is Auto Scaling? | 441 |
| 4.2. Notes | 442 |
| 4.2.1. Precautions | 442 |
| 4.2.2. Manual intervention | 443 |
| 4.2.3. Limits | 444 |
| 4.2.4. Scaling group status | 444 |
| 4.2.5. Scaling processes | 445 |
| 4.2.6. Remove unhealthy ECS instances | 446 |
| 4.2.7. Instance rollback after a failed scaling activity | 446 |
| 4.2.8. Instance lifecycle management | 446 |
| 4.3. Quick start | 447 |
| 4.3.1. Overview | 447 |
| 4.3.2. Log on to the Auto Scaling console | 447 |
| 4.3.3. Create a scaling group | 448 |
| 4.3.4. Create a scaling configuration | 450 |
| 4.3.5. Enable a scaling group | 452 |
| 4.3.6. Create a scaling rule | 453 |
| 4.3.7. Create a scheduled task | 453 |
| 4.3.8. Create an event-triggered task | 454 |
| 4.4. Scaling groups | 456 |
| 4.4.1. Create a scaling group | 456 |
| 4.4.2. Enable a scaling group | 458 |
| 4.4.3. View scaling groups | 458 |
| 4.4.4. Modify a scaling group | 459 |

| | |
|---|-----|
| 4.4.5. Disable a scaling group | 460 |
| 4.4.6. Delete a scaling group | 460 |
| 4.4.7. Query ECS instances | 460 |
| 4.4.8. Put an ECS instance into the Standby state | 461 |
| 4.4.9. Remove an ECS instance from the Standby state | 462 |
| 4.4.10. Put an ECS instance into the Protected state | 462 |
| 4.4.11. Remove an ECS instance from the Protected state | 463 |
| 4.5. Scaling configurations | 463 |
| 4.5.1. Create a scaling configuration | 463 |
| 4.5.2. View scaling configurations | 465 |
| 4.5.3. Modify a scaling configuration | 466 |
| 4.5.4. Apply a scaling configuration | 466 |
| 4.5.5. Delete a scaling configuration | 466 |
| 4.6. Scaling rules | 467 |
| 4.6.1. Create a scaling rule | 467 |
| 4.6.2. View scaling rules | 467 |
| 4.6.3. Modify a scaling rule | 468 |
| 4.6.4. Delete a scaling rule | 468 |
| 4.7. Scaling tasks | 468 |
| 4.7.1. Manually execute a scaling rule | 468 |
| 4.7.2. Manually add an ECS instance | 469 |
| 4.7.3. Manually remove an ECS instance | 470 |
| 4.8. Scheduled tasks | 470 |
| 4.8.1. Create a scheduled task | 470 |
| 4.8.2. View scheduled tasks | 471 |
| 4.8.3. Modify a scheduled task | 472 |
| 4.8.4. Disable a scheduled task | 472 |
| 4.8.5. Enable a scheduled task | 472 |

| | |
|--|-----|
| 4.8.6. Delete a scheduled task | 473 |
| 4.9. Event-triggered tasks | 473 |
| 4.9.1. Create an event-triggered task | 473 |
| 4.9.2. View event-triggered tasks | 474 |
| 4.9.3. Modify an event-triggered task | 475 |
| 4.9.4. Disable an event-triggered task | 475 |
| 4.9.5. Enable an event-triggered task | 475 |
| 4.9.6. Delete an event-triggered task | 476 |
| 5.Resource Orchestration Service (ROS) | 477 |
| 5.1. What is ROS? | 477 |
| 5.2. Log on to the ROS console | 477 |
| 5.3. Create a stack | 478 |
| 5.4. Template syntax | 478 |
| 5.4.1. Template structure | 478 |
| 5.4.2. Parameters | 480 |
| 5.4.3. Resources | 483 |
| 5.4.4. Outputs | 487 |
| 5.4.5. Functions | 489 |
| 5.4.6. Mappings | 510 |
| 5.4.7. Conditions | 511 |
| 5.5. Resource types | 513 |
| 5.5.1. ECS | 513 |
| 5.5.1.1. ALIYUN::ECS::AutoSnapshotPolicy | 514 |
| 5.5.1.2. ALIYUN::ECS::BandwidthPackage | 517 |
| 5.5.1.3. ALIYUN::ECS::Command | 518 |
| 5.5.1.4. ALIYUN::ECS::CustomImage | 521 |
| 5.5.1.5. ALIYUN::ECS::DedicatedHost | 526 |
| 5.5.1.6. ALIYUN::ECS::Disk | 533 |

| | |
|---|-----|
| 5.5.1.7. ALIYUN::ECS::DiskAttachment | 537 |
| 5.5.1.8. ALIYUN::ECS::ForwardEntry | 539 |
| 5.5.1.9. ALIYUN::ECS::Instance | 541 |
| 5.5.1.10. ALIYUN::ECS::InstanceClone | 549 |
| 5.5.1.11. ALIYUN::ECS::InstanceGroup | 556 |
| 5.5.1.12. ALIYUN::ECS::InstanceGroupClone | 566 |
| 5.5.1.13. ALIYUN::ECS::Invocation | 576 |
| 5.5.1.14. ALIYUN::ECS::JoinSecurityGroup | 578 |
| 5.5.1.15. ALIYUN::ECS::LaunchTemplate | 579 |
| 5.5.1.16. ALIYUN::ECS::NatGateway | 589 |
| 5.5.1.17. ALIYUN::ECS::NetworkInterface | 591 |
| 5.5.1.18. ALIYUN::ECS::NetworkInterfaceAttachment | 595 |
| 5.5.1.19. ALIYUN::ECS::NetworkInterfacePermission | 596 |
| 5.5.1.20. ALIYUN::ECS::Route | 598 |
| 5.5.1.21. ALIYUN::ECS::SNatEntry | 601 |
| 5.5.1.22. ALIYUN::ECS::SecurityGroup | 602 |
| 5.5.1.23. ALIYUN::ECS::SecurityGroupClone | 614 |
| 5.5.1.24. ALIYUN::ECS::SecurityGroupEgress | 617 |
| 5.5.1.25. ALIYUN::ECS::SecurityGroupIngress | 622 |
| 5.5.1.26. ALIYUN::ECS::Snapshot | 627 |
| 5.5.1.27. ALIYUN::ECS::SSHKeyPair | 629 |
| 5.5.1.28. ALIYUN::ECS::SSHKeyPairAttachment | 631 |
| 5.5.1.29. ALIYUN::ECS::VPC | 632 |
| 5.5.1.30. ALIYUN::ECS::VSwitch | 635 |
| 5.5.2. ESS | 638 |
| 5.5.2.1. ALIYUN::ESS::AlarmTask | 638 |
| 5.5.2.2. ALIYUN::ESS::AlarmTaskEnable | 643 |
| 5.5.2.3. ALIYUN::ESS::LifecycleHook | 644 |

| | |
|---|-----|
| 5.5.2.4. ALIYUN::ESS::ScalingConfiguration | 649 |
| 5.5.2.5. ALIYUN::ESS::ScalingGroup | 657 |
| 5.5.2.6. ALIYUN::ESS::ScalingGroupEnable | 665 |
| 5.5.2.7. ALIYUN::ESS::ScalingRule | 667 |
| 5.5.2.8. ALIYUN::ESS::ScheduledTask | 670 |
| 5.5.3. OSS | 674 |
| 5.5.3.1. ALIYUN::OSS::Bucket | 674 |
| 5.5.4. RDS | 682 |
| 5.5.4.1. ALIYUN::RDS::Account | 683 |
| 5.5.4.2. ALIYUN::RDS::AccountPrivilege | 685 |
| 5.5.4.3. ALIYUN::RDS::DBInstance | 688 |
| 5.5.4.4. ALIYUN::RDS::DBInstanceParameterGroup | 696 |
| 5.5.4.5. ALIYUN::RDS::DBInstanceSecurityIps | 698 |
| 5.5.4.6. ALIYUN::RDS::PrepayDBInstance | 700 |
| 5.5.5. ROS | 713 |
| 5.5.5.1. ALIYUN::ROS::WaitCondition | 713 |
| 5.5.5.2. ALIYUN::ROS::WaitConditionHandle | 715 |
| 5.5.5.3. ALIYUN::ROS::Stack | 718 |
| 5.5.6. SLB | 726 |
| 5.5.6.1. ALIYUN::SLB::AccessControl | 726 |
| 5.5.6.2. ALIYUN::SLB::BackendServerAttachment | 730 |
| 5.5.6.3. ALIYUN::SLB::BackendServerToVServerGroupAddition | 732 |
| 5.5.6.4. ALIYUN::SLB::Certificate | 734 |
| 5.5.6.5. ALIYUN::SLB::DomainExtension | 737 |
| 5.5.6.6. ALIYUN::SLB::Listener | 738 |
| 5.5.6.7. ALIYUN::SLB::LoadBalancer | 751 |
| 5.5.6.8. ALIYUN::SLB::LoadBalancerClone | 757 |
| 5.5.6.9. ALIYUN::SLB::MasterSlaveServerGroup | 761 |

| | |
|--|-----|
| 5.5.6.10. ALIYUN::SLB::Rule | 764 |
| 5.5.6.11. ALIYUN::SLB::VServerGroup | 767 |
| 5.5.7. VPC | 769 |
| 5.5.7.1. ALIYUN::VPC::EIP | 769 |
| 5.5.7.2. ALIYUN::VPC::EIPAssociation | 772 |
| 5.5.7.3. ALIYUN::VPC::PeeringRouterInterfaceBinding | 775 |
| 5.5.7.4. ALIYUN::VPC::PeeringRouterInterfaceConnection | 776 |
| 5.5.7.5. ALIYUN::VPC::RouterInterface | 777 |
| 6.Object Storage Service (OSS) | 783 |
| 6.1. What is OSS? | 783 |
| 6.2. Usage notes | 783 |
| 6.3. Quick start | 783 |
| 6.3.1. Log on to the OSS console | 784 |
| 6.3.2. Create buckets | 784 |
| 6.3.3. Upload objects | 786 |
| 6.3.4. Obtain object URLs | 787 |
| 6.4. Buckets | 787 |
| 6.4.1. View bucket information | 787 |
| 6.4.2. Delete a bucket | 788 |
| 6.4.3. Modify bucket ACLs | 788 |
| 6.4.4. Configure static website hosting | 789 |
| 6.4.5. Configure hotlink protection | 790 |
| 6.4.6. Configure logging | 791 |
| 6.4.7. Configure CORS | 791 |
| 6.4.8. Configure lifecycle rules | 793 |
| 6.4.9. Configure storage quota | 794 |
| 6.4.10. Configure back-to-origin rules | 795 |
| 6.4.11. Configure server-side encryption | 801 |

| | |
|--|-----|
| 6.4.12. Bind a bucket to a VPC network | 802 |
| 6.4.13. Configure CRR | 802 |
| 6.5. Objects | 804 |
| 6.5.1. Search for objects | 804 |
| 6.5.2. Configure object ACLs | 804 |
| 6.5.3. Create folders | 805 |
| 6.5.4. Delete objects | 806 |
| 6.5.5. Manage parts | 806 |
| 6.5.6. Configure object tagging | 806 |
| 6.6. Create single tunnels | 807 |
| 6.7. Add OSS paths | 808 |
| 7. Tablestore | 809 |
| 7.1. What is Tablestore? | 809 |
| 7.2. Precautions | 809 |
| 7.3. Quick start | 810 |
| 7.3.1. Log on to the Tablestore console | 810 |
| 7.3.2. Create an instance | 811 |
| 7.3.3. Create tables | 812 |
| 7.3.4. Read and write data in the console | 815 |
| 7.3.5. Bind a VPC to a Tablestore instance | 817 |
| 8. ApsaraDB RDS for MySQL | 819 |
| 8.1. What is ApsaraDB RDS? | 819 |
| 8.2. Log on to the ApsaraDB RDS console | 819 |
| 8.3. Quick start | 820 |
| 8.3.1. Limits | 820 |
| 8.3.2. Procedure | 821 |
| 8.3.3. Create an instance | 822 |
| 8.3.4. Initialization settings | 824 |

| | |
|--|-----|
| 8.3.4.1. Configure a whitelist | 824 |
| 8.3.4.2. Create an account | 826 |
| 8.3.4.3. Create a database | 830 |
| 8.3.5. Connect to an ApsaraDB RDS for MySQL instance | 831 |
| 8.4. Instances | 832 |
| 8.4.1. Create an instance | 832 |
| 8.4.2. Create an ApsaraDB RDS for MySQL instance with stan.. | 834 |
| 8.4.3. View basic information of an instance | 837 |
| 8.4.4. Restart an instance | 837 |
| 8.4.5. Change the specifications of an instance | 837 |
| 8.4.6. Set a maintenance window | 838 |
| 8.4.7. Change the data replication mode | 838 |
| 8.4.8. Release an instance | 839 |
| 8.4.9. Upgrade the minor version of an instance | 839 |
| 8.4.10. Modify parameters of an instance | 840 |
| 8.4.11. Read-only instances | 842 |
| 8.4.11.1. Overview of read-only instances | 842 |
| 8.4.11.2. Create a read-only instance | 843 |
| 8.4.11.3. View details of read-only instances | 844 |
| 8.5. Accounts | 845 |
| 8.5.1. Create an account | 845 |
| 8.5.2. Reset the password | 849 |
| 8.5.3. Modify account permissions | 849 |
| 8.5.4. Delete an account | 850 |
| 8.6. Databases | 850 |
| 8.6.1. Create a database | 850 |
| 8.6.2. Delete a database | 851 |
| 8.7. Database connection | 851 |

| | |
|--|-----|
| 8.7.1. Change the endpoint and port number of an instance | 851 |
| 8.7.2. Log on to an ApsaraDB RDS instance by using DMS | 852 |
| 8.7.3. Hybrid access from both the classic network and VPCs | 854 |
| 8.7.4. Change the network type of an instance | 856 |
| 8.7.5. Switch an ApsaraDB RDS for MySQL instance to a new... | 857 |
| 8.8. Database proxy | 857 |
| 8.8.1. Dedicated proxy | 857 |
| 8.8.2. Short-lived connection optimization | 861 |
| 8.8.3. Transaction splitting | 861 |
| 8.8.4. Read/write splitting | 863 |
| 8.8.4.1. Enable read/write splitting | 863 |
| 8.8.4.2. Configure read/write splitting | 866 |
| 8.8.4.3. Disable read/write splitting | 867 |
| 8.8.4.4. Upgrade an ApsaraDB RDS for MySQL instance fro... | 867 |
| 8.9. Monitoring and alerts | 869 |
| 8.9.1. View resource and engine monitoring data | 869 |
| 8.9.2. Set a monitoring frequency | 871 |
| 8.10. Data security | 872 |
| 8.10.1. Configure a whitelist | 872 |
| 8.10.2. Configure SSL encryption | 874 |
| 8.10.3. Configure TDE | 877 |
| 8.10.4. SQL audit | 880 |
| 8.11. Service availability | 881 |
| 8.11.1. Configure automatic or manual switchover | 881 |
| 8.11.2. Change the data replication mode | 882 |
| 8.12. Database backup and restoration | 883 |
| 8.12.1. Automatic backup | 883 |
| 8.12.2. Manual backup | 885 |

| | |
|--|-----|
| 8.12.3. Restore individual databases and tables for an Apsara... | 885 |
| 8.12.4. Download data and log backup files | 887 |
| 8.12.5. Upload binlogs | 889 |
| 8.12.6. Restore data to a new instance (formerly known as cl... | 889 |
| 8.13. CloudDBA | 891 |
| 8.13.1. Introduction to CloudDBA | 892 |
| 8.13.2. Diagnostics | 892 |
| 8.13.3. Session management | 893 |
| 8.13.4. Real-time monitoring | 893 |
| 8.13.5. Storage analysis | 893 |
| 8.13.6. Deadlock analysis | 894 |
| 8.13.7. Dashboard | 894 |
| 8.13.8. Slow query logs | 894 |
| 8.13.9. Diagnostic reports | 895 |
| 8.14. Logs | 895 |
| 8.15. Use mysqldump to migrate MySQL data | 895 |
| 9. ApsaraDB RDS for SQL Server | 898 |
| 9.1. What is ApsaraDB RDS? | 898 |
| 9.2. Log on to the ApsaraDB RDS console | 898 |
| 9.3. Quick Start | 898 |
| 9.3.1. Procedure | 898 |
| 9.3.2. Create an instance | 899 |
| 9.3.3. Configure an IP address whitelist | 902 |
| 9.3.4. Connect to an instance | 903 |
| 9.3.5. Create an account | 904 |
| 9.3.6. Create a database | 906 |
| 9.4. Instances | 906 |
| 9.4.1. Create an instance | 906 |

| | |
|---|-----|
| 9.4.2. View basic information of an instance | 909 |
| 9.4.3. Restart an instance | 909 |
| 9.4.4. Change the specifications of an instance | 909 |
| 9.4.5. Set a maintenance window | 910 |
| 9.4.6. Configure primary/secondary switchover | 910 |
| 9.4.7. Release an instance | 911 |
| 9.4.8. Read-only instances | 911 |
| 9.4.8.1. Overview of read-only ApsaraDB RDS for SQL Serve... | 911 |
| 9.4.8.2. Create a read-only ApsaraDB RDS for SQL Server in.. | 913 |
| 9.4.8.3. View details of read-only instances | 914 |
| 9.5. Accounts | 915 |
| 9.5.1. Create an account | 915 |
| 9.5.2. Reset the password | 916 |
| 9.6. Databases | 917 |
| 9.6.1. Create a database | 917 |
| 9.6.2. Delete a database | 917 |
| 9.6.3. Change the character set collation and the time zone ... | 919 |
| 9.7. Database connection | 922 |
| 9.7.1. Change the endpoint and port number of an instance | 922 |
| 9.7.2. Connect to an instance | 923 |
| 9.8. Monitoring and alerting | 924 |
| 9.8.1. Set a monitoring frequency | 924 |
| 9.8.2. View resource and engine monitoring data | 924 |
| 9.9. Data security | 925 |
| 9.9.1. Configure an IP address whitelist | 925 |
| 9.9.2. Configure SSL encryption | 927 |
| 9.9.3. Configure TDE | 929 |
| 9.10. Database backup and restoration | 930 |

| | |
|---|-----|
| 9.10.1. Configure an automatic backup policy | 930 |
| 9.10.2. Manually back up an instance | 931 |
| 9.10.3. Shrink transaction logs | 931 |
| 9.11. Migrate full backup data to ApsaraDB RDS for SQL Server | 932 |
| 10. ApsaraDB RDS for PostgreSQL | 936 |
| 10.1. What is ApsaraDB RDS? | 936 |
| 10.2. Limits on ApsaraDB RDS for PostgreSQL | 936 |
| 10.3. Log on to the ApsaraDB RDS console | 936 |
| 10.4. Quick Start | 937 |
| 10.4.1. Procedure | 937 |
| 10.4.2. Create an instance | 938 |
| 10.4.3. Configure an IP address whitelist | 940 |
| 10.4.4. Create a database and an account | 941 |
| 10.4.5. Connect to an ApsaraDB RDS for PostgreSQL instance | 945 |
| 10.5. Instances | 946 |
| 10.5.1. Create an instance | 946 |
| 10.5.2. Create an ApsaraDB RDS for PostgreSQL instance tha... | 949 |
| 10.5.3. View basic information of an instance | 951 |
| 10.5.4. Restart an instance | 952 |
| 10.5.5. Change the specifications of an instance | 952 |
| 10.5.6. Set a maintenance window | 952 |
| 10.5.7. Configure primary/secondary switchover | 953 |
| 10.5.8. Release an instance | 954 |
| 10.5.9. Modify parameters of an instance | 954 |
| 10.5.10. Read-only instances | 955 |
| 10.5.10.1. Overview of read-only ApsaraDB RDS for PostgreS... | 955 |
| 10.5.10.2. Create a read-only ApsaraDB RDS for PostgreSQL... | 957 |
| 10.5.10.3. View a read-only ApsaraDB RDS for PostgreSQL i... | 958 |

| | |
|--|-----|
| 10.6. Database connection | 959 |
| 10.6.1. Connect to an ApsaraDB RDS for PostgreSQL instance | 959 |
| 10.6.2. Use DMS to log on to an ApsaraDB RDS instance | 960 |
| 10.6.3. View and modify the internal endpoint and port num... .. | 961 |
| 10.7. Accounts | 962 |
| 10.7.1. Create an account | 962 |
| 10.7.2. Reset the password | 966 |
| 10.8. Databases | 966 |
| 10.8.1. Create a database | 966 |
| 10.8.2. Delete a database | 968 |
| 10.9. Networks, VPCs, and vSwitches | 969 |
| 10.9.1. Change the network type of an ApsaraDB RDS for Po... .. | 969 |
| 10.9.2. Configure hybrid access from both the classic network.. .. | 971 |
| 10.10. Monitoring | 973 |
| 10.10.1. View monitored resources | 973 |
| 10.10.2. Set a monitoring frequency | 974 |
| 10.11. Data security | 974 |
| 10.11.1. Switch to the enhanced whitelist mode | 974 |
| 10.11.2. Configure an IP address whitelist | 975 |
| 10.11.3. Configure SSL encryption | 976 |
| 10.11.4. Configure data encryption | 977 |
| 10.12. Logs and audit | 979 |
| 10.12.1. Configure SQL audit | 979 |
| 10.12.2. Manage logs | 980 |
| 10.13. Backup | 981 |
| 10.13.1. Back up an ApsaraDB RDS for PostgreSQL instance | 981 |
| 10.13.2. Download data and log backup files | 982 |
| 10.13.3. Create a logical backup for an ApsaraDB RDS for Po... .. | 983 |

| | |
|--|------|
| 10.13.4. Create a full backup of an ApsaraDB RDS for Postgr... | 987 |
| 10.14. Restoration | 988 |
| 10.14.1. Restore data of an ApsaraDB RDS for PostgreSQL ins... | 988 |
| 10.14.2. Restore data from a logical backup file | 990 |
| 10.15. Plug-ins | 993 |
| 10.15.1. Plug-ins supported | 993 |
| 10.15.2. Use mysql_fdw to read data from and write data to... | 1001 |
| 10.15.3. Use oss_fdw to read and write foreign data files | 1003 |
| 10.16. Use Pgpool for read/write splitting in ApsaraDB RDS fo... | 1007 |
| 10.17. Use ShardingSphere to develop ApsaraDB RDS for Postg... | 1020 |
| 11. Cloud Native Distributed Database PolarDB-X | 1026 |
| 11.1. What is PolarDB-X? | 1026 |
| 11.2. Quick start | 1026 |
| 11.3. Log on to the PolarDB-X console | 1027 |
| 11.4. Instance management | 1027 |
| 11.4.1. Create an instance | 1027 |
| 11.4.2. Change instance specifications | 1029 |
| 11.4.3. Read-only PolarDB-X instances | 1029 |
| 11.4.3.1. Overview | 1029 |
| 11.4.3.2. Create a read-only PolarDB-X instance | 1030 |
| 11.4.3.3. Manage a read-only PolarDB-X instance | 1031 |
| 11.4.3.4. Release a read-only PolarDB-X instance | 1031 |
| 11.4.4. Restart a PolarDB-X instance | 1032 |
| 11.4.5. Release a PolarDB-X instance | 1032 |
| 11.4.6. Recover data | 1033 |
| 11.4.6.1. Backup and restoration | 1033 |
| 11.4.6.2. Configure an automatic backup policy | 1034 |
| 11.4.6.3. Configure local logs | 1034 |

| | |
|--|------|
| 11.4.6.4. Manual backup | 1035 |
| 11.4.6.5. Recover data | 1035 |
| 11.4.6.6. SQL flashback | 1036 |
| 11.4.6.6.1. Overview | 1036 |
| 11.4.6.6.2. Generate a recovery file | 1036 |
| 11.4.6.6.3. Rollback SQL statements and original SQL stat... | 1038 |
| 11.4.6.6.4. Exact match and fuzzy match | 1038 |
| 11.4.6.7. Table recycle bin | 1039 |
| 11.4.6.7.1. Overview | 1039 |
| 11.4.6.7.2. Enable the table recycle bin | 1040 |
| 11.4.6.7.3. Recover tables | 1040 |
| 11.4.6.7.4. Delete tables from the recycle bin | 1040 |
| 11.4.6.7.5. Disable the table recycle bin | 1041 |
| 11.4.7. Set parameters | 1041 |
| 11.4.8. SQL audit and analysis | 1043 |
| 11.4.8.1. Description | 1043 |
| 11.4.8.2. Enable SQL audit and analysis | 1044 |
| 11.4.8.3. Log fields | 1046 |
| 11.4.8.4. Log analysis | 1047 |
| 11.4.8.5. Log reports | 1052 |
| 11.4.9. Monitor PolarDB-X instances | 1057 |
| 11.4.9.1. View monitoring information | 1057 |
| 11.4.9.2. Monitoring metrics | 1058 |
| 11.4.9.3. How metrics work | 1059 |
| 11.4.9.4. Prevent performance problems | 1060 |
| 11.4.9.4.1. Example 1: PolarDB-X CPU utilization | 1060 |
| 11.4.9.4.2. Example 2: Logical RT and physical RT | 1061 |
| 11.4.9.4.3. Example 3: Logical QPS and physical QPS | 1063 |

| | |
|--|------|
| 11.4.9.4.4. Example 4: High memory usage | 1064 |
| 11.4.10. View the instance version | 1065 |
| 11.5. Account management | 1065 |
| 11.5.1. Terms | 1065 |
| 11.5.2. Create an account | 1067 |
| 11.5.3. Reset the password | 1068 |
| 11.5.4. Modify account permissions | 1069 |
| 11.5.5. Delete an account | 1071 |
| 11.6. Database management | 1072 |
| 11.6.1. Create a database | 1072 |
| 11.6.2. View a database | 1073 |
| 11.6.3. Perform smooth scale-out | 1074 |
| 11.6.4. View database monitoring information | 1076 |
| 11.6.5. Set the IP address whitelist | 1076 |
| 11.6.6. Delete a database | 1077 |
| 11.6.7. Fix database shard connections | 1077 |
| 11.7. Custom control commands | 1078 |
| 11.7.1. Overview | 1078 |
| 11.7.2. Help statements | 1078 |
| 11.7.3. Statements for viewing rules and node topologies | 1078 |
| 11.7.4. SQL tuning statements | 1083 |
| 11.7.5. Statistics query statements | 1089 |
| 11.7.6. SHOW PROCESSLIST and KILL commands | 1093 |
| 11.7.7. SHOW PROCESSLIST and KILL commands in earlier ve... .. | 1096 |
| 11.8. Custom hints | 1098 |
| 11.8.1. Introduction to hints | 1098 |
| 11.8.2. Read/write splitting | 1100 |
| 11.8.3. Specify a timeout period for an SQL statement | 1101 |

| | |
|--|------|
| 11.8.4. Specify a database shard to run an SQL statement | 1101 |
| 11.8.5. Scan all or some of database shards and table shards | 1104 |
| 11.8.6. INDEX HINT | 1106 |
| 11.9. PolarDB-X 5.2 hints | 1107 |
| 11.9.1. Introduction to hints | 1107 |
| 11.9.2. Read/write splitting | 1108 |
| 11.9.3. Prevent the delay from a read-only ApsaraDB RDS for... | 1109 |
| 11.9.4. Specify a timeout period for an SQL statement | 1110 |
| 11.9.5. Specify a database shard to run an SQL statement | 1111 |
| 11.9.6. Scan all database shards and table shards | 1115 |
| 11.10. Distributed transactions | 1116 |
| 11.10.1. Distributed transactions based on MySQL 5.7 | 1116 |
| 11.10.2. Distributed transactions based on MySQL 5.6 | 1117 |
| 11.11. DDL operations | 1118 |
| 11.11.1. DDL statements | 1118 |
| 11.11.2. CREATE TABLE statement | 1118 |
| 11.11.2.1. Overview | 1118 |
| 11.11.2.2. Create a single-database non-partition table | 1119 |
| 11.11.2.3. Create a non-partition table in database shards | 1120 |
| 11.11.2.4. Create table shards in database shards | 1120 |
| 11.11.2.5. Use the primary key as the shard key | 1130 |
| 11.11.2.6. Create a broadcast table | 1131 |
| 11.11.2.7. Other attributes of the MySQL CREATE TABLE sta... | 1131 |
| 11.11.3. ALTER TABLE statement | 1131 |
| 11.11.4. DROP TABLE statement | 1132 |
| 11.11.5. FAQ about DDL statements | 1132 |
| 11.11.6. DDL functions for sharding | 1133 |
| 11.11.6.1. Overview | 1133 |

| | |
|---|------|
| 11.11.6.2. HASH | 1135 |
| 11.11.6.3. UNI_HASH | 1136 |
| 11.11.6.4. RIGHT_SHIFT | 1138 |
| 11.11.6.5. RANGE_HASH | 1139 |
| 11.11.6.6. MM | 1139 |
| 11.11.6.7. DD | 1140 |
| 11.11.6.8. WEEK | 1141 |
| 11.11.6.9. MMDD | 1141 |
| 11.11.6.10. YYYYMM | 1142 |
| 11.11.6.11. YYYYWEEK | 1143 |
| 11.11.6.12. YYYYDD | 1144 |
| 11.11.6.13. YYYYMM_OPT | 1145 |
| 11.11.6.14. YYYYWEEK_OPT | 1147 |
| 11.11.6.15. YYYYDD_OPT | 1147 |
| 11.12. Automatic protection of important SQL statements | 1148 |
| 11.13. PolarDB-X sequence | 1149 |
| 11.13.1. Overview | 1149 |
| 11.13.2. Explicit sequence usage | 1151 |
| 11.13.3. Implicit sequence usage | 1155 |
| 11.13.4. Limits and precautions for sequences | 1157 |
| 11.14. Best practices | 1158 |
| 11.14.1. Select a shard key | 1158 |
| 11.14.2. Select the number of shards | 1159 |
| 11.14.3. Basic concepts of SQL optimization | 1160 |
| 11.14.4. SQL optimization methods | 1164 |
| 11.14.4.1. Overview | 1164 |
| 11.14.4.2. Single-table SQL optimization | 1165 |
| 11.14.4.3. JOIN query optimization | 1169 |

| | |
|---|------|
| 11.14.4.4. Subquery optimization | 1172 |
| 11.14.5. Select connection pools for an application | 1172 |
| 11.14.6. Connections to PolarDB-X instances | 1173 |
| 11.14.7. Perform instance upgrade | 1175 |
| 11.14.8. Perform scale-out | 1176 |
| 11.14.9. Troubleshoot slow SQL statements in DRDS | 1178 |
| 11.14.9.1. Details about a low SQL statement | 1178 |
| 11.14.9.2. Locate slow SQL statements | 1180 |
| 11.14.9.3. Locate nodes with performance loss | 1182 |
| 11.14.9.4. Troubleshoot the performance loss | 1183 |
| 11.14.10. Handle DDL exceptions | 1184 |
| 11.14.11. Efficiently scan DRDS data | 1187 |
| 11.15. Appendix: PolarDB-X terms | 1189 |
| 12. AnalyticDB for PostgreSQL | 1196 |
| 12.1. What is AnalyticDB for PostgreSQL? | 1196 |
| 12.2. Quick start | 1196 |
| 12.2.1. Overview | 1196 |
| 12.2.2. Log on to the AnalyticDB for PostgreSQL console | 1196 |
| 12.2.3. Create an instance | 1197 |
| 12.2.4. Configure a whitelist | 1198 |
| 12.2.5. Create an initial account | 1199 |
| 12.2.6. Obtain client tools | 1199 |
| 12.2.7. Connect to a database | 1200 |
| 12.3. Instances | 1206 |
| 12.3.1. Reset the password | 1206 |
| 12.3.2. View monitoring information | 1206 |
| 12.3.3. Switch the network type of an instance | 1206 |
| 12.3.4. Restart an instance | 1207 |

| | |
|---|------|
| 12.3.5. Import data | 1207 |
| 12.3.5.1. Import data from or export data to OSS in parallel | 1207 |
| 12.3.5.2. Import data from MySQL | 1215 |
| 12.3.5.3. Import data from PostgreSQL | 1217 |
| 12.3.5.4. Use the \COPY statement to import data | 1218 |
| 12.4. Databases | 1219 |
| 12.4.1. Overview | 1219 |
| 12.4.2. Create a database | 1219 |
| 12.4.3. Create a distribution key | 1219 |
| 12.4.4. Construct data | 1220 |
| 12.4.5. Query data | 1220 |
| 12.4.6. Manage extensions | 1221 |
| 12.4.7. Manage users and permissions | 1221 |
| 12.4.8. Manage JSON data | 1222 |
| 12.4.9. Use HyperLogLog | 1229 |
| 12.4.10. Use the CREATE LIBRARY statement | 1230 |
| 12.4.11. Create and use a PL/Java UDF | 1231 |
| 12.5. Table | 1233 |
| 12.5.1. Create a table | 1233 |
| 12.5.2. Principles and scenarios of row store, column store, h... | 1238 |
| 12.5.3. Enable the column store and compression features | 1239 |
| 12.5.4. Add a field to a column store table and set the defa... | 1240 |
| 12.5.5. Configure table partitions | 1242 |
| 12.5.6. Configure the sort key | 1243 |
| 12.6. Best practices | 1244 |
| 12.6.1. Configure memory and load parameters | 1245 |
| 13.KVStore for Redis | 1252 |
| 13.1. What is KVStore for Redis? | 1252 |

| | |
|--|------|
| 13.2. Quick Start | 1252 |
| 13.2.1. Get started with KVStore for Redis | 1252 |
| 13.2.2. Log on to the KVStore for Redis console | 1253 |
| 13.2.3. Create a KVStore for Redis instance | 1254 |
| 13.2.4. Configure a whitelist | 1255 |
| 13.2.5. Connect to an instance | 1257 |
| 13.2.5.1. Use a Redis client | 1257 |
| 13.2.5.2. Use redis-cli | 1269 |
| 13.3. Instance management | 1270 |
| 13.3.1. Change a password | 1270 |
| 13.3.2. Configure a whitelist | 1270 |
| 13.3.3. Change specifications | 1272 |
| 13.3.4. Specify a maintenance window | 1273 |
| 13.3.5. Upgrade the minor version | 1273 |
| 13.3.6. Configure SSL encryption | 1274 |
| 13.3.7. Clear data | 1274 |
| 13.3.8. Release an instance | 1275 |
| 13.3.9. Manage database accounts | 1275 |
| 13.3.10. Restart an instance | 1276 |
| 13.3.11. Export the list of instances | 1276 |
| 13.3.12. Use a Lua script | 1277 |
| 13.4. Connection management | 1277 |
| 13.4.1. View endpoints | 1277 |
| 13.4.2. Apply for a public endpoint | 1278 |
| 13.4.3. Modify the endpoint of an KVStore for Redis instance | 1278 |
| 13.5. Performance monitoring | 1279 |
| 13.5.1. Query monitoring data | 1279 |
| 13.5.2. Select metrics | 1280 |

| | |
|--|------|
| 13.5.3. Modify the data collection interval | 1281 |
| 13.5.4. Understand metrics | 1281 |
| 13.6. Parameter configuration | 1284 |
| 13.7. Backup and recovery | 1289 |
| 13.7.1. Automatically back up data | 1289 |
| 13.7.2. Back up an instance | 1289 |
| 13.7.3. Download backup files | 1289 |
| 13.7.4. Restore data | 1290 |
| 13.7.5. Clone an instance | 1291 |
| 13.8. CloudDBA | 1291 |
| 13.8.1. Performance trends | 1291 |
| 13.8.2. Add a performance trend chart | 1292 |
| 13.8.3. View performance metrics in real time | 1293 |
| 13.8.4. Instance sessions | 1294 |
| 13.8.5. Slow queries | 1295 |
| 14. ApsaraDB for MongoDB | 1296 |
| 14.1. Usage notes | 1296 |
| 14.2. Log on to the ApsaraDB for MongoDB console | 1296 |
| 14.3. Quick start | 1297 |
| 14.3.1. Use ApsaraDB for MongoDB | 1297 |
| 14.3.2. Create an ApsaraDB for MongoDB instance | 1297 |
| 14.3.3. Reset the password for an ApsaraDB for MongoDB instance | 1301 |
| 14.3.4. Configure a whitelist for an ApsaraDB for MongoDB instance | 1302 |
| 14.3.5. Connect to an instance | 1303 |
| 14.3.5.1. Use DMS to log on to an ApsaraDB for MongoDB instance | 1303 |
| 14.3.5.2. Use the mongo shell to connect to an ApsaraDB for MongoDB instance | 1304 |
| 14.3.5.3. Introduction to connection strings and URIs | 1307 |
| 14.3.5.3.1. Overview of replica set instance connections | 1307 |

| | |
|--|------|
| 14.3.5.3.2. Overview of sharded cluster instance connectio... | 1308 |
| 14.4. Instances | 1310 |
| 14.4.1. Create an ApsaraDB for MongoDB instance | 1310 |
| 14.4.2. View the details of an ApsaraDB for MongoDB instan... | 1313 |
| 14.4.3. Restart an ApsaraDB for MongoDB instance | 1314 |
| 14.4.4. Change the specifications of an ApsaraDB for Mongo... | 1314 |
| 14.4.5. Change the name of an ApsaraDB for MongoDB insta... | 1315 |
| 14.4.6. Reset the password for an ApsaraDB for MongoDB in... | 1316 |
| 14.4.7. Switch node roles | 1316 |
| 14.4.8. Migrate an ApsaraDB for MongoDB instance across zo... | 1319 |
| 14.4.9. Release an ApsaraDB for MongoDB instance | 1321 |
| 14.4.10. Primary/secondary failover | 1322 |
| 14.4.10.1. Trigger a primary/secondary failover for a replica... | 1322 |
| 14.4.10.2. Trigger a primary/secondary failover for a sharde... | 1323 |
| 14.4.11. Monitoring | 1324 |
| 14.5. Backup and restoration | 1326 |
| 14.5.1. Configure automatic backup for an ApsaraDB for Mon... | 1326 |
| 14.5.2. Manually back up an ApsaraDB for MongoDB instance | 1327 |
| 14.5.3. Restore data to the current ApsaraDB for MongoDB i... | 1327 |
| 14.6. Database connections | 1328 |
| 14.6.1. Modify a public or internal endpoint of an ApsaraDB ... | 1328 |
| 14.6.2. Use DMS to log on to an ApsaraDB for MongoDB ins... | 1330 |
| 14.6.3. Use the mongo shell to connect to an ApsaraDB for ... | 1331 |
| 14.6.4. Apply for a public endpoint for an ApsaraDB for Mon.. | 1333 |
| 14.6.5. Release a public endpoint | 1335 |
| 14.6.6. Overview of replica set instance connections | 1337 |
| 14.6.7. Overview of sharded cluster instance connections | 1338 |
| 14.7. Data security | 1340 |

| | |
|---|------|
| 14.7.1. Configure a whitelist for an ApsaraDB for MongoDB in... | 1340 |
| 14.7.2. Create or delete a whitelist | 1341 |
| 14.7.3. Audit logs | 1342 |
| 14.7.4. Configure SSL encryption for an ApsaraDB for Mongo... | 1343 |
| 14.7.5. Configure TDE for an ApsaraDB for MongoDB instance | 1344 |
| 14.7.6. Use the mongo shell to connect to an ApsaraDB for ... | 1346 |
| 14.8. Zone-disaster recovery | 1347 |
| 14.8.1. Create a dual-zone replica set instance | 1347 |
| 14.8.2. Create a dual-zone sharded cluster instance | 1347 |
| 14.9. CloudDBA | 1348 |
| 14.9.1. Performance trends | 1348 |
| 14.9.2. Real-time performance | 1348 |
| 14.9.3. Instance sessions | 1349 |
| 14.9.4. Storage analysis | 1351 |
| 14.9.5. Slow query logs | 1353 |
| 15. Data Management (DMS) | 1355 |
| 15.1. What is DMS? | 1355 |
| 15.2. Quick start | 1355 |
| 15.2.1. Log on to the DMS console | 1355 |
| 15.2.2. Customize the top navigation bar | 1356 |
| 15.2.3. Add an instance | 1356 |
| 15.2.4. Add a user | 1358 |
| 15.2.5. Use the sharing feature | 1359 |
| 15.3. Control modes | 1361 |
| 15.4. Features that are supported by each role | 1362 |
| 15.5. Apply for permissions | 1365 |
| 15.6. SQLConsole | 1369 |
| 15.6.1. Single database query | 1369 |

| | |
|--|------|
| 15.6.2. Cmd Tab | 1371 |
| 15.6.3. Super SQL mode | 1372 |
| 15.6.4. Cross-database query | 1372 |
| 15.7. Data plans | 1374 |
| 15.7.1. Change data | 1374 |
| 15.7.2. Import data | 1376 |
| 15.7.3. Data export | 1378 |
| 15.7.4. Generate test data | 1380 |
| 15.7.5. Clone databases | 1382 |
| 15.8. Data factory | 1384 |
| 15.8.1. Task orchestration | 1384 |
| 15.8.2. Data warehouse development | 1389 |
| 15.8.2.1. Overview | 1389 |
| 15.8.2.2. Create a data warehouse project | 1390 |
| 15.8.2.3. Create or import an internal table | 1392 |
| 15.8.2.4. Manage task flows | 1393 |
| 15.8.2.5. Use the data service feature | 1393 |
| 15.8.3. Data service | 1394 |
| 15.8.3.1. Overview | 1394 |
| 15.8.3.2. Develop an API | 1395 |
| 15.8.3.3. Unpublish or test an API | 1399 |
| 15.8.3.4. Test an API | 1399 |
| 15.8.3.5. Call an API | 1400 |
| 15.9. Schemas | 1401 |
| 15.9.1. Schema design | 1401 |
| 15.9.2. Schema synchronization | 1404 |
| 15.9.3. Synchronize shadow tables | 1406 |
| 15.9.4. Initialize empty databases | 1407 |

| | |
|---|------|
| 15.9.5. Repair table consistency | 1409 |
| 15.10. SQL review | 1410 |
| 15.11. System management | 1412 |
| 15.11.1. Manage instances | 1412 |
| 15.11.2. Database management | 1413 |
| 15.11.3. Manage users | 1414 |
| 15.11.4. Enable metadata access control | 1415 |
| 15.11.5. Manage tasks | 1417 |
| 15.11.6. Configuration management | 1417 |
| 15.11.7. Database grouping | 1418 |
| 15.11.8. Security management | 1420 |
| 15.11.8.1. Manage security rules | 1420 |
| 15.11.8.2. DSL syntax for security rules | 1421 |
| 15.11.8.3. Configure security rules for a database instance | 1426 |
| 15.11.8.4. Customize approval processes | 1426 |
| 15.11.8.5. Operation audit | 1429 |
| 15.11.8.6. Configure IP whitelists | 1431 |
| 15.11.8.7. Row-level control | 1432 |
| 15.11.8.8. Manage sensitive data | 1434 |
| 15.11.8.9. Data protection | 1437 |
| 15.11.9. Security rules | 1439 |
| 15.11.9.1. Overview of security rule sets | 1439 |
| 15.11.9.2. Manage security rules under checkpoints | 1439 |
| 15.11.9.3. SQLConsole for relational databases | 1440 |
| 15.11.9.4. SQLConsole for MongoDB | 1445 |
| 15.11.9.5. SQLConsole for Redis | 1449 |
| 15.11.9.6. Data change | 1453 |
| 15.11.9.7. Permission application | 1457 |

| | |
|--|------|
| 15.11.9.8. Data export | 1459 |
| 15.11.9.9. Schema design | 1460 |
| 15.11.9.10. Database and table synchronization | 1464 |
| 15.11.9.11. Sensitive field change | 1466 |
| 15.11.9.12. Test data generation | 1467 |
| 15.11.9.13. Database cloning | 1467 |
| 16. Server Load Balancer (SLB) | 1469 |
| 16.1. What is SLB? | 1469 |
| 16.2. Log on to the SLB console | 1470 |
| 16.3. Quick start | 1470 |
| 16.3.1. Overview | 1470 |
| 16.3.2. Make preparations | 1471 |
| 16.3.3. Create an SLB instance | 1473 |
| 16.3.4. Configure an SLB instance | 1474 |
| 16.3.5. Release an SLB instance | 1476 |
| 16.4. SLB instances | 1476 |
| 16.4.1. SLB instance overview | 1476 |
| 16.4.2. Create an SLB instance | 1479 |
| 16.4.3. Start and stop an SLB instance | 1480 |
| 16.4.4. Tags | 1480 |
| 16.4.4.1. Tag overview | 1480 |
| 16.4.4.2. Add tags | 1481 |
| 16.4.4.3. Query SLB instances by tag | 1482 |
| 16.4.4.4. Remove a tag | 1482 |
| 16.4.5. Release an SLB instance | 1483 |
| 16.5. Listeners | 1484 |
| 16.5.1. Listener overview | 1484 |
| 16.5.2. Add a TCP listener | 1484 |

| | |
|---|------|
| 16.5.3. Add a UDP listener | 1487 |
| 16.5.4. Add an HTTP listener | 1489 |
| 16.5.5. Add an HTTPS listener | 1492 |
| 16.5.6. Configure forwarding rules | 1495 |
| 16.5.7. Enable access control | 1497 |
| 16.5.8. Disable access control | 1497 |
| 16.6. Backend servers | 1498 |
| 16.6.1. Backend server overview | 1498 |
| 16.6.2. Default server groups | 1499 |
| 16.6.2.1. Add a default backend server | 1499 |
| 16.6.2.2. Add IDC servers to the default server group | 1500 |
| 16.6.2.3. Change the weight of a backend server | 1501 |
| 16.6.2.4. Remove a backend server | 1502 |
| 16.6.3. VServer groups | 1502 |
| 16.6.3.1. Add ECS instances to a VServer group | 1502 |
| 16.6.3.2. Add IDC servers to a VServer group | 1503 |
| 16.6.3.3. Modify a VServer group | 1504 |
| 16.6.3.4. Delete a VServer group | 1505 |
| 16.6.4. Active/standby server groups | 1505 |
| 16.6.4.1. Add ECS instances to a primary/secondary server g...----- | 1505 |
| 16.6.4.2. Add IDC servers to a primary/secondary server gro...----- | 1508 |
| 16.6.4.3. Delete a primary/secondary server group | 1510 |
| 16.7. Health check | 1510 |
| 16.7.1. Health check overview | 1510 |
| 16.7.2. Configure health checks | 1518 |
| 16.7.3. Disable the health check feature | 1520 |
| 16.8. Certificate management | 1520 |
| 16.8.1. Certificate overview | 1520 |

| | |
|--|------|
| 16.8.2. Certificate requirements | 1520 |
| 16.8.3. Upload certificates | 1522 |
| 16.8.4. Generate a CA certificate | 1523 |
| 16.8.5. Convert the certificate format | 1526 |
| 16.8.6. Replace a certificate | 1526 |
| 17.Virtual Private Cloud (VPC) | 1527 |
| 17.1. What is a VPC? | 1527 |
| 17.2. Log on to the VPC console | 1528 |
| 17.3. Quick start | 1528 |
| 17.3.1. Plan and design a VPC | 1528 |
| 17.3.2. Create an IPv4 VPC | 1531 |
| 17.3.3. Create an IPv6 VPC | 1535 |
| 17.4. VPCs and VSwitches | 1540 |
| 17.4.1. Overview | 1540 |
| 17.4.2. VPC management | 1542 |
| 17.4.2.1. Create a VPC | 1542 |
| 17.4.2.2. Add a secondary IPv4 CIDR block | 1543 |
| 17.4.2.3. Delete a secondary IPv4 CIDR block | 1545 |
| 17.4.2.4. Modify the name and description of a VPC | 1545 |
| 17.4.2.5. Delete a VPC | 1545 |
| 17.4.3. VSwitch management | 1546 |
| 17.4.3.1. Create a vSwitch | 1546 |
| 17.4.3.2. Create cloud resources in a vSwitch | 1547 |
| 17.4.3.3. Modify a vSwitch | 1548 |
| 17.4.3.4. Delete a vSwitch | 1548 |
| 17.5. Route tables | 1548 |
| 17.5.1. Overview | 1548 |
| 17.5.2. Add a custom route entry | 1553 |

| | |
|--|------|
| 17.5.3. Export route entries | 1555 |
| 17.5.4. Modify a route table | 1556 |
| 17.5.5. Delete a custom route entry | 1556 |
| 17.6. HAVIPs | 1556 |
| 17.6.1. Overview | 1556 |
| 17.6.2. Create HAVIPs | 1560 |
| 17.6.3. Associate HAVIPs with backend cloud resources | 1560 |
| 17.6.3.1. Associate an HAVIP with an ECS instance | 1561 |
| 17.6.3.2. Associate an HAVIP with an ENI | 1562 |
| 17.6.4. Associate HAVIPs with EIPs | 1562 |
| 17.6.5. Disassociate HAVIPs from backend cloud resources | 1563 |
| 17.6.5.1. Disassociate an HAVIP from an ECS instance | 1563 |
| 17.6.5.2. Disassociate an HAVIP from an ENI | 1563 |
| 17.6.6. Disassociate an HAVIP from an EIP | 1563 |
| 17.6.7. Delete an HAVIP | 1564 |
| 17.7. Network ACLs | 1564 |
| 17.7.1. Overview | 1564 |
| 17.7.2. Scenarios | 1567 |
| 17.7.3. Create a network ACL | 1570 |
| 17.7.4. Associate a network ACL with a vSwitch | 1571 |
| 17.7.5. Add network ACL rules | 1571 |
| 17.7.5.1. Add an inbound rule | 1571 |
| 17.7.5.2. Add an outbound rule | 1572 |
| 17.7.5.3. Change the priority of a network ACL rule | 1573 |
| 17.7.6. Disassociate a network ACL from a vSwitch | 1574 |
| 17.7.7. Delete a network ACL | 1574 |
| 18. NAT Gateway | 1575 |
| 18.1. What is NAT Gateway? | 1575 |

| | |
|--|------|
| 18.2. Log on to the NAT Gateway console | 1575 |
| 18.3. Quick Start | 1576 |
| 18.3.1. Overview | 1576 |
| 18.3.2. Create a NAT gateway | 1577 |
| 18.3.3. Associate an EIP with a NAT gateway | 1578 |
| 18.3.4. Create a DNAT entry | 1578 |
| 18.3.5. Create an SNAT entry | 1579 |
| 18.4. Manage a NAT gateway | 1580 |
| 18.4.1. Overview | 1581 |
| 18.4.2. Create a NAT gateway | 1581 |
| 18.4.3. Modify a NAT gateway | 1582 |
| 18.4.4. Delete a NAT gateway | 1583 |
| 18.5. Manage EIPs | 1583 |
| 18.5.1. Associate an EIP with a NAT gateway | 1583 |
| 18.5.2. Disassociate an EIP from a NAT gateway | 1584 |
| 18.6. Manage a DNAT table | 1584 |
| 18.6.1. DNAT overview | 1584 |
| 18.6.2. Create a DNAT entry | 1585 |
| 18.6.3. Modify a DNAT entry | 1586 |
| 18.6.4. Delete a DNAT entry | 1586 |
| 18.7. Manage an SNAT table | 1587 |
| 18.7.1. SNAT table overview | 1587 |
| 18.7.2. Create an SNAT entry | 1587 |
| 18.7.3. Modify an SNAT entry | 1589 |
| 18.7.4. Delete a SNAT entry | 1589 |
| 18.8. NAT service plan | 1589 |
| 18.8.1. Create a NAT service plan | 1589 |
| 18.8.2. Modify the bandwidth of a NAT service plan | 1590 |

| | |
|---|------|
| 18.8.3. Add an IP address | 1590 |
| 18.8.4. Release an IP address | 1591 |
| 18.8.5. Delete a NAT service plan | 1591 |
| 18.9. Anti-DDoS Origin Basic | 1591 |
| 19.VPN Gateway | 1593 |
| 19.1. What is VPN Gateway? | 1593 |
| 19.2. Log on to the VPN Gateway console | 1593 |
| 19.3. Get started with IPsec-VPN | 1594 |
| 19.3.1. Connect a data center to a VPC | 1594 |
| 19.4. Get started with SSL-VPN | 1597 |
| 19.4.1. Connect a Linux client to a VPC | 1597 |
| 19.4.2. Connect a Windows client to a VPC | 1600 |
| 19.4.3. Connect a macOS client to a VPC | 1602 |
| 19.5. Manage a VPN Gateway | 1605 |
| 19.5.1. Create a VPN gateway | 1605 |
| 19.5.2. Modify a VPN gateway | 1606 |
| 19.5.3. Configure routes of a VPN Gateway | 1607 |
| 19.5.3.1. Route overview | 1607 |
| 19.5.3.2. Work with a policy-based route | 1607 |
| 19.5.3.3. Manage destination-based routes | 1609 |
| 19.5.4. Delete a VPN gateway | 1610 |
| 19.5.5. View the monitoring information about a VPN gatew... .. | 1611 |
| 19.6. Manage a customer gateway | 1611 |
| 19.6.1. Create a customer gateway | 1611 |
| 19.6.2. Modify a customer gateway | 1612 |
| 19.6.3. Delete a customer gateway | 1613 |
| 19.7. Configure IPsec-VPN connections | 1613 |
| 19.7.1. Configuration overview | 1613 |

| | |
|--|------|
| 19.7.2. Manage an IPsec-VPN connection | 1614 |
| 19.7.2.1. Create an IPsec-VPN connection | 1614 |
| 19.7.2.2. Modify an IPsec-VPN connection | 1616 |
| 19.7.2.3. Download the configuration file of an IPsec-VPN c... | 1617 |
| 19.7.2.4. Configure a security group | 1617 |
| 19.7.2.5. View IPsec-VPN connection logs | 1618 |
| 19.7.2.6. Delete an IPsec-VPN connection | 1619 |
| 19.7.3. View the monitoring information about an IPsec-VPN ... | 1619 |
| 19.7.4. MTU considerations | 1620 |
| 19.8. Configure SSL-VPN | 1620 |
| 19.8.1. SSL-VPN configuration overview | 1620 |
| 19.8.2. Manage an SSL server | 1621 |
| 19.8.2.1. Create an SSL server | 1621 |
| 19.8.2.2. Modify an SSL server | 1622 |
| 19.8.2.3. Configure a security group | 1622 |
| 19.8.2.4. Delete an SSL server | 1624 |
| 19.8.3. Manage an SSL client certificate | 1624 |
| 19.8.3.1. Create an SSL client certificate | 1624 |
| 19.8.3.2. Download an SSL client certificate | 1625 |
| 19.8.3.3. Delete an SSL client certificate | 1625 |
| 19.8.4. Query SSL-VPN connection logs | 1625 |
| 20. Elastic IP Address | 1627 |
| 20.1. What is Elastic IP Address? | 1627 |
| 20.2. Log on to the EIP console | 1627 |
| 20.3. Quick start | 1628 |
| 20.3.1. Tutorial overview | 1628 |
| 20.3.2. Apply for EIPs | 1628 |
| 20.3.3. Associate an EIP with an ECS instance | 1629 |

| | |
|---|------|
| 20.3.4. Disassociate an EIP from a cloud resource | 1630 |
| 20.3.5. Release an EIP | 1630 |
| 20.4. Manage EIPs | 1631 |
| 20.4.1. Apply for EIPs | 1631 |
| 20.4.2. Bind an EIP to a cloud instance | 1631 |
| 20.4.2.1. Associate an EIP with an ECS instance | 1631 |
| 20.4.2.2. Associate an EIP with an SLB instance | 1632 |
| 20.4.2.3. Associate an EIP with a NAT gateway | 1633 |
| 20.4.2.4. Bind an EIP to a secondary ENI | 1633 |
| 20.4.2.4.1. Overview | 1633 |
| 20.4.2.4.2. Associate an EIP with a secondary ENI in NAT... | 1635 |
| 20.4.3. Resize the maximum bandwidth | 1636 |
| 20.4.4. Disassociate an EIP from a cloud resource | 1636 |
| 20.4.5. Release an EIP | 1636 |
| 21. Apsara Stack Security | 1637 |
| 21.1. What is Apsara Stack Security | 1637 |
| 21.2. Usage notes | 1637 |
| 21.3. Quick start | 1638 |
| 21.3.1. User roles and permissions | 1638 |
| 21.3.2. Log on to Apsara Stack Security Center | 1639 |
| 21.4. Threat Detection Service | 1639 |
| 21.4.1. Overview | 1639 |
| 21.4.2. Security overview | 1640 |
| 21.4.2.1. View security overview information | 1640 |
| 21.4.3. Security alerts | 1641 |
| 21.4.3.1. View security alerts | 1641 |
| 21.4.3.2. Manage quarantined files | 1641 |
| 21.4.3.3. Configure security alerts | 1642 |

| | |
|---|------|
| 21.4.4. Attack analysis | 1645 |
| 21.4.5. Cloud service check | 1646 |
| 21.4.5.1. Overview | 1646 |
| 21.4.5.2. Run cloud service checks | 1648 |
| 21.4.5.3. View the check results of configuration assessment... | 1649 |
| 21.4.6. Application whitelist | 1651 |
| 21.4.7. Assets | 1654 |
| 21.4.7.1. View the security status of a server | 1655 |
| 21.4.7.2. View the security status of cloud services | 1657 |
| 21.4.7.3. View the details of a single asset | 1659 |
| 21.4.7.4. Enable and disable server protection | 1662 |
| 21.4.7.5. Perform a quick security check | 1663 |
| 21.4.7.6. Manage server groups | 1663 |
| 21.4.7.7. Manage asset tags | 1665 |
| 21.4.8. Vulnerability scan | 1668 |
| 21.4.8.1. Quick start | 1668 |
| 21.4.8.2. View the information on the Overview page | 1669 |
| 21.4.8.3. Asset management | 1670 |
| 21.4.8.3.1. View the results of asset analysis | 1670 |
| 21.4.8.3.2. Import assets | 1671 |
| 21.4.8.3.3. Manage assets | 1673 |
| 21.4.8.3.4. Manage asset availability | 1675 |
| 21.4.8.3.5. Manage custom update detection tasks | 1678 |
| 21.4.8.4. Risk management | 1680 |
| 21.4.8.4.1. Manage vulnerabilities | 1680 |
| 21.4.8.4.2. Manage host compliance risks | 1681 |
| 21.4.8.4.3. Manage external risks | 1682 |
| 21.4.8.4.4. Create a custom risk detection task | 1682 |

| | |
|---|------|
| 21.4.8.5. Report management | 1683 |
| 21.4.8.5.1. Create a report | 1683 |
| 21.4.8.5.2. Delete multiple reports at a time | 1684 |
| 21.4.8.6. Configuration management | 1685 |
| 21.4.8.6.1. Configure overall monitoring | 1685 |
| 21.4.8.6.2. Configure basic monitoring | 1688 |
| 21.4.8.6.3. Configure web monitoring | 1690 |
| 21.4.8.6.4. Configure a whitelist | 1692 |
| 21.4.8.6.5. Configure a scan engine for internal assets | 1692 |
| 21.4.9. Create a security report | 1693 |
| 21.5. Server security | 1694 |
| 21.5.1. Server security overview | 1694 |
| 21.5.2. Server fingerprints | 1695 |
| 21.5.2.1. Manage listening ports | 1695 |
| 21.5.2.2. Manage software versions | 1696 |
| 21.5.2.3. Manage processes | 1696 |
| 21.5.2.4. Manage account information | 1696 |
| 21.5.2.5. Manage scheduled tasks | 1697 |
| 21.5.2.6. Set the fingerprint collection frequency | 1697 |
| 21.5.3. Threat protection | 1697 |
| 21.5.3.1. Vulnerability management | 1697 |
| 21.5.3.1.1. Manage Linux vulnerabilities | 1697 |
| 21.5.3.1.2. Manage Windows vulnerabilities | 1698 |
| 21.5.3.1.3. Manage Web CMS vulnerabilities | 1699 |
| 21.5.3.1.4. Manage urgent vulnerabilities | 1700 |
| 21.5.3.1.5. Configure vulnerability management policies | 1701 |
| 21.5.3.2. Baseline check | 1702 |
| 21.5.3.2.1. Baseline check overview | 1702 |

| | |
|--|------|
| 21.5.3.2.2. Configure baseline check policies | 1707 |
| 21.5.3.2.3. View baseline check results and manage failed...----- | 1708 |
| 21.5.4. Intrusion prevention | 1710 |
| 21.5.4.1. Intrusion events | 1711 |
| 21.5.4.1.1. Intrusion event types | 1711 |
| 21.5.4.1.2. View and handle detected intrusion events | 1712 |
| 21.5.4.1.3. View exceptions related to an alert | 1713 |
| 21.5.4.1.4. Use the file quarantine feature | 1714 |
| 21.5.4.1.5. Configure security alerts | 1714 |
| 21.5.4.1.6. Cloud threat detection | 1716 |
| 21.5.4.2. Website tamper-proofing | 1717 |
| 21.5.4.2.1. Overview | 1717 |
| 21.5.4.2.2. Configure tamper protection | 1718 |
| 21.5.4.2.3. View protection status | 1721 |
| 21.5.4.3. Configure the anti-virus feature | 1722 |
| 21.5.5. Log retrieval | 1723 |
| 21.5.5.1. Log retrieval overview | 1723 |
| 21.5.5.2. Query logs | 1723 |
| 21.5.5.3. Supported log sources and fields | 1724 |
| 21.5.5.4. Logical operators | 1728 |
| 21.5.6. Settings | 1728 |
| 21.5.6.1. Install the Server Guard agent | 1729 |
| 21.5.6.2. Manage protection modes | 1729 |
| 21.6. Physical server security | 1730 |
| 21.6.1. Create and grant permissions to a security administrat...----- | 1730 |
| 21.6.2. View the information on the Overview page | 1731 |
| 21.6.3. Physical servers | 1731 |
| 21.6.3.1. Manage physical server groups | 1731 |

| | |
|--|------|
| 21.6.3.2. Manage physical servers | 1732 |
| 21.6.4. Intrusion detection | 1733 |
| 21.6.4.1. Configure policies to identify unusual logons | 1733 |
| 21.6.4.2. Handle unusual logons | 1735 |
| 21.6.5. Server fingerprints | 1735 |
| 21.6.5.1. Configure data refresh frequencies | 1735 |
| 21.6.5.2. View listening ports | 1736 |
| 21.6.5.3. View running processes | 1736 |
| 21.6.5.4. View account information | 1737 |
| 21.6.5.5. View software versions | 1737 |
| 21.6.6. Log retrieval | 1737 |
| 21.6.6.1. Supported log sources and fields | 1738 |
| 21.6.6.2. Logical operators | 1741 |
| 21.6.6.3. Query logs | 1742 |
| 21.6.7. Configure security settings for physical servers | 1743 |
| 21.7. Application security | 1743 |
| 21.7.1. Quick start | 1743 |
| 21.7.2. Detection overview | 1744 |
| 21.7.2.1. View protection overview | 1744 |
| 21.7.2.2. View access information | 1745 |
| 21.7.3. Protection logs | 1746 |
| 21.7.3.1. View attack detection logs | 1746 |
| 21.7.3.2. View HTTP flood protection logs | 1746 |
| 21.7.3.3. View system operation logs | 1747 |
| 21.7.3.4. View access logs | 1747 |
| 21.7.4. Protection configuration | 1747 |
| 21.7.4.1. Configure protection policies | 1747 |
| 21.7.4.2. Create a custom rule | 1749 |

| | |
|---|------|
| 21.7.4.3. Configure an HTTP flood protection rule | 1750 |
| 21.7.4.4. Configure the HTTP flood whitelist | 1753 |
| 21.7.4.5. Manage SSL certificates | 1754 |
| 21.7.4.6. Add Internet websites for protection | 1755 |
| 21.7.4.7. Add VPC websites for protection | 1759 |
| 21.7.4.8. Verify the configurations of a website on your on-... | 1764 |
| 21.7.4.9. Modify DNS resolution settings | 1764 |
| 21.7.5. System management | 1765 |
| 21.7.5.1. View the load status of nodes | 1765 |
| 21.7.5.2. View the network status of nodes | 1766 |
| 21.7.5.3. View the disk status of nodes | 1767 |
| 21.7.5.4. Configure alerts | 1768 |
| 21.7.5.5. Configure alert thresholds | 1769 |
| 21.8. Security Operations Center (SOC) | 1770 |
| 21.8.1. View the dashboard | 1770 |
| 21.8.2. Security Monitoring | 1771 |
| 21.8.2.1. View security monitoring data of tenants | 1771 |
| 21.8.2.2. View security monitoring data of the Apsara Stack... | 1773 |
| 21.8.2.3. View the global traffic | 1775 |
| 21.8.3. Asset Management | 1776 |
| 21.8.3.1. View tenant assets | 1776 |
| 21.8.3.2. View platform assets | 1776 |
| 21.8.4. Log Analysis | 1777 |
| 21.8.4.1. View the Log Overview page | 1777 |
| 21.8.4.2. View global logs | 1777 |
| 21.8.4.3. Log configurations | 1779 |
| 21.8.4.3.1. Manage log sources | 1779 |
| 21.8.4.3.2. Create a log collection task | 1779 |

| | |
|--|------|
| 21.8.4.3.3. Manage log collectors | 1782 |
| 21.8.4.3.4. Manage storage policies | 1783 |
| 21.8.4.4. Security Audit | 1784 |
| 21.8.4.4.1. Overview | 1784 |
| 21.8.4.4.2. View security audit overview | 1784 |
| 21.8.4.4.3. Query audit events | 1785 |
| 21.8.4.4.4. View raw logs | 1786 |
| 21.8.4.4.5. Manage log sources | 1787 |
| 21.8.4.4.6. Policy settings | 1787 |
| 21.8.4.4.6.1. Manage audit rules | 1787 |
| 21.8.4.4.6.2. Configure alert recipients | 1789 |
| 21.8.4.4.6.3. Manage archives of events and logs | 1790 |
| 21.8.4.4.6.4. Manage export tasks | 1790 |
| 21.8.4.4.6.5. Modify system settings | 1791 |
| 21.8.5. Rules | 1791 |
| 21.8.5.1. Create an IPS rule for traffic monitoring | 1791 |
| 21.8.5.2. Manage IPS rules of Cloud Firewall | 1792 |
| 21.8.5.3. Create IDS rules for traffic monitoring | 1793 |
| 21.8.5.4. Manage IDS rules for traffic monitoring | 1794 |
| 21.8.5.5. Specify custom thresholds for DDoS traffic scrubbi... .. | 1795 |
| 21.8.5.6. View Server Guard rules | 1795 |
| 21.8.6. Threat intelligence | 1796 |
| 21.8.6.1. Enable the service configuration feature | 1796 |
| 21.8.6.2. View the Overview page | 1796 |
| 21.8.6.3. Search for and view the information about a susp... .. | 1797 |
| 21.8.7. Create a report task | 1798 |
| 21.8.8. System Configurations | 1799 |
| 21.8.8.1. View and manage metrics | 1799 |

| | |
|---|------|
| 21.8.8.2. Alert settings | 1801 |
| 21.8.8.2.1. Configure alert contacts | 1802 |
| 21.8.8.2.2. Configure alert notifications | 1802 |
| 21.8.8.3. Updates | 1803 |
| 21.8.8.3.1. Overview of the system updates feature | 1803 |
| 21.8.8.3.2. Enable automatic update check and update ru... .. | 1803 |
| 21.8.8.3.3. Manually import an update package and upda... .. | 1804 |
| 21.8.8.3.4. Roll back a rule library | 1804 |
| 21.8.8.3.5. View the update history of a rule library | 1805 |
| 21.8.8.4. Global configuration | 1805 |
| 21.8.8.4.1. Set CIDR blocks for traffic monitoring | 1805 |
| 21.8.8.4.1.1. Add a CIDR block for traffic monitoring | 1805 |
| 21.8.8.4.1.2. Manage CIDR blocks for traffic monitoring | 1806 |
| 21.8.8.4.2. Region settings | 1806 |
| 21.8.8.4.2.1. Add a CIDR block for a region | 1806 |
| 21.8.8.4.2.2. Manage CIDR blocks for a region | 1807 |
| 21.8.8.4.3. Configure whitelists | 1807 |
| 21.8.8.4.4. Configure attack blocking policies | 1808 |
| 21.8.8.4.5. Block IP addresses | 1809 |
| 21.8.8.4.6. Configure custom IP addresses and locations | 1810 |
| 21.8.8.4.6.1. Add custom IP addresses and locations | 1810 |
| 21.8.8.4.6.2. Manage custom IP addresses and locations | 1811 |
| 21.8.8.5. System Monitoring | 1811 |
| 21.8.8.5.1. Configure CIDR blocks for traffic redirection in | 1811 |
| 21.8.8.6. Inspect services | 1811 |
| 21.8.8.7. Remote operations | 1812 |
| 21.8.8.7.1. Enable remote operations | 1812 |
| 21.8.8.8. Account management | 1813 |

| | |
|--|------|
| 21.8.8.8.1. View and modify an Apsara Stack tenant accou.. | 1813 |
| 21.8.8.8.2. Add an Alibaba Cloud account | 1814 |
| 21.9. Optional security products | 1815 |
| 21.9.1. Anti-DDoS settings | 1815 |
| 21.9.1.1. Overview | 1815 |
| 21.9.1.2. View and configure DDoS mitigation policies | 1815 |
| 21.9.1.3. View DDoS events | 1817 |
| 21.9.2. Cloud Firewall | 1818 |
| 21.9.2.1. Policy configuration | 1818 |
| 21.9.2.1.1. Synchronize assets for the Internet firewall | 1818 |
| 21.9.2.1.2. Create a VPC firewall | 1819 |
| 21.9.2.1.3. Create an IDC-VPC firewall | 1820 |
| 21.9.2.2. Access control | 1823 |
| 21.9.2.2.1. Manage address books | 1823 |
| 21.9.2.2.2. Configure access control policies on the Interne.. | 1824 |
| 21.9.2.2.3. Create a policy group | 1826 |
| 21.9.2.2.4. Configure access control policies on an interna... | 1827 |
| 21.9.2.2.5. Configure access control policies on a VPC fire... | 1829 |
| 21.9.2.2.6. Configure access control policies on an IDC-VP... | 1832 |
| 21.9.2.3. Intrusion prevention | 1834 |
| 21.9.2.3.1. Configure intrusion prevention policies | 1834 |
| 21.9.2.3.2. View the traffic blocked by IPS | 1836 |
| 21.9.2.4. View security groups | 1837 |
| 21.9.2.5. Log audit | 1839 |
| 21.9.2.5.1. View event logs | 1839 |
| 21.9.2.5.2. View traffic logs | 1839 |
| 21.9.3. Data Encryption Service | 1840 |
| 21.9.3.1. Data Encryption Service overview | 1840 |

| | |
|---|------|
| 21.9.3.2. Management of Data Encryption Service instances | 1841 |
| 21.9.3.2.1. Create a Data Encryption Service instance | 1841 |
| 21.9.3.2.2. Configure a Data Encryption Service instance | 1841 |
| 21.9.3.2.3. Release a Data Encryption Service instance | 1841 |
| 21.10. Restrictions | 1842 |
| 21.11. Log on to Cloud Security Operations Center | 1842 |
| 21.12. Services | 1843 |
| 21.12.1. Data Encryption Service | 1844 |
| 21.12.1.1. Manage Data Encryption Service instances | 1844 |
| 21.12.1.1.1. Create an instance | 1844 |
| 21.12.1.1.2. Configure a VPC | 1844 |
| 21.12.1.1.3. Manage an instance | 1845 |
| 21.12.1.2. Manage HSMs | 1846 |
| 21.12.1.2.1. Add an HSM | 1846 |
| 21.12.1.2.2. Configure the network information for an HSM | 1848 |
| 21.12.1.2.3. Migrate an HSM | 1848 |
| 21.12.1.2.4. Update an HSM | 1849 |
| 21.12.1.2.5. Manage an HSM | 1850 |
| 21.12.1.3. Manage VSMs | 1851 |
| 21.12.1.3.1. Configure the network information for a VSM | 1851 |
| 21.12.1.3.2. Update a VSM | 1852 |
| 21.12.1.3.3. Export snapshots | 1852 |
| 21.12.1.3.4. Manage a VSM | 1853 |
| 21.12.1.4. Manage manufacturers | 1854 |
| 21.12.1.4.1. Add a manufacturer | 1854 |
| 21.12.1.4.2. Manage a manufacturer | 1854 |
| 21.12.1.5. Manage HSM models | 1855 |
| 21.12.1.5.1. Add an HSM model | 1855 |

| | |
|---|------|
| 21.12.1.5.2. Manage an HSM model | 1856 |
| 21.12.1.6. View the information about snapshots | 1857 |
| 21.12.1.7. Manage update files | 1857 |
| 21.12.1.7.1. Upload an update file | 1857 |
| 21.12.1.7.2. Delete an update file | 1858 |
| 21.12.1.8. Manage tasks | 1858 |
| 21.12.1.8.1. View task details | 1858 |
| 21.12.1.8.2. Terminate a task | 1859 |
| 21.12.1.9. Query the configurations of an HSM | 1859 |
| 22.Key Management Service (KMS) | 1860 |
| 22.1. Manage keys in the KMS console | 1860 |
| 22.1.1. Log on to the KMS console | 1860 |
| 22.1.2. Create a CMK | 1860 |
| 22.1.3. View the details of a CMK | 1861 |
| 22.1.4. Enable a CMK | 1861 |
| 22.1.5. Disable a CMK | 1861 |
| 22.1.6. Schedule the deletion of a CMK | 1862 |
| 22.1.7. Configure the rotation policy of a CMK | 1862 |
| 22.1.8. Generate a data key | 1863 |
| 22.1.9. Generate a CSR | 1864 |
| 22.2. Use RAM for access control | 1865 |
| 22.3. Use aliases | 1868 |
| 22.4. Use CMKs | 1872 |
| 22.5. Use symmetric keys | 1873 |
| 22.5.1. Overview | 1873 |
| 22.5.2. EncryptionContext | 1874 |
| 22.5.3. Import and delete key material | 1875 |
| 22.6. Use asymmetric keys | 1879 |

| | |
|--|------|
| 22.6.1. Overview | 1879 |
| 22.6.2. Encrypt and decrypt data by using an asymmetric CMK | 1881 |
| 22.6.3. Generate and verify a digital signature by using an asymmetric CMK | 1883 |
| 22.7. Use managed HSMs | 1886 |
| 22.7.1. Overview | 1886 |
| 22.7.2. Use managed HSMs to create and use keys | 1887 |
| 22.8. Key rotation | 1888 |
| 22.8.1. Overview | 1888 |
| 22.8.2. Automatic key rotation | 1889 |
| 22.8.3. Manual CMK rotation | 1891 |
| 23. Log Service | 1894 |
| 23.1. What is Log Service? | 1894 |
| 23.2. Quick start | 1894 |
| 23.2.1. Procedure | 1894 |
| 23.2.2. Log on to the Log Service console | 1895 |
| 23.2.3. Obtain an AccessKey pair | 1896 |
| 23.2.4. Manage projects | 1897 |
| 23.2.5. Manage Logstores | 1900 |
| 23.2.6. Manage shards | 1903 |
| 23.3. Data collection | 1906 |
| 23.3.1. Collection by Logtail | 1906 |
| 23.3.1.1. Overview | 1906 |
| 23.3.1.1.1. Logtail overview | 1906 |
| 23.3.1.1.2. Log collection process of Logtail | 1909 |
| 23.3.1.1.3. Logtail configuration files and record files | 1911 |
| 23.3.1.2. Installation | 1918 |
| 23.3.1.2.1. Install Logtail in Linux | 1918 |
| 23.3.1.2.2. Install Logtail in Windows | 1920 |

| | |
|--|------|
| 23.3.1.2.3. Set Logtail startup parameters | 1922 |
| 23.3.1.3. Logtail machine group | 1925 |
| 23.3.1.3.1. Overview | 1925 |
| 23.3.1.3.2. Create a machine group based on a server IP | 1926 |
| 23.3.1.3.3. Create a machine group based on a custom ID | 1927 |
| 23.3.1.3.4. View server groups | 1930 |
| 23.3.1.3.5. Modify a server group | 1931 |
| 23.3.1.3.6. View the status of a server group | 1931 |
| 23.3.1.3.7. Delete a server group | 1931 |
| 23.3.1.3.8. Manage server group configurations | 1932 |
| 23.3.1.3.9. Manage a Logtail configuration | 1933 |
| 23.3.1.3.10. Configure an account ID on a server | 1934 |
| 23.3.1.4. Text logs | 1935 |
| 23.3.1.4.1. Configure text log collection | 1935 |
| 23.3.1.4.2. Collect logs by line | 1939 |
| 23.3.1.4.3. Use regular expressions to collect logs | 1942 |
| 23.3.1.4.4. Collect DSV formatted logs | 1947 |
| 23.3.1.4.5. Collect JSON logs | 1953 |
| 23.3.1.4.6. Collect NGINX logs | 1957 |
| 23.3.1.4.7. Collect IIS logs | 1962 |
| 23.3.1.4.8. Collect Apache logs | 1967 |
| 23.3.1.4.9. Configure parsing scripts | 1973 |
| 23.3.1.4.10. Configure the time format | 1974 |
| 23.3.1.4.11. Import historical logs | 1977 |
| 23.3.1.4.12. Generate a topic | 1979 |
| 23.3.1.5. Custom plug-ins | 1981 |
| 23.3.1.5.1. Collect MySQL binary logs | 1981 |
| 23.3.1.5.2. Collect MySQL query results | 1991 |

| | |
|---|------|
| 23.3.1.5.3. Collect syslogs | 1996 |
| 23.3.1.5.4. Configure data processing methods | 2000 |
| 23.3.1.6. Collect container logs | 2020 |
| 23.3.1.6.1. Collect standard Docker logs | 2020 |
| 23.3.1.6.2. Collect Kubernetes logs | 2023 |
| 23.3.1.6.3. Collect container text logs | 2027 |
| 23.3.1.6.4. Collect stdout and stderr logs from containers | 2032 |
| 23.3.1.7. Limits | 2041 |
| 23.3.2. Other collection methods | 2044 |
| 23.3.2.1. WebTracking | 2044 |
| 23.3.2.2. Use SDKs to collect logs | 2048 |
| 23.3.2.2.1. Producer Library | 2048 |
| 23.3.2.2.2. Log4j Appender | 2048 |
| 23.3.2.2.3. Logback Appender | 2048 |
| 23.3.2.2.4. Golang Producer Library | 2049 |
| 23.3.2.2.5. Python logging | 2049 |
| 23.3.2.3. Common log formats | 2052 |
| 23.3.2.3.1. Log4j logs | 2052 |
| 23.3.2.3.2. Python logs | 2053 |
| 23.3.2.3.3. Node.js logs | 2059 |
| 23.3.2.3.4. WordPress logs | 2060 |
| 23.3.2.3.5. Unity3D logs | 2061 |
| 23.4. Query and analysis | 2063 |
| 23.4.1. Overview | 2063 |
| 23.4.2. Real-time analysis | 2064 |
| 23.4.3. Enable the indexing feature and configure indexes fo... | 2066 |
| 23.4.4. Query logs | 2070 |
| 23.4.5. Export logs | 2073 |

| | |
|--|------|
| 23.4.6. Index data type | 2073 |
| 23.4.6.1. Overview | 2073 |
| 23.4.6.2. Query text data | 2075 |
| 23.4.6.3. Numeric type | 2076 |
| 23.4.6.4. JSON indexes | 2077 |
| 23.4.7. Query syntax and functions | 2080 |
| 23.4.7.1. Search syntax | 2080 |
| 23.4.7.2. LiveTail | 2084 |
| 23.4.7.3. LogReduce | 2088 |
| 23.4.7.4. Contextual query | 2092 |
| 23.4.7.5. Saved search | 2094 |
| 23.4.7.6. Quick analysis | 2095 |
| 23.4.8. Analysis grammar | 2098 |
| 23.4.8.1. General aggregate functions | 2098 |
| 23.4.8.2. Security check functions | 2099 |
| 23.4.8.3. Map functions | 2102 |
| 23.4.8.4. Approximate functions | 2103 |
| 23.4.8.5. Mathematical statistics functions | 2104 |
| 23.4.8.6. Mathematical calculation functions | 2105 |
| 23.4.8.7. String functions | 2106 |
| 23.4.8.8. Date and time functions | 2108 |
| 23.4.8.9. URL functions | 2112 |
| 23.4.8.10. Regular expression functions | 2114 |
| 23.4.8.11. JSON functions | 2115 |
| 23.4.8.12. Type conversion functions | 2116 |
| 23.4.8.13. IP functions | 2116 |
| 23.4.8.14. GROUP BY syntax | 2119 |
| 23.4.8.15. Window functions | 2120 |

| | |
|---|------|
| 23.4.8.16. HAVING syntax | 2122 |
| 23.4.8.17. ORDER BY syntax | 2122 |
| 23.4.8.18. LIMIT syntax | 2122 |
| 23.4.8.19. Syntax for CASE statements and if() functions | 2123 |
| 23.4.8.20. Nested subqueries | 2124 |
| 23.4.8.21. Array functions | 2125 |
| 23.4.8.22. Binary string functions | 2127 |
| 23.4.8.23. Bitwise operations | 2127 |
| 23.4.8.24. Interval-valued comparison and periodicity-valued... | 2128 |
| 23.4.8.25. Comparison functions and operators | 2131 |
| 23.4.8.26. Lambda functions | 2132 |
| 23.4.8.27. Logical functions | 2134 |
| 23.4.8.28. Field aliases | 2135 |
| 23.4.8.29. JOIN operations between Logstores and Relationa.. | 2136 |
| 23.4.8.30. Geospatial functions | 2138 |
| 23.4.8.31. Geography functions | 2141 |
| 23.4.8.32. JOIN syntax | 2142 |
| 23.4.8.33. UNNEST function | 2143 |
| 23.4.9. Machine learning syntax and functions | 2144 |
| 23.4.9.1. Overview | 2144 |
| 23.4.9.2. Smooth functions | 2146 |
| 23.4.9.3. Multi-period estimation functions | 2150 |
| 23.4.9.4. Change point detection functions | 2152 |
| 23.4.9.5. Maximum value detection function | 2155 |
| 23.4.9.6. Prediction and anomaly detection functions | 2156 |
| 23.4.9.7. Time series decomposition function | 2162 |
| 23.4.9.8. Time series clustering functions | 2163 |
| 23.4.9.9. Frequent pattern statistics function | 2167 |

| | |
|--|------|
| 23.4.9.10. Differential pattern statistics function | 2168 |
| 23.4.9.11. Root cause analysis function | 2169 |
| 23.4.9.12. Correlation analysis functions | 2172 |
| 23.4.9.13. Kernel density estimation function | 2175 |
| 23.4.10. Advanced analysis | 2176 |
| 23.4.10.1. Optimize queries | 2176 |
| 23.4.10.2. Use cases | 2178 |
| 23.4.10.3. Time field conversion examples | 2178 |
| 23.4.11. Visual analysis | 2179 |
| 23.4.11.1. Analysis graph | 2179 |
| 23.4.11.1.1. Overview | 2179 |
| 23.4.11.1.2. Display query results on a table | 2180 |
| 23.4.11.1.3. Display query results on a line chart | 2181 |
| 23.4.11.1.4. Display query results on a column chart | 2183 |
| 23.4.11.1.5. Display query results on a bar chart | 2184 |
| 23.4.11.1.6. Display query results on a pie chart | 2186 |
| 23.4.11.1.7. Display query results on an area chart | 2189 |
| 23.4.11.1.8. Display query results on a single value chart | 2190 |
| 23.4.11.1.9. Display query results on a progress bar | 2195 |
| 23.4.11.1.10. Display query results on a map | 2197 |
| 23.4.11.1.11. Flow chart | 2200 |
| 23.4.11.1.12. Display query results in a Sankey diagram | 2202 |
| 23.4.11.1.13. Display query results on a word cloud | 2203 |
| 23.4.11.1.14. Display query results on a treemap chart | 2204 |
| 23.4.11.2. Dashboard | 2205 |
| 23.4.11.2.1. Overview | 2205 |
| 23.4.11.2.2. Create and delete a dashboard | 2206 |
| 23.4.11.2.3. Configure the display mode of a dashboard | 2209 |

| | |
|---|------|
| 23.4.11.2.4. Edit mode | 2211 |
| 23.4.11.2.5. Drill-down analysis | 2213 |
| 23.4.11.2.6. Configure and use a filter on a dashboard of... .. | 2219 |
| 23.4.11.2.7. Markdown chart | 2223 |
| 23.5. Alerts | 2225 |
| 23.5.1. Overview | 2225 |
| 23.5.2. Configure an alarm | 2227 |
| 23.5.2.1. Configure alerts | 2227 |
| 23.5.2.2. Grant permissions on alerts to a RAM user | 2229 |
| 23.5.2.3. Configure alert notification methods | 2230 |
| 23.5.3. Modify and view an alarm | 2234 |
| 23.5.3.1. Modify an alert | 2234 |
| 23.5.3.2. View history alerts | 2235 |
| 23.5.3.3. Manage an alert | 2236 |
| 23.5.4. Relevant syntax and fields for reference | 2237 |
| 23.5.4.1. Conditional expression syntax of an alert | 2237 |
| 23.5.4.2. Fields in alert log entries | 2241 |
| 23.6. Real-time consumption | 2243 |
| 23.6.1. Overview | 2243 |
| 23.6.2. Consume log data | 2244 |
| 23.6.3. Consumption by consumer groups | 2245 |
| 23.6.3.1. Use consumer groups to consume log data | 2245 |
| 23.6.3.2. View the status of a consumer group | 2251 |
| 23.6.4. Use LogHub Storm to consume log data | 2253 |
| 23.6.5. Use Flume to consume log data | 2257 |
| 23.6.6. Use open source Flink to consume log data | 2260 |
| 23.6.7. Use Logstash to consume log data | 2265 |
| 23.6.8. Use Spark Streaming to consume log data | 2266 |

| | |
|---|------|
| 23.6.9. Use Realtime Compute to consume log data | 2270 |
| 23.7. Data shipping | 2273 |
| 23.7.1. Ship logs to OSS | 2273 |
| 23.7.1.1. Overview | 2274 |
| 23.7.1.2. Ship log data from Log Service to OSS | 2274 |
| 23.7.1.3. Obtain the ARN of a RAM role | 2278 |
| 23.7.1.4. Storage Formats | 2279 |
| 23.7.1.5. Decompress Snappy compressed files | 2281 |
| 23.8. RAM | 2283 |
| 23.8.1. Overview | 2283 |
| 23.8.2. Create a RAM role | 2283 |
| 23.8.3. Create a user | 2284 |
| 23.8.4. Create a RAM user group | 2284 |
| 23.8.5. Add a RAM user to a RAM user group | 2285 |
| 23.8.6. Create a permission policy | 2286 |
| 23.8.7. Grant a RAM user the permissions to manage a proje... | 2286 |
| 23.8.8. Grant permissions to a RAM role | 2287 |
| 23.8.9. Use custom policies to grant RAM user the required p... | 2287 |
| 23.9. FAQ | 2292 |
| 23.9.1. Log collection | 2292 |
| 23.9.1.1. How do I troubleshoot Logtail collection errors? | 2292 |
| 23.9.1.2. What can I do if Log Service does not receive hea... | 2293 |
| 23.9.1.3. How do I query the local log collection statuses? | 2295 |
| 23.9.1.4. How do I test a regular expression? | 2307 |
| 23.9.1.5. How do I optimize regular expressions? | 2309 |
| 23.9.1.6. How do I use the full regex mode to collect log e... | 2309 |
| 23.9.1.7. How do I set the time format for logs? | 2310 |
| 23.9.1.8. How do I configure non-printable characters in a ... | 2310 |

| | |
|--|------|
| 23.9.1.9. How do I troubleshoot errors during container log... | 2311 |
| 23.9.2. Log search and analysis | 2314 |
| 23.9.2.1. FAQ about log query | 2314 |
| 23.9.2.2. What can I do if no log data is retrieved? | 2315 |
| 23.9.2.3. What are the differences between log consumptio... | 2316 |
| 23.9.2.4. How do I resolve common errors returned in log d.. | 2317 |
| 23.9.2.5. Why data queries are inaccurate? | 2319 |
| 23.9.3. Alarm | 2319 |
| 23.9.3.1. FAQ about alerts | 2319 |
| 24.Apsara Stack DNS | 2321 |
| 24.1. What is Apsara Stack DNS? | 2321 |
| 24.2. User roles and permissions | 2321 |
| 24.3. Log on to the Apsara Stack DNS console | 2322 |
| 24.4. Internal DNS resolution management | 2322 |
| 24.4.1. Global internal domain names | 2322 |
| 24.4.1.1. Overview | 2322 |
| 24.4.1.2. View an internal domain name | 2322 |
| 24.4.1.3. Add a domain name | 2322 |
| 24.4.1.4. Add a description for a domain name | 2323 |
| 24.4.1.5. Delete a domain name | 2323 |
| 24.4.1.6. Delete multiple domain names | 2323 |
| 24.4.1.7. Configure DNS records | 2324 |
| 24.4.1.8. View a resolution policy | 2324 |
| 24.4.2. Global forwarding configurations | 2324 |
| 24.4.2.1. Global forwarding domain names | 2324 |
| 24.4.2.1.1. Overview | 2325 |
| 24.4.2.1.2. View global forwarding domain names | 2325 |
| 24.4.2.1.3. Add a domain name | 2325 |

| | |
|---|------|
| 24.4.2.1.4. Add a description for a domain name | 2325 |
| 24.4.2.1.5. Modify the forwarding configurations of a dom...----- | 2326 |
| 24.4.2.1.6. Delete a domain name | 2326 |
| 24.4.2.1.7. Delete multiple domain names | 2326 |
| 24.4.2.2. Global default forwarding configurations | 2326 |
| 24.4.2.2.1. Enable default forwarding | 2327 |
| 24.4.2.2.2. Modify default forwarding configurations | 2327 |
| 24.4.2.2.3. Disable default forwarding | 2327 |
| 24.4.3. Global recursive resolution | 2327 |
| 24.4.3.1. Enable global recursive resolution | 2327 |
| 24.4.3.2. Disable global recursive resolution | 2328 |
| 24.5. PrivateZone (DNS Standard Edition only) | 2328 |
| 24.5.1. Tenant internal domain name | 2328 |
| 24.5.1.1. View a domain name | 2328 |
| 24.5.1.2. Add a domain name | 2328 |
| 24.5.1.3. Bind an organization to a VPC | 2328 |
| 24.5.1.4. Unbind a domain name from a VPC | 2329 |
| 24.5.1.5. Add a description for a domain name | 2329 |
| 24.5.1.6. Delete a domain name | 2329 |
| 24.5.1.7. Delete multiple domain names | 2329 |
| 24.5.1.8. Configure DNS records | 2330 |
| 24.5.1.9. View a resolution policy | 2334 |
| 24.5.2. Tenant forwarding configurations | 2335 |
| 24.5.2.1. Tenant forwarding domain names | 2335 |
| 24.5.2.1.1. View a tenant forwarding domain name | 2335 |
| 24.5.2.1.2. Add a tenant forwarding domain name | 2335 |
| 24.5.2.1.3. Bind an organization to a VPC | 2336 |
| 24.5.2.1.4. Unbind a domain name from a VPC | 2336 |

| | |
|--|------|
| 24.5.2.1.5. Modify the forwarding configurations of a dom... | 2337 |
| 24.5.2.1.6. Add a description for a tenant forwarding dom... | 2337 |
| 24.5.2.1.7. Delete a tenant forwarding domain name | 2337 |
| 24.5.2.1.8. Delete multiple tenant forwarding domain nam... | 2337 |
| 24.5.2.2. Tenant default forwarding configurations | 2338 |
| 24.5.2.2.1. View default forwarding configurations | 2338 |
| 24.5.2.2.2. Add a default forwarding configuration | 2338 |
| 24.5.2.2.3. Bind an organization to a VPC | 2338 |
| 24.5.2.2.4. Unbind a domain name from a VPC | 2339 |
| 24.5.2.2.5. Modify a default forwarding configuration | 2339 |
| 24.5.2.2.6. Add a default forwarding configuration | 2339 |
| 24.5.2.2.7. Delete a default forwarding configuration | 2340 |
| 24.5.2.2.8. Delete multiple default forwarding configuratio... | 2340 |
| 24.6. Internal Global Traffic Manager (internal GTM Standard E.. | 2340 |
| 24.6.1. Scheduling instance management | 2340 |
| 24.6.1.1. Scheduling Instance | 2340 |
| 24.6.1.1.1. Create a scheduling instance | 2341 |
| 24.6.1.1.2. Modify a scheduling instance | 2341 |
| 24.6.1.1.3. Configure a scheduling instance | 2341 |
| 24.6.1.1.3.1. Create an access policy for a scheduling inst... | 2343 |
| 24.6.1.1.3.2. Modify the access policy of a scheduling ins... | 2345 |
| 24.6.1.1.3.3. Delete the access policy of a scheduling ins... | 2346 |
| 24.6.1.1.4. Delete a scheduling instance | 2346 |
| 24.6.1.2. Address Pool | 2346 |
| 24.6.1.2.1. Create an address pool | 2347 |
| 24.6.1.2.2. Modify the configurations of an address pool | 2347 |
| 24.6.1.2.3. Delete an address pool | 2347 |
| 24.6.1.2.4. Enable health check | 2348 |

| | |
|---|------|
| 24.6.1.3. Scheduling Domain | 2349 |
| 24.6.1.3.1. Create a scheduling domain | 2349 |
| 24.6.1.3.2. Add a description for a scheduling domain | 2349 |
| 24.6.1.3.3. Delete a scheduling domain | 2349 |
| 24.6.1.4. View alert logs | 2349 |
| 24.6.2. Scheduling line management | 2350 |
| 24.6.2.1. IP Address Line Configuration | 2350 |
| 24.6.2.1.1. Add a line | 2350 |
| 24.6.2.1.2. Sort lines | 2350 |
| 24.6.2.1.3. Modify the configurations of a line | 2350 |
| 24.6.2.1.4. Delete a line | 2350 |
| 24.6.3. Data synchronization management | 2350 |
| 24.6.3.1. Synchronization cluster management | 2350 |

1. Apsara Uni-manager Management Console

1.1. What is the Apsara Uni-manager Management Console?

The Apsara Uni-manager Management Console is a service capability platform based on the Alibaba Cloud Apsara Stack platform and designed for government and enterprise customers. This platform improves IT management and troubleshooting and is dedicated to providing a leading service capability platform of the cloud computing industry. It provides large-scale and cost-efficient end-to-end cloud computing and big data services for customers in industries such as government, education, healthcare, finance, and enterprise.

Overview

The Apsara Uni-manager Management Console simplifies the management and deployment of physical and virtual resources by building an Apsara Stack platform that supports various business types of government and enterprise customers. The console helps you build your business systems in a simple and quick manner, fully improve resource utilization, and reduce O&M costs. This allows you to shift your focus from O&M to business. The console brings the Internet economy model to government and enterprise customers, and builds a new ecosystem chain based on cloud computing.

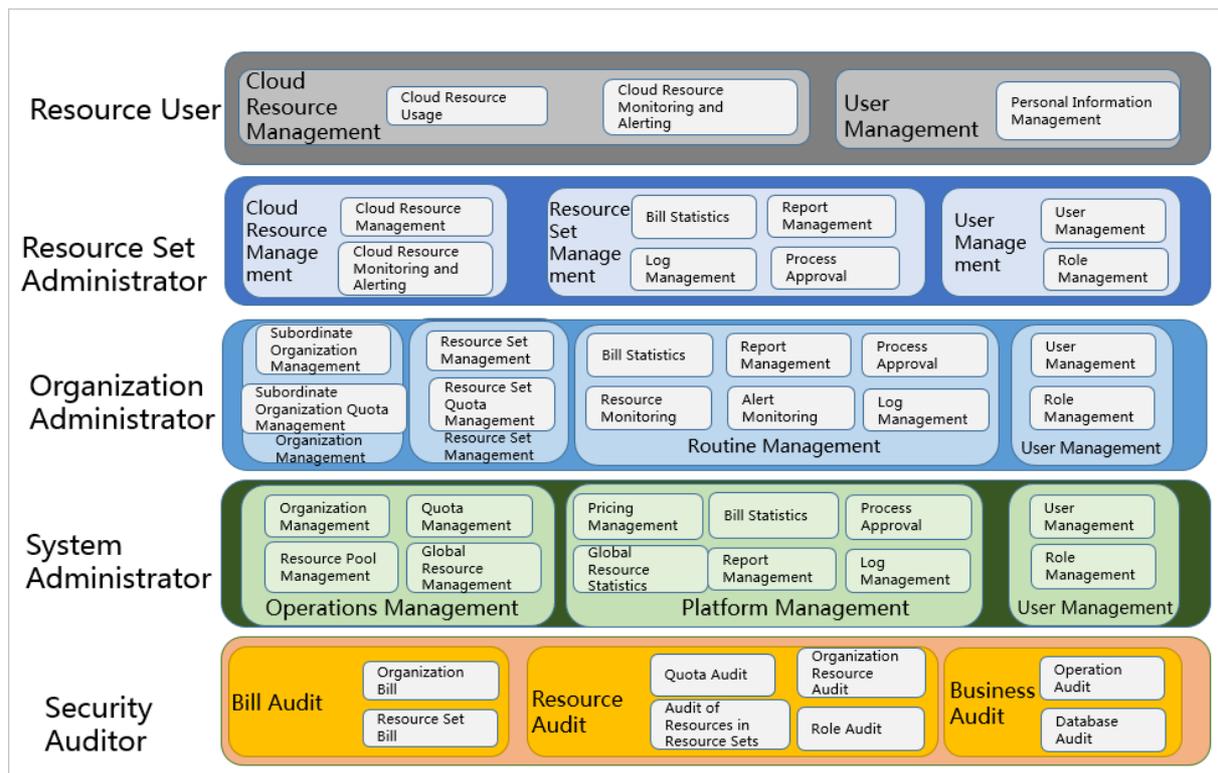
Workflow

Operations in the Apsara Uni-manager Management Console are divided into the following parts:

- **System initialization:** This part is designed to complete basic system configurations, such as creating organizations, resource sets, and users, creating basic resources such as VPCs, and creating contacts and contact groups in Cloud Monitor.
- **Cloud resource creation:** This part is designed to create resources.
- **Cloud resource management:** This part is designed to complete resource management operations, such as starting, using, and releasing resources, and changing resource configurations and resource quotas.

1.2. User roles and permissions

This topic describes roles and their permissions.



Roles and permissions

| Role | Role permission |
|-------------------------------------|--|
| Resource user | This role has the permissions to view and modify resources in a resource set and create alert rules. |
| Resource set administrator | This role has the permissions to create, modify, and delete resources in a resource set and manage the users of the resource set. |
| Organization administrator | This role has the permissions to manage an organization and its subordinate organizations, create, modify, and delete the resources of organizations, create and view alert rules for resources, and export reports. |
| Operations administrator | This role has read and write permissions on all resources. |
| Security auditor | This role performs security audit on the Apsara Uni-manager Management Console and has the read-only permissions on operation logs of the Apsara Uni-manager Management Console. |
| Platform administrator | This role has the permissions to initialize the system and create operations administrators. |
| Resource auditor | This role has the read-only permissions on all resources in the Apsara Uni-manager Management Console. |
| Organization security administrator | This role manages the security of an organization, including the security of hosts, applications, and networks. This role has the read-only permissions on operation logs of the Apsara Uni-manager Management Console and read and write permissions on ApsaraDB RDS, ECS, and Apsara Stack Security. |

| Role | Role permission |
|---|--|
| Security system configuration administrator | This role configures system security features such as the upgrade center and global configurations. This role has read and write permissions on the upgrade, protection, and configuration features of Apsara Stack Security. |
| Global organization security administrator | This role manages the security of global tenants by using Cloud Security Operation Center (SOC). This role has read and write permissions on all features of Apsara Stack Security. |
| Platform security administrator | This role manages the security of the Apsara Uni-manager Management Console by using SOC. |
| Global organization security auditor | This role checks the security conditions of all organizations by using SOC. This role has the read-only permissions on operation logs of the Apsara Uni-manager Management Console and all features of Apsara Stack Security. |
| Platform security auditor | This role checks the security conditions of the Apsara Uni-manager Management Console by using SOC. This role has the read-only permissions on operation logs of the Apsara Uni-manager Management Console, Server Guard, Cloud Firewall, Sensitive Data Discovery and Protection, SOC, system configurations, and Web Application Firewall (WAF) configurations as well as read and write permissions on Anti-DDoS, Threat Detection, and Update Center of Apsara Stack Security. |
| Platform security configuration administrator | This role configures and has read and write permissions on security services in the Apsara Uni-manager Management Console, such as Server Guard and WAF. |
| Organization resource auditor | This role has the read-only permissions on all resources in an organization to which it belongs. |

1.3. Log on to the Apsara Uni-manager Management Console

This topic describes how to log on to the Apsara Uni-manager Management Console.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

Note When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

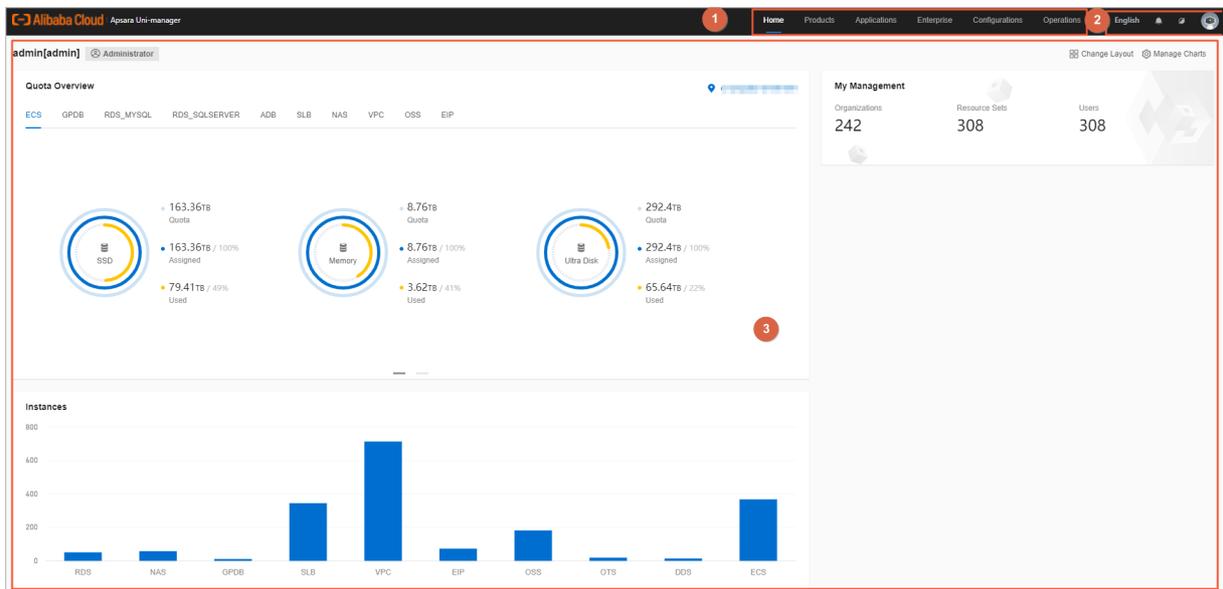
- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click Login.

1.4. Web page introduction

The web page of the Apsara Uni-manager Management Console consists of the search box, top navigation bar, information section of the current logon user, and operation section.

Apsara Uni-manager Management Console page



Functional sections of the web page

| Section | Description |
|--------------------------|---|
| 1 Top navigati on bar | <p>This section includes the following modules:</p> <ul style="list-style-type: none"> • Home: uses charts to show the usage and monitoring data of existing system resources in each region. • Products: manages all types of basic cloud services and resources. • Applications: manages cloud apps of enterprises. • Enterprise: manages organizations, resource sets, roles, users, logon policies, user groups, ownership, and resource pools. • Configurations: manages resource pools, password policies, specifications, menus, and Resource Access Management (RAM) roles. • Operations: manages the daily operations of cloud resources, including usage statistics and quotas. • Security: provides operation and system logs. |

| Section | | Description |
|---------|---|--|
| 2 | Information section of the current logon user | <ul style="list-style-type: none"> •  English : allows you to switch between English, simplified Chinese, and traditional Chinese. •  : provides message notifications. •  : allows you to switch between day and night modes. • User Information: When you click the  icon of the current logon user, the User Information, View Version, and Exit menu items are displayed. If you click User Information, you can perform the following operations on the User Information page: <ul style="list-style-type: none"> ◦ View basic information. ◦ Modify personal information. ◦ Change the logon password. ◦ View the AccessKey pair of your Apsara Stack tenant account. ◦ Switch the current role. ◦ Enable or disable alert notification. |
| 3 | Operation section | <p>Operation section: shows the information and operations.</p> |

1.5. Initial configuration

1.5.1. Configuration description

Before you use the Apsara Uni-manager Management Console, you must complete a series of basic configuration operations as an administrator, such as creating organizations, resource sets, users, and roles and initializing resources. This is the initial system configuration.

The Apsara Uni-manager Management Console manages the organizations, resource sets, users, and roles of cloud data centers in a centralized and service-oriented manner to grant different resource access permissions to different users.

- **Organization**

After the Apsara Uni-manager Management Console is deployed, a root organization is automatically generated. You can create other organizations under the root organization.

Organizations are displayed in a hierarchical structure. You can create subordinate organizations under each organization level.
- **Resource Set**

A resource set is a container used to store resources. Each resource must belong to a resource set.
- **User**

A user is a resource manager and user.
- **Role**

A role is a set of access permissions. You can assign different roles to different users to implement system access control to meet a variety of different requirements.

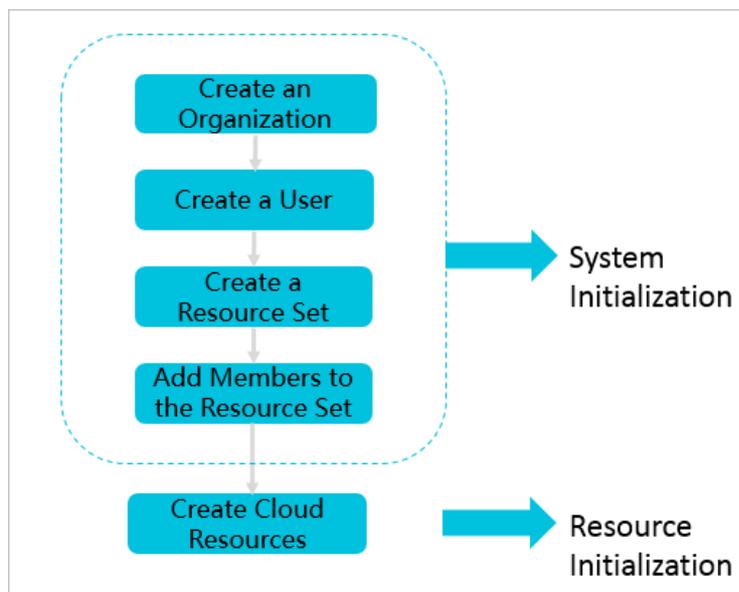
The following table describes the relationships among organizations, resource sets, users, roles, and cloud resources.

| Relationship between two items | Relationship type | Description |
|--------------------------------|-------------------|---|
| Organization and resource set | One-to-many | An organization can have multiple resource sets, but each resource set can belong to only a single organization. |
| Organization and user | One-to-many | An organization can have multiple users, but each user can belong to only a single organization. |
| Resource set and user | Many-to-many | A user can have multiple resource sets, and a resource set can be assigned to multiple users under the same level-1 organization. |
| User and role | Many-to-many | A user can have multiple roles, and a role can be assigned to multiple users. |
| Resource set and resource | One-to-many | A resource set can have multiple resources, but each cloud resource can belong to only a single resource set. |

1.5.2. Configuration process

This topic describes the initial configuration process.

Before you use the Apsara Uni-manager Management Console, you must complete the initial system configurations as an administrator based on the process shown in the following figure.



- Create an organization**
 Create an organization to store resource sets and their resources.
- Create a user**
 Create a user and assign the user different roles to meet different requirements for system access control.
- Create a resource set**
 Create a resource set before you apply for resources.
- Add a member to a resource set**
 Add users to the resource set.

5. Create cloud resources

Create instances in each service console based on project requirements. For more information about how to create cloud service instances, see the user guide of each cloud service.

1.6. Monitoring

1.6.1. View the workbench

The Apsara Uni-manager Management Console uses charts to keep you up to date on the current usage and monitoring information of resources.

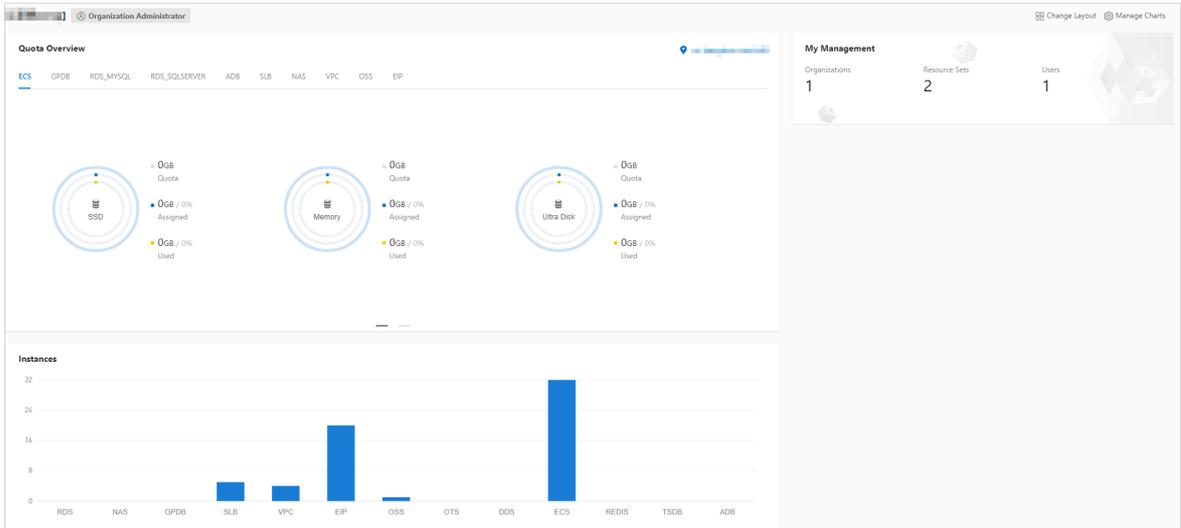
Context

The resource types displayed may vary with region types. See your dashboard for available resource types.

Procedure

1. **Log on to the Apsara Uni-manager Management Console.**

By default, the **workbench** page appears when you log on to the Apsara Uni-manager Management Console. To return to the workbench page from other pages, click **Home** in the top navigation bar.



2. On the **workbench** page, you can view the instance summary information for all regions of the Apsara Stack environment.

You can click **Manage Charts** in the upper-right corner of the page to select all or individual modules to view relevant information. You can also click **Change Layout** in the upper-right corner of the page and drag a specific module to the desired location.

- o **Quota Overview**
Shows the usage and quotas of Elastic Compute Service (ECS), ApsaraDB RDS, Object Storage Service (OSS), and Server Load Balancer (SLB) resources.
- o **Instances**
Shows the numbers of ECS instances, ApsaraDB RDS instances, OSS buckets, and SLB instances in each region.
- o **Instance Trends**
Shows the numbers of recent ECS instances, ApsaraDB RDS instances, OSS buckets, and SLB instances.
- o **Resource Load**

Shows the top five ECS and ApsaraDB RDS instances in terms of disk usage, CPU utilization, and memory usage.

- **Alert Rules**

Shows the number of alerts and details of the alerts.

- **My Management**

Shows the numbers of organizations, resource sets, and users.

- **Multi-cloud Regions**

Shows the information of all primary and secondary nodes in Apsara Stack. The network connection status and related alerts are displayed for each secondary node.

- **Multi-cloud Resources**

Shows the cloud services and the number of instances in each secondary node.

1.6.2. CloudMonitor

1.6.2.1. Cloud Monitor overview

Cloud Monitor provides real-time monitoring, alerting, and notification services for resources to protect your services and businesses.

Cloud Monitor can monitor metrics for a variety of services such as ECS, ApsaraDB for RDS, SLB, OSS, KVStore for Redis, VPN Gateway, AnalyticDB for PostgreSQL, ApsaraDB for MongoDB, EIP, and API Gateway.

You can use the metrics of cloud services to configure alert rules and notification policies. This way, you can stay up to date on the running status and performance of your service instances and scale resources in a timely manner when resources are insufficient.

1.6.2.2. Metrics

This topic describes the metrics available for each service.

CloudMonitor checks the availability of services based on their metrics. You can configure alert rules and notification policies for these metrics to stay up to date on the running status and performance of monitored service instances.

CloudMonitor can monitor resources of other services, including Elastic Compute Service (ECS), ApsaraDB RDS, Server Load Balancer (SLB), Object Storage Service (OSS), KVStore for Redis, VPN Gateway, AnalyticDB for PostgreSQL, ApsaraDB for MongoDB, Elastic IP Address (EIP), and API Gateway. The following tables list the metrics for each service.

Operating system metrics for ECS

| Metric | Description | Unit |
|--------------------------|---|------|
| Host.cpu.total | The total CPU utilization of an ECS instance. | % |
| Host.mem.usedutilization | The memory usage of an ECS instance. | % |
| Host.load1 | The system loads over the last 1 minute. This metric is unavailable for Windows operating systems. | N/A |
| Host.load5 | The system loads over the last 5 minutes. This metric is unavailable for Windows operating systems. | N/A |

| Metric | Description | Unit |
|-----------------------|--|---------|
| Host.load15 | The system loads over the last 15 minutes. This metric is unavailable for Windows operating systems. | N/A |
| Host.disk.utilization | The disk usage of an ECS instance. | % |
| Host.disk.readbytes | The number of bytes read from the disk per second. | byte/s |
| Host.disk.writebytes | The number of bytes written to the disk per second. | byte/s |
| Host.disk.readlops | The number of read requests received by the disk per second. | count/s |
| Host.disk.writelops | The number of write requests received by the disk per second. | count/s |
| Host.fs.inode | The inode usage. | % |

Basic metrics for ECS

| Metric | Description | Unit |
|--|--|---------|
| CPU utilization | The CPU utilization of an ECS instance. | % |
| Inbound bandwidth to the Internet | The average rate of inbound traffic to the Internet. | bit/s |
| Inbound bandwidth to the internal network | The average rate of inbound traffic to the internal network. | bit/s |
| Outbound bandwidth from the Internet | The average rate of outbound traffic from the Internet. | bit/s |
| Outbound bandwidth from the internal network | The average rate of outbound bandwidth from the internal network. | bit/s |
| System disk BPS | The number of bytes read from and written to the system disk per second. | byte/s |
| System disk IOPS | The number of reads from and writes to the system disk per second. | count/s |
| Advance CPU credits | The changes in advance CPU credits. Advance CPU credits can be used only when the unlimited mode is enabled. | N/A |
| CPU credit consumption | The changes in CPU credit consumption. Consumption trends are consistent with CPU utilization. | N/A |

| Metric | Description | Unit |
|-----------------------|--|------|
| Overdrawn CPU credits | The changes in overdrawn CPU credits. Overdrawn CPU credits can be used only when the unlimited mode is enabled. | N/A |
| CPU credit balance | The changes in CPU credit balance. The CPU credit balance is used to maintain CPU credit usage. | N/A |

 **Note**

For ECS instances, you must install a monitoring plug-in to collect metric data at the operating system level.
 Installation method: On the **Cloud Monitor** page, select the instance that you want to monitor from the ECS instance list and click **Batch Install** in the lower part of the page.
 Metric data is displayed in the monitoring chart within 5 to 10 minutes after the monitoring plug-in is installed.

Metrics for ApsaraDB RDS for PostgreSQL

| Metric | Description | Apsara Stack service | Formula |
|------------------|---|-----------------------------|---|
| CPU utilization | The CPU utilization of an ApsaraDB RDS for PostgreSQL instance. Unit: %. | ApsaraDB RDS for PostgreSQL | Used CPU cores of an ApsaraDB RDS for PostgreSQL instance/Total CPU cores of the ApsaraDB RDS for PostgreSQL instance |
| Memory usage | The memory usage of an ApsaraDB RDS for PostgreSQL instance. Unit: %. | ApsaraDB RDS for PostgreSQL | Used memory of an ApsaraDB RDS for PostgreSQL instance/Total memory of the ApsaraDB RDS for PostgreSQL instance |
| Disk usage | The disk usage of an ApsaraDB RDS for PostgreSQL instance. Unit: %. | ApsaraDB RDS for PostgreSQL | None |
| IOPS usage | The number of I/O requests for an ApsaraDB RDS for PostgreSQL instance per second. Unit: %. | ApsaraDB RDS for PostgreSQL | Number of I/O requests for an ApsaraDB RDS for PostgreSQL instance/Statistical period |
| Connection usage | The number of connections between an application and an ApsaraDB RDS for PostgreSQL instance per second. Unit: %. | ApsaraDB RDS for PostgreSQL | Number of connections between an application and an ApsaraDB RDS for PostgreSQL instance/Statistical period |

Metrics for ApsaraDB RDS for MySQL

| Metric | Description | Apsara Stack service | Formula |
|-----------------|---|------------------------|---|
| CPU utilization | The CPU utilization of an ApsaraDB RDS for MySQL instance. Unit: %. | ApsaraDB RDS for MySQL | Used CPU cores of an ApsaraDB RDS for MySQL instance/Total CPU cores of the ApsaraDB RDS for MySQL instance |
| Memory usage | The memory usage of an ApsaraDB RDS for MySQL instance. Unit: %. | ApsaraDB RDS for MySQL | Used memory of an ApsaraDB RDS for MySQL instance/Total memory of the ApsaraDB RDS for MySQL instance |

| Metric | Description | Apsara Stack service | Formula |
|--|--|------------------------|--|
| Disk usage | The disk usage of an ApsaraDB RDS for MySQL instance. Unit: %. | ApsaraDB RDS for MySQL | None |
| IOPS usage | The number of I/O requests for an ApsaraDB RDS for MySQL instance per second. Unit: %. | ApsaraDB RDS for MySQL | Number of I/O requests for an ApsaraDB RDS for MySQL instance/Statistical period |
| Connection usage | The number of connections between an application and an ApsaraDB RDS for MySQL instance per second. Unit: %. | ApsaraDB RDS for MySQL | Number of connections between an application and an ApsaraDB RDS for MySQL instance/Statistical period |
| Inbound bandwidth to ApsaraDB RDS for MySQL | The inbound traffic to an ApsaraDB RDS for MySQL instance per second. | ApsaraDB RDS for MySQL | None |
| Outbound bandwidth from ApsaraDB RDS for MySQL | The outbound traffic from an ApsaraDB RDS for MySQL instance per second. | ApsaraDB RDS for MySQL | None |

Metrics for ApsaraDB RDS for SQL Server

| Metric | Description | Apsara Stack service | Formula |
|--|---|-----------------------------|---|
| CPU utilization | The CPU utilization of an ApsaraDB RDS for SQL Server instance. Unit: %. | ApsaraDB RDS for SQL Server | Used CPU cores of an ApsaraDB RDS for SQL Server instance/Total CPU cores of the ApsaraDB RDS for SQL Server instance |
| Memory usage | The memory usage of an ApsaraDB RDS for SQL Server instance. Unit: %. | ApsaraDB RDS for SQL Server | Used memory of an ApsaraDB RDS for SQL Server instance/Total memory of the ApsaraDB RDS for SQL Server instance |
| Disk usage | The disk usage of an ApsaraDB RDS for SQL Server instance. Unit: %. | ApsaraDB RDS for SQL Server | None |
| IOPS usage | The number of I/O requests for an ApsaraDB RDS for SQL Server instance per second. Unit: %. | ApsaraDB RDS for SQL Server | Number of I/O requests for an ApsaraDB RDS for SQL Server instance/Statistical period |
| Connection usage | The number of connections between an application and an ApsaraDB RDS for SQL Server instance per second. Unit: %. | ApsaraDB RDS for SQL Server | Number of connections between an application and an ApsaraDB RDS for SQL Server instance/Statistical period |
| Inbound bandwidth to ApsaraDB RDS for SQL Server | The inbound traffic to an ApsaraDB RDS for SQL Server instance per second. | ApsaraDB RDS for SQL Server | None |

| Metric | Description | Apsara Stack service | Formula |
|---|---|-----------------------------|---------|
| Outbound bandwidth from ApsaraDB RDS for SQL Server | The outbound traffic from an ApsaraDB RDS for SQL Server instance per second. | ApsaraDB RDS for SQL Server | None |

Metrics for PolarDB

| Metric | Description | Apsara Stack service | Formula |
|------------------|--|----------------------|--|
| CPU utilization | The CPU utilization of a PolarDB instance. Unit: %. | PolarDB | Used CPU cores of a PolarDB instance/Total CPU cores of the PolarDB instance |
| Memory usage | The memory usage of a PolarDB instance. Unit: %. | PolarDB | Used memory of a PolarDB instance/Total memory of the PolarDB instance |
| Disk usage | The disk usage of a PolarDB instance. Unit: %. | PolarDB | None |
| IOPS usage | The number of I/O requests for a PolarDB instance per second. Unit: %. | PolarDB | Number of I/O requests for a PolarDB instance/Statistical period |
| Connection usage | The number of connections between an application and a PolarDB instance per second. Unit: %. | PolarDB | Number of connections between an application and a PolarDB instance/Statistical period |

Metrics for SLB

| Metric | Description | Unit |
|--|---|---------|
| Inbound bandwidth on a port | The average rate of inbound traffic on a port. | bit/s |
| Outbound bandwidth on a port | The average rate of outbound traffic on a port. | bit/s |
| Number of new connections on a port | The average number of new TCP connections established between clients and SLB instances in a statistical period. | N/A |
| Number of inbound packets received on a port | The number of packets received by an SLB instance per second. | count/s |
| Number of outbound packets sent on a port | The number of packets sent by an SLB instance per second. | count/s |
| Number of active connections on a port | The number of TCP connections in the ESTABLISHED state. If persistent connections are used, a connection can transfer multiple file requests at one time. | N/A |

| Metric | Description | Unit |
|--|---|---------|
| Number of inactive connections on a port | The number of TCP connections that are not in the ESTABLISHED state. You can run the <code>netstat -an</code> command to view the connections for both Windows and Linux instances. | N/A |
| Number of concurrent connections on a port | The number of established TCP connections. | count/s |
| Number of dropped connections on a port | The number of connections dropped per second. | count/s |
| Number of dropped inbound packets on a port | The number of inbound packets dropped per second. | count/s |
| Number of dropped outbound packets on a port | The number of outbound packets dropped per second. | count/s |
| Dropped inbound bandwidth on a port | The amount of inbound traffic dropped per second. | bit/s |
| Dropped outbound bandwidth on a port | The amount of outbound traffic dropped per second. | bit/s |

Metrics for monitoring service overview of OSS

| Metric | Description | Unit |
|--|--|------|
| Availability | The metric that describes the system availability of OSS. You can obtain the metric value based on the following formula: Metric value = 1 - Server error requests with the returned HTTP status code 5xx/All requests . | % |
| Valid request percentage | The percentage of valid requests out of all requests. | % |
| Total number of requests | The total number of requests that are received and processed by the OSS server. | N/A |
| Number of valid requests | The total number of requests with HTTP status codes 2xx and 3xx returned. | N/A |
| Outbound traffic from the Internet | The amount of outbound traffic from the Internet. | byte |
| Inbound traffic to the Internet | The amount of inbound traffic to the Internet. | byte |
| Outbound traffic from the internal network | The amount of outbound traffic from the internal network. | byte |
| Inbound traffic to the internal network | The amount of inbound traffic to the internal network. | byte |

| Metric | Description | Unit |
|--|--|------|
| CDN outbound traffic | The amount of outbound traffic sent over CDN after CDN is activated. Such outbound traffic over CDN is back-to-origin traffic. | byte |
| CDN inbound traffic | The amount of inbound traffic received over CDN after CDN is activated. | byte |
| Outbound traffic of cross-region replication | The amount of outbound traffic generated during data replication after cross-region replication is enabled. | byte |
| Inbound traffic of cross-region replication | The amount of inbound traffic generated during data replication after cross-region replication is enabled. | byte |
| Storage size | The amount of total storage occupied by the buckets of a specified user before the statistics collection deadline. | byte |
| Number of PUT requests | The total number of PUT requests made by the user between 00:00:00 on the first day of the current month and the statistics collection deadline. | N/A |
| Number of GET requests | The total number of GET requests made by the user between 00:00:00 on the first day of the current month and the statistics collection deadline. | N/A |

Metrics for request status details of OSS

| Metric | Description | Unit |
|--|---|------|
| Number of requests with server-side errors | The total number of system-level error requests with the returned HTTP status code 5xx. | N/A |
| Percentage of requests with server-side errors | The percentage of requests with server-side errors out of all requests. | % |
| Number of requests with network errors | The total number of requests with the returned HTTP status code 499. | N/A |
| Percentage of requests with network errors | The percentage of requests with network errors out of all requests. | % |
| Number of requests with client-side authorization errors | The total number of requests with the returned HTTP status code 403. | N/A |
| Percentage of requests with client-side authorization errors | The percentage of requests with authorization errors out of all requests. | % |

| Metric | Description | Unit |
|---|--|------|
| Number of requests with client-side errors indicating resources not found | The total number of requests with the returned HTTP status code 404. | N/A |
| Percentage of requests with client-side errors indicating resources not found | The percentage of requests with errors indicating resources not found out of all requests. | % |
| Number of requests with client-side timeout errors | The total number of requests with the returned HTTP status code 408 or OSS error code RequestTimeout. | N/A |
| Percentage of requests with client-side timeout errors | The percentage of requests with client-side timeout errors out of all requests. | % |
| Number of requests with other client-side errors | The total number of requests other than the foregoing client-side error requests with the returned HTTP status code 4xx. | N/A |
| Percentage of requests with other client-side errors | The percentage of requests with other client-side errors out of all requests. | % |
| Number of successful requests | The total number of requests with the returned HTTP status code 2xx. | N/A |
| Percentage of successful requests | The percentage of successful requests out of all requests. | % |
| Number of redirected requests | The total number of requests with the returned HTTP status code 3xx. | N/A |
| Percentage of redirected requests | The percentage of redirected requests out of all requests. | % |

Metrics for maximum latency of OSS

| Metric | Description | Unit |
|---|---|------|
| Maximum end-to-end latency of GetObject requests | The maximum end-to-end latency of successful GetObject requests. | ms |
| Maximum server latency of GetObject requests | The maximum server latency of successful GetObject requests. | ms |
| Maximum end-to-end latency of HeadObject requests | The maximum end-to-end latency of successful HeadObject requests. | ms |
| Maximum server latency of HeadObject requests | The maximum server latency of successful HeadObject requests. | ms |
| Maximum end-to-end latency of PutObject requests | The maximum end-to-end latency of successful PutObject requests. | ms |
| Maximum server latency of PutObject requests | The maximum server latency of successful PutObject requests. | ms |

| Metric | Description | Unit |
|---|---|------|
| Maximum end-to-end latency of PostObject requests | The maximum end-to-end latency of successful PostObject requests. | ms |
| Maximum server latency of PostObject requests | The maximum server latency of successful PostObject requests. | ms |
| Maximum end-to-end latency of AppendObject requests | The maximum end-to-end latency of successful AppendObject requests. | ms |
| Maximum server latency of AppendObject requests | The maximum server latency of successful AppendObject requests. | ms |
| Maximum end-to-end latency of UploadPart requests | The maximum end-to-end latency of successful UploadPart requests. | ms |
| Maximum server latency of UploadPart requests | The maximum server latency of successful UploadPart requests. | ms |
| Maximum end-to-end latency of UploadPartCopy requests | The maximum end-to-end latency of successful UploadPartCopy requests. | ms |
| Maximum server latency of UploadPartCopy requests | The maximum server latency of successful UploadPartCopy requests. | ms |

Metrics for successful request category of OSS

| Metric | Description | Unit |
|--|---|------|
| Number of successful GetObject requests | The number of successful GetObject requests. | N/A |
| Number of successful HeadObject requests | The number of successful HeadObject requests. | N/A |
| Number of successful PostObject requests | The number of successful PostObject requests. | N/A |
| Number of successful AppendObject requests | The number of successful AppendObject requests. | N/A |
| Number of successful UploadPart requests | The number of successful UploadPart requests. | N/A |
| Number of successful UploadPartCopy requests | The number of successful UploadPartCopy requests. | N/A |
| Number of successful DeleteObject requests | The number of successful DeleteObject requests. | N/A |
| Number of successful DeleteObjects requests | The number of successful DeleteObjects requests. | N/A |

Metrics for KVStore for Redis

| Metric | Description | Apsara Stack service | Unit |
|-----------------|--|----------------------|------|
| CPU utilization | The CPU utilization of a KVStore for Redis instance. | KVStore for Redis | % |

| Metric | Description | Apsara Stack service | Unit |
|--|---|----------------------|---------|
| Memory usage | The percentage of memory that is in use. | KVStore for Redis | % |
| Used memory | The amount of memory that is in use. | KVStore for Redis | byte |
| Number of used connections | The total number of client connections that are in use. | KVStore for Redis | N/A |
| Percentage of used connections | The percentage of connections that are in use. | KVStore for Redis | % |
| Write bandwidth | The write traffic per second. | KVStore for Redis | byte/s |
| Read bandwidth | The read traffic per second. | KVStore for Redis | byte/s |
| Number of failed operations per second | The number of failed operations on a KVStore for Redis instance per second. | KVStore for Redis | count/s |
| Write bandwidth usage | The percentage of total bandwidth used by write operations. | KVStore for Redis | % |
| Read bandwidth usage | The percentage of total bandwidth used by read operations. | KVStore for Redis | % |
| Used QPS | The number of queries per second (QPS). | KVStore for Redis | count/s |
| QPS usage | The QPS usage. | KVStore for Redis | % |
| Average response time | The average response time. | KVStore for Redis | ms |
| Maximum response time | The maximum response time. | KVStore for Redis | ms |
| Number of failed commands | The number of failed commands. | KVStore for Redis | N/A |
| Hit rate | The current hit rate. | KVStore for Redis | % |
| Inbound traffic | The inbound traffic to a KVStore for Redis instance. | KVStore for Redis | byte |
| Inbound bandwidth usage | The inbound bandwidth usage of a KVStore for Redis instance. | KVStore for Redis | % |
| Outbound traffic | The outbound traffic from a KVStore for Redis instance. | KVStore for Redis | byte |
| Outbound bandwidth usage | The outbound bandwidth usage of a KVStore for Redis instance. | KVStore for Redis | % |

Metrics for VPN Gateway

| Metric | Dimension | Monitoring period | Unit |
|--|-------------------|-------------------|------|
| Number of inbound packets in a connection per second | User and instance | 1 minute | pps |

| Metric | Dimension | Monitoring period | Unit |
|---|-------------------|-------------------|-------|
| Number of outbound packets in a connection per second | User and instance | 1 minute | pps |
| Inbound bandwidth of a connection | User and instance | 1 minute | bit/s |
| Outbound bandwidth of a connection | User and instance | 1 minute | bit/s |
| Number of connections | User and instance | 1 minute | N/A |

Metrics for AnalyticDB for PostgreSQL

| Metric | Description | Unit |
|------------------|--|------|
| Connection usage | The number of connections between an application and an AnalyticDB for PostgreSQL instance per second. | % |
| CPU utilization | The CPU utilization of an AnalyticDB for PostgreSQL instance. | % |
| Disk usage | The disk usage of an AnalyticDB for PostgreSQL instance. | % |
| IOPS usage | The number of I/O requests for an AnalyticDB for PostgreSQL instance per second. | % |
| Memory usage | The memory usage of an AnalyticDB for PostgreSQL instance. | % |

Metrics for ApsaraDB for MongoDB

| Tab | Metric | Description | Unit |
|--------------|-----------------|--|------|
| Basic metric | CPU utilization | The CPU utilization of an ApsaraDB for MongoDB instance. | % |
| | Memory usage | The memory usage of an ApsaraDB for MongoDB instance. | % |
| | Disk usage | The disk usage of an ApsaraDB for MongoDB instance. | % |
| | IOPS usage | The percentage of the IOPS used by an ApsaraDB for MongoDB instance out of the maximum available IOPS. | % |

| Tab | Metric | Description | Unit |
|----------------------|--|---|------|
| | Connection usage | The number of connections between an application and an ApsaraDB for MongoDB instance per second. | % |
| | QPS | The number of queries per second. | N/A |
| | Number of used connections | The number of current connections to an ApsaraDB for MongoDB instance. | N/A |
| Disk capacity | Disk space occupied by an instance | The total used space. | byte |
| | Disk space occupied by data | The disk space occupied by data. | byte |
| | Disk space occupied by logs | The disk space occupied by logs. | byte |
| Network request | Inbound traffic to the internal network | The inbound traffic. | byte |
| | Outbound traffic from the internal network | The outbound traffic. | byte |
| | Number of requests | The number of processed requests. | N/A |
| Number of operations | Number of Insert operations | None | N/A |
| | Number of Query operations | None | N/A |
| | Number of Update operations | None | N/A |
| | Number of Delete operations | None | N/A |
| | Number of Getmore operations | None | N/A |
| | Number of Command operations | None | N/A |

Metrics for EIP

| Metric | Description | Dimension | Monitoring period | Unit |
|-------------------|--|-----------|-------------------|-------|
| Inbound bandwidth | The traffic that passes through EIP to ECS per second. | Instance | 1 minute | bit/s |

| Metric | Description | Dimension | Monitoring period | Unit |
|---------------------------------------|---|-----------|-------------------|-------|
| Outbound bandwidth | The traffic that passes through EIP from ECS per second. | Instance | 1 minute | bit/s |
| Number of inbound packets per second | The number of packets that pass through EIP to ECS per second. | Instance | 1 minute | pps |
| Number of outbound packets per second | The number of packets that pass through EIP from ECS per second. | Instance | 1 minute | pps |
| Packet loss rate due to throttling | The packet loss rate when the actually used bandwidth exceeds the configured upper limit. | Instance | 1 minute | pps |

Metrics for API Gateway

| Metric | Description | Dimension | Unit | Monitoring period |
|--------------------------|---|--------------|------|-------------------|
| Error distribution | The number of 2xx, 4xx, and 5xx status codes returned for an API in the monitoring period. | User and API | N/A | 1 minute |
| Inbound traffic | The total traffic of requests received by an API in the monitoring period. | User and API | byte | 1 minute |
| Outbound traffic | The total traffic of responses sent by an API in the monitoring period. | User and API | byte | 1 minute |
| Response time | The latency between the time when API Gateway calls the backend service of an API and the time when the result is received from the backend service in the monitoring period. | User and API | s | 1 minute |
| Number of total requests | The total number of requests received by an API in the monitoring period. | User and API | N/A | 1 minute |

1.6.2.3. View monitoring charts

You can view monitoring charts to obtain up-to-date information about each instance.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, choose **Products > Monitoring and O&M > CloudMonitor**.
3. In the left-side navigation pane of the CloudMonitor page, click **Cloud Service Monitoring**.
4. Click a cloud service.
5. Click **Monitoring Charts** in the **Actions** column corresponding to an instance.

On the Monitoring Charts page that appears, you can select a date and time to view the monitoring data of each metric.

1.6.3. Alerts

1.6.3.1. View alert overview

On the **Overview** page in Cloud Monitor, you can view the alert status statistics and alert logs.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, choose **Products > Monitoring and O&M > CloudMonitor**.
3. In the left-side navigation pane of the Cloud Monitor page, click **Overview**.
4. On the **Overview** page, view the alert status statistics and alert logs that are generated in the last 24 hours.

1.6.3.2. Enable or disable alert notification

You can choose whether to enable alert notification by SMS, email, or DingTalk.

Prerequisites

Valid contact information is specified when you create a user. If your contact information is changed, you must modify personal information. For more information, see [Modify personal information](#).

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#).
2. In the upper-right corner of the homepage, move the pointer over the profile picture and click **User Information**.
3. In the **Messages** section, select **Email** or **DingTalk** to enable alert notification.

To disable alert notification, clear the corresponding check box.

1.6.3.3. View alert logs

You can view alert information to stay up to date on the running status of Elastic Compute Service (ECS), ApsaraDB RDS, Server Load Balancer (SLB), KVStore for Redis, VPN Gateway, AnalyticDB for PostgreSQL, ApsaraDB for MongoDB, Elastic IP Address (EIP), API Gateway, and Object Storage Service (OSS).

Context

Alert information contains information for all items that do not comply with your configured alert rules.

 **Note**

- The system can retain up to one million alert items generated within the last three months.
- This topic describes how to view alert information for ECS. You can view the alert information for other cloud resources in a similar manner.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, choose **Products > Monitoring and O&M > Cloud Monitor**.
3. In the left-side navigation pane of the Cloud Monitor page, choose **Alerts > Alert History**.
4. On the **Alert Rule History List** page, filter alert information by rule ID, rule name, service, and date.

The following table describes the fields in the query result. Alert information fields

| Parameter | Description |
|-----------------------------|---|
| Product | The service for which the alert was triggered. |
| Fault Instance | The instance for which the alert was triggered. |
| Occurred At | The time when the alert was triggered. |
| Rule Name | The name of the alert rule. |
| Status | The status of the alert rule. |
| Notification Contact | The recipient of the alert notification. |

1.6.3.4. Alarm rules

1.6.3.4.1. View alert rules

After you create alert rules, you can view your alert rules on the Alert Rules page.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#).
2. In the top navigation bar, choose **Products > Monitoring and O&M > CloudMonitor**.
3. In the left-side navigation pane of the Cloud Monitor page, click **Cloud Service Monitoring**.
4. Click a cloud service.
5. Click **Alert Rules** in the **Actions** column corresponding to an instance to go to its **Alert Rules** page.

On the **Alert Rules** page, view the detailed information of alert rules.

1.6.3.4.2. Create an alert rule

You can create an alert rule to monitor an instance.

Prerequisites

For Elastic Compute Service (ECS) instances, you must install a monitoring plug-in to collect metric data at the operating system level.

You can use the following method to install a monitoring plug-in:

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, choose **Products > Monitoring and O&M > Cloud Monitor**.
3. In the left-side navigation pane, choose **Cloud Service Monitoring > ECS**.
4. In the ECS instance list, select the instances that you want to monitor and click **Batch Install**.

 **Note**

The monitoring chart displays monitoring data 5 to 10 minutes after the monitoring plug-in is installed.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, choose **Products > Monitoring and O&M > Cloud Monitor**.
3. In the left-side navigation pane of the Cloud Monitor page, click **Cloud Service Monitoring**.
4. Click a cloud service.
5. Click **Alert Rules** in the **Actions** column corresponding to an instance.

 **Note** You can also use the search feature to query specific instances for which you want to create alert rules.

6. On the **Alert Rules** page, click **Create Alert Rule**.

Parameters for creating an alert rule

| Parameter | Description |
|-----------------------------|---|
| Product | The monitored cloud service. |
| Resource Range | The range of resources that is associated with the alert rule. |
| Rule Description | The description of the alert rule. |
| Add Rule Description | Click Add Rule Description to go to the rule configuration panel. For more information, see Parameters for adding rule description . |
| Effective Time | Only a single alert is sent during each mute duration, even if the metric value exceeds the alert rule threshold several times in a row. |
| Effective Period | An alert is sent only when the threshold is crossed during the effective period. |
| HTTP Callback | The callback URL when the alert conditions are met. |
| Alert Contact Group | The group to which alerts are sent. |

Parameters for adding rule description

| Parameter | Description |
|--------------------|--|
| Rule Name | The name of the alert rule. The name must be 1 to 64 characters in length and can contain letters and digits. |
| Metric Name | Different products have different monitoring metrics. For more information, see Metrics . |
| Comparison | The comparison between thresholds and observed values. The comparison operators include > , >= , < , and <= . When the comparison rule is satisfied, an alert rule is triggered. |

| Parameter | Description |
|---------------------------|--|
| Threshold and Alert Level | Different metrics have different reference thresholds. |

7. Click **OK**.

1.6.3.4.3. Disable an alert rule

You can disable one or more alert rules.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, choose **Products > Monitoring and O&M > CloudMonitor**.
3. In the left-side navigation pane of the Cloud Monitor page, click **Cloud Service Monitoring**.
4. Click a cloud service.
5. Click **Alert Rules** in the **Actions** column corresponding to an instance.
6. On the **Alert Rules** page, choose **More > Disable** in the **Actions** column corresponding to the alert rule to be disabled.
7. In the Disable Alert Rule message, click **Confirm**.

1.6.3.4.4. Enable an alert rule

After an alert rule is disabled, you can re-enable it

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, choose **Products > Monitoring and O&M > CloudMonitor**.
3. In the left-side navigation pane of the Cloud Monitor page, click **Cloud Service Monitoring**.
4. Click a cloud service.
5. Click **Alert Rules** in the **Actions** column corresponding to an instance.
6. On the **Alert Rules** page, choose **More > Enable** in the **Actions** column corresponding to the alert rule to be enabled.
7. In the Enable Alert Rule message, click **Confirm**.

1.6.3.4.5. Delete an alert rule

You can delete alert rules that are no longer needed.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, choose **Products > Monitoring and O&M > CloudMonitor**.
3. In the left-side navigation pane of the Cloud Monitor page, click **Cloud Service Monitoring**.
4. Click a cloud service.
5. Click **Alert Rules** in the **Actions** column corresponding to an instance.
6. On the **Alert Rules** page, click **Delete** in the **Actions** column corresponding to the alert rule to be deleted.
7. In the Delete Alert message, click **Confirm**.

1.7. VMware Cloud on Alibaba Cloud

1.7.1. Log on to the VMware Cloud on Alibaba Cloud console

This topic describes how to log on to the VMware Cloud on Alibaba Cloud console.

Prerequisites

- Before you log on to the Apsara Uni-manager Management Console, you must obtain the endpoint of the console from the deployment personnel.
- A browser is available. We recommend that you use Google Chrome.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

 **Note** When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top navigation bar, choose **Products > Elastic Computing > VMware Cloud on Alibaba Cloud**.

1.7.2. Bind a VMware Cloud on Alibaba Cloud region

Before you use VMware Cloud on Alibaba Cloud, you must bind a VMware Cloud on Alibaba Cloud region to an organization.

Prerequisites

A VMware Cloud on Alibaba Cloud region is managed. For more information, see [Add a VMware node](#).

Procedure

1. Log on to the Apsara Uni-manager Management Console.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, click **Resource Pools**.
4. In the organization navigation tree, click an organization. In the **Regions** section, select the region that you want to bind.
5. Click **Update Association**.

1.7.3. Instructions

1.7.3.1. Limits

Before you use VMware Cloud on Alibaba Cloud virtual machine (VM) templates, you must familiarize yourself with the limits of instances.

General limits

- You must select appropriate operating systems for VMware Cloud on Alibaba Cloud VM templates.

The following operating systems are verified to be available in the Apsara Uni-manager Management Console:

- CentOS8.2.2004
 - CentOS7.2003
 - CentOS6.10
 - Ubuntu-20.04.1
 - Ubuntu-18.04.5
 - Ubuntu-16.04.7
 - Windows Server 2016
 - Windows Server 2019
- The Apsara Uni-manager Management Console supports VMware vSphere 6.x. Other versions of VMware vSphere, such as 5.x or 7.x, can in theory be supported. However, the specific support depends on the compatibility of the VMware Cloud on Alibaba Cloud API and must be evaluated by the R&D team of the Apsara Uni-manager Management Console.
 - You must install VMware Tools.

For more information, see the VMware documentation. Select Full Installation in the installation process.

- You must modify network interface controller configurations in the operating system of the VM.

When you create a VM in the Apsara Uni-manager Management Console, you can specify the IP address of the operating system. This feature is supported by valid network interface controller configurations.

Operating systems of VM templates must be in DHCP mode. Information such as the MAC address and universally unique identifier (UUID) in the network interface controller configurations must be removed. The following information can be retained.

```
TYPE=Ethernet
BOOTPROTO=dhcp
DEFROUTE=yes
NAME=eth0
DEVICE=eth0
ONBOOT=yes
```

 **Note**

Some configurations are required for the following operating systems:

- CentOS 6: You must clear the content in the network interface controller configuration file named `70-persistent-net.rules`. The file is stored in the `/etc/udev/rules.d/` directory.
- CentOS 7: The system generates the name for a network interface controller, such as `ifcfg-ens160`. You must modify the name to `ifcfg-eth0` to make the name take effect.
- Ubuntu 18.04, 20.04, and later: You must run the `sudo rm /etc/netplan/*.yaml` command to remove the network interface controller configurations.

1.7.3.2. Suggestions

Consider the following operation suggestions to make more efficient use of VMware Cloud on Alibaba Cloud virtual machine (VM) templates.

- Select the latest version of VM hardware.
- Select thin provision for VM disks.

Disk replication is required when you create VMs based on templates. Files of disks of the thin provision type are small in size. This can help accelerate the creation of VMs.

 **Note**

Large sizes of disk files in VM templates or slow storage write speeds may cause VM creation to time out and fail. The maximum timeout period supported by the Apsara Uni-manager Management Console is 10 minutes.

1.7.4. Instances

1.7.4.1. Create a VMware Cloud on Alibaba Cloud instance

A VMware Cloud on Alibaba Cloud instance is a virtual machine (VM) that contains the basic computing components of a server, such as CPU, memory, operating system, network, and disks.

Prerequisites

- The region where VMware Cloud on Alibaba Cloud is deployed is managed. For more information, see [Add a VMware node](#).
- The region where VMware Cloud on Alibaba Cloud is deployed is bound to an organization. For more information, see [Bind a VMware Cloud on Alibaba Cloud region](#).

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Click **Create Instance** in the upper-right corner.
5. Configure parameters listed in the following table to create an instance.

| Section | Parameter | Required | Description |
|----------------|----------------------------------|----------|--|
| Basic Settings | Organization | Yes | The organization in which to create the instance. |
| | Resource Set | Yes | The resource set in which to create the instance. |
| Region | Region | Yes | The region in which to create the instance. |
| | Zone | Yes | The zone in which to create the instance. |
| | VPC | Yes | The VPC in which to create the instance. |
| | vSwitch | Yes | Select the vSwitch to which the instance belongs. The vSwitch corresponds to the port group of a VMware ESXi host or a distributed switch, and maps to the VLAN of a physical switch. |
| | Private IP Address | Yes | The private IPv4 address of the instance. The private IPv4 address must be within the CIDR block of the vSwitch. |
| | Private Subnet Mask | Yes | The private subnet mask. Example: 255.255.255.0. The specified subnet mask must be within the CIDR block of the selected vSwitch. |
| | Private IP Address of Gateway | Yes | The private IP address of the gateway. Example: 192.168.100.1. The IP address must be within the CIDR block of the selected vSwitch. |
| | Private IP Address of DNS Server | No | The private IP address of the DNS server. Example: 114.114.114.114. The IP address must be within the CIDR block of the selected vSwitch. |
| | | | |

| Section | Parameter | Required | Description |
|----------|------------------|----------|--|
| Instance | Instance Family | No | The instance family of the instance. Valid values: <ul style="list-style-type: none"> Memory Optimized Compute Optimized General Purpose |
| | Instance Type | Yes | The instance type of the instance. You can specify the vCPUs and memory. |
| Image | Image Type | No | The type of the image. Default value: Public Image . |
| | Public image | Yes | The public image of the instance. |
| | System Disk (GB) | No | <p>The system disk to which the operating system is installed.</p> <p>You can configure different storage types for the disk. Valid values:</p> <ul style="list-style-type: none"> Shared Storage: All: The system selects an available shared storage. We recommend that you select this type. Shared Storage: storageA: The storage named storageA of the VMware Cloud on Alibaba Cloud instance is used. Administrators must make sure the storage is appropriate. If the storage capacity is insufficient, the instance fails to be created. |

| Storage Section | Parameter | Required | Description |
|-----------------|------------------|----------|---|
| | Data Disk (GB) | No | <p>You can also add data disks after the instance is created.</p> <p>You can configure different storage types for the disk. Valid values:</p> <ul style="list-style-type: none"> ◦ Shared Storage: All: The system selects an available shared storage. We recommend that you select this type. ◦ Shared Storage: storageA: The storage named storageA of the VMware Cloud on Alibaba Cloud instance is used. Administrators must make sure the storage is appropriate. If the storage capacity is insufficient, the instance fails to be created. <p>You must also specify the provision type when you create the instance. Valid values:</p> <ul style="list-style-type: none"> ◦ Thin Provision: Storage space increases with the use of the disk. ◦ Thick Provision Lazy Zeroed: Storage space is equal to the size of the disk and does not increase. The disk is formatted when data is written. ◦ Thick Provision Eager Zeroed: Storage space is equal to the size of the disk and does not increase. The storage of the disk is immediately formatted when the disk is created. |
| Password | Password Setting | No | Select Set after Purchase . |
| Instance Name | Instance Name | Yes | <p>The name of the instance.</p> <p>The name must be 2 to 128 characters in length and can contain letters, periods (.), underscores (_), hyphens (-), and colons (:). It must start with a letter and cannot start with http:// or https://.</p> |

6. Click **Submit**.

1.7.4.2. View instance information

You can view the list of created instances as well as details of individual instances, such as their basic configurations, disks, and elastic network interfaces (ENIs).

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console.](#)
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
You can view the list of VMware Cloud on Alibaba Cloud instances that are deployed in the current region.
4. Use one of the following methods to go to the details page of an instance:
 - In the **Instance ID/Name** column, click the instance ID.
 - Click **Manage** in the **Actions** column corresponding to the instance.
 - Choose **More > Show Details** in the **Actions** column corresponding to the instance.

1.7.4.3. Modify an instance

You can modify the name and description of a created VMware Cloud on Alibaba Cloud instance.

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console.](#)
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Find the instance that you want to modify and choose **More > Modify** in the **Actions** column.
5. Modify the name and description of the instance.
6. Click **OK**.

1.7.4.4. Remotely connect to an instance

You can remotely connect to and manage added VMware Cloud on Alibaba Cloud instances.

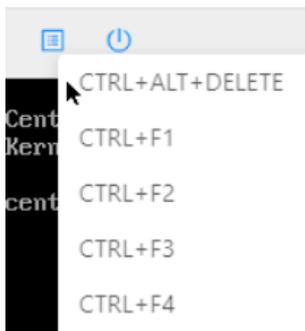
Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console.](#)
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Find the instance that you want to manage and click **Remote Connection** in the **Actions** column.
5. Enter the username and password.
 - For a Linux instance, enter the username *root* and the logon password.

Note

When you log on to the Linux instance, the password is not displayed as you enter it. Press the Enter key after you enter the password.

- For a Windows instance, to use a key combination such as Ctrl+Alt+Delete, click the List icon in the upper-right corner of the page and select the corresponding composite key from the drop-down list.



Enter the username and password, and click the Log On icon.

1.7.4.5. Stop an instance

You can stop VMware Cloud on Alibaba Cloud instances that are not in use. The stop operation interrupts services that are running on the instances. Exercise caution when you perform this operation.

Prerequisites

The instance is in the **Running** state.

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console.](#)
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Use one of the following methods to stop instances:
 - To stop a single instance, find the instance and choose **More > Instance Status > Stop** in the **Actions** column.
 - To stop one or more instances at a time, select the instances and click **Stop** in the lower part of the **Instances** page.
5. Click **OK**.

Execution results

When the instance is being stopped, its state in the **Status** column changes from **Running** to **Stopping**. After the instance is stopped, its state changes to **Stopped**.

1.7.4.6. Start an instance

You can start a stopped instance.

Prerequisites

The instance is in the **Stopped** state.

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console.](#)
2. In the left-side navigation pane, click **Instances**.

3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Use one of the following methods to start instances:
 - To start a single instance, find the instance and choose **More > Instance Status > Start** in the **Actions** column.
 - To start one or more instances at a time, select the instances and click **Start** in the lower part of the Instances page.
5. Click **OK**.

Execution results

When the instance is being started, its state in the **Status** column changes from **Stopped** to **Starting**. After the instance is started, its state changes to **Running**.

1.7.4.7. Restart an instance

After you change the logon password of an instance or install system updates, you must restart the instance. The restart operation stops the instances for a short period of time and interrupts services that are running on the instance. Exercise caution when you perform this operation.

Prerequisites

The instance is in the **Running** state.

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console.](#)
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Use one of the following methods to restart the instance:
 - To restart a single instance, find the instance and choose **More > Instance Status > Restart** in the **Actions** column.
 - To restart one or more instances at a time, select the instances and click **Restart** in the lower part of the Instances page.
5. In the Restart Instance dialog box, select a restart mode.
 - **Restart**: restarts the instance normally.
 - **Force Restart**: forces the instance to restart. This may result in the loss of unsaved data.
6. Click **OK**.

1.7.4.8. Delete an instance

You can delete instances that are no longer needed to release their resources. Deleted instances cannot be recovered. We recommend that you back up data before you delete an instance. If data disks are released with the instances, the disk data cannot be recovered.

Prerequisites

The instance is in the **Stopped** state.

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console.](#)

2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Select the instance and click **Delete** in the lower part of the Instances page.
5. Click **OK**.

1.7.4.9. Change the instance type of an instance

You can change the instance types of instances to suit your business needs. This eliminates the need to create VMware Cloud on Alibaba Cloud instances.

Prerequisites

The instance is in the **Stopped** state.

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the instance whose instance type you want to change and click **Upgrade/Downgrade** in the **Actions** column.
5. On the page that appears, select a new instance type and click **Submit**.
The page that appears shows the instance types available for selection.
6. Start the instance to make the new instance type take effect.
For more information, see [Start an instance](#).

1.7.5. Images

1.7.5.1. Create a custom image

You can create a custom image and use it to create identical instances or replace the system disks of existing instances. This way, you can configure many instances that have identical operating systems and data environments.

Create a custom image from an instance

You can create a custom image from an instance to replicate the data of all system and data disks on the instance.

Note

To avoid data security risks, we recommend that you delete sensitive data from an instance before you use the instance to create a custom image.

1. [Log on to the VMware Cloud on Alibaba Cloud console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Find the instance from which you want to create a custom image and choose **More > Create Custom Image** in the **Actions** column.
5. Set the name, sharing scope, and description for the custom image, and click **OK**.

The name must be 2 to 128 characters in length and can contain letters, digits, periods (.), underscores (_), hyphens (-), and colons (:). It cannot start with a special character or digit.

You can set the sharing scope to the permission scope of the image.

The description must be 2 to 256 characters in length and cannot start with `http://` or `https://`.

1.7.5.2. View images

You can view the list of created images.

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console.](#)
2. In the left-side navigation pane, choose **Images > Images**.
3. In the top navigation bar, move the pointer over **Region** and select the region where the image is created.
4. Select a filter option, enter the corresponding information in the search box, and then click **Search**.

You can select multiple filter options to narrow down search results.

| Parameter | Description |
|------------|--|
| Image Name | The image name used to search for the image. |
| Image ID | The image ID used to search for the image. |

1.7.6. Snapshots

1.7.6.1. Create a snapshot

You can manually create a snapshot for a disk to back up disk data.

Prerequisites

- The instance to which the disk is attached is in the **Running** or **Stopped** state.
- The disk is in the **Running** state.

Background information

A snapshot of a disk can be used to roll back data of the disk.

When you create a snapshot, take note of the following items:

- For each disk, the first snapshot taken is a full snapshot and subsequent snapshots are incremental snapshots. It takes longer to create the first full snapshot than it does subsequent incremental snapshots. The amount of taken time depends on the amount of data that has been changed since the previous snapshot. The more data that has been changed, the longer it takes to create an incremental snapshot.
- Avoid creating snapshots during peak hours.

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console.](#)
2. In the left-side navigation pane, click **Instances**.

3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Click the instance from which you want to create a snapshot. On the page that appears, click the **Snapshots** tab.
5. Click **Create and Bind Snapshot**.
6. Set the name, type, and description for the snapshot, and click **Submit**.

| Parameter | Description |
|----------------------|--|
| Snapshot Name | The name of the snapshot. |
| Snapshot Type | The snapshot type. Valid values: <ul style="list-style-type: none"> ◦ Disk Snapshot: stores information in disks. ◦ Memory Snapshot: stores information in memory. |
| Snapshot Description | The description of the snapshot. |

1.7.6.2. Delete a snapshot

You can delete a snapshot that is no longer needed. After the snapshot is deleted, it cannot be recovered.

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Find the instance whose snapshot is to be deleted and click the **Snapshots** tab.
5. Use one of the following methods to delete the snapshot:
 - To delete a single snapshot, find the snapshot and click **Delete** in the **Actions** column.
 - To delete one or more snapshots at a time, select the snapshots and click **Delete** in the lower part of the **Snapshots** tab.
6. Click **OK**.

1.7.6.3. View snapshots

You can view the list of created snapshots.

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Find the instance in which you want to view snapshots and click the **Snapshots** tab.
5. Select a filter option, enter the corresponding information in the search box, and then click **Search**.

You can select multiple filter options to narrow down search results.

| Parameter | Description |
|---------------|--|
| Snapshot Name | The snapshot name used to search for the snapshot. |
| Snapshot ID | The snapshot ID used to search for the snapshot. |

1.7.7. Disks

1.7.7.1. Create a disk

To increase the storage space of VMware Cloud on Alibaba Cloud instances, you can create standalone data disks and then attach them to the instances. This topic describes how to create an empty data disk. You cannot create standalone system disks.

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console.](#)
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Find the instance for which you want to create a disk and click the **Disks** tab.
5. Click **Create and Attach Disk**.
6. Configure parameters listed in the following table to create a disk.

| Section | Parameter | Required | Description |
|---------|----------------|----------|---------------------------------------|
| Region | Zone | Yes | The zone in which to create the disk. |
| | Specifications | Yes | The disk category and the disk size. |

| Section | Parameter | Required | Description |
|----------------|----------------|----------|---|
| Basic Settings | Provision Type | Yes | The provision type. Valid values: <ul style="list-style-type: none"> ◦ Thin Provision: Storage space increases with the use of the disk. ◦ Thick Provision Lazy Zeroed: Storage space is equal to the size of the disk and does not increase. The disk is formatted when data is written. ◦ Thick Provision Eager Zeroed: Storage space is equal to the size of the disk and does not increase. The storage of the disk is immediately formatted when the disk is created. |

7. Click **Submit**.

Execution results

The created disk is displayed in the disk list and in the **Running** state.

1.7.7.2. View disks

You can view the list of created disks and the details of individual disks.

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Click the instance for which you want to view disks. On the page that appears, click the **Disks** tab.
5. Select a filter option from the drop-down list, enter the relevant information in the search box, and then click **Search**.

You can select multiple filter options to narrow down search results.

| Parameter | Description |
|-----------|--|
| Disk Name | The disk name used to search for the disk. |
| Disk ID | The disk ID used to search for the disk. |

| Parameter | Description |
|-----------------|--|
| Disk Properties | The disk type used to search for disks of that type. Valid values: <ul style="list-style-type: none"> All System Disk Data Disk |

1.7.7.3. Detach a data disk

You can detach data disks. System disks cannot be detached.

Procedure

Warning

Resources are released after disks are detached. Make sure that the data of a disk is backed up before you detach it.

1. [Log on to the VMware Cloud on Alibaba Cloud console.](#)
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Click the instance from which you want to detach a data disk. On the page that appears, click the **Disks** tab.
5. Find the data disk that you want to detach and choose **More > Detach** in the **Actions** column.
6. Click **OK**.

1.7.8. ENIs

1.7.8.1. Create an ENI

You can create and bind elastic network interfaces (ENIs) to VMware Cloud on Alibaba Cloud instances.

Prerequisites

A virtual private cloud (VPC) and a vSwitch are created. For more information, see [Create a VPC](#) and [Create a vSwitch](#) in *Apsara Stack VPC User Guide*.

Background information

ENIs are classified into primary and secondary ENIs.

A primary ENI is created by default when an instance is created in a VPC. This primary ENI has the same lifecycle as the instance and cannot be unbound from the instance.

ENIs that are separately created are secondary ENIs. This topic describes how to create a secondary ENI.

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console.](#)
2. In the left-side navigation pane, click **Instances**.

3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Click the instance for which you want to create an ENI. On the page that appears, click the **ENIs** tab.
5. Click **Create and Bind ENI**.
6. Configure parameters listed in the following table to create an ENI.

| Section | Parameter | Required | Description |
|----------------|--------------|----------|---|
| Region | Organization | Yes | The organization in which to create the ENI. |
| | Resource Set | Yes | The resource set in which to create the ENI. |
| | Region | Yes | The region in which to create the ENI. |
| | Zone | Yes | The zone in which to create the ENI. |
| Basic Settings | VPC | Yes | <p>The VPC in which to create the ENI. The secondary ENI can be bound only to an instance in the same VPC.</p> <div style="background-color: #e0f2f7; padding: 5px;"> <p> Note After an ENI is created, you cannot change its VPC.</p> </div> |
| | vSwitch | Yes | <p>The vSwitch in which to create the ENI. The secondary ENI can be bound only to an instance in the same VPC. Select a vSwitch that is deployed within the same zone as the instance to which the ENI is bound. The vSwitch of the ENI can be different from that of the instance.</p> <div style="background-color: #e0f2f7; padding: 5px;"> <p> Note After an ENI is created, you cannot change its vSwitch.</p> </div> |

7. Click **Submit**.

Execution results

The created ENI is displayed on the ENIs page and is in the **Bound** state.

1.7.8.2. View ENIs

You can view the list of created elastic network interfaces (ENIs).

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console.](#)
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Click the instance for which you want to view ENIs. On the page that appears, click the **ENIs** tab.
5. Select a filter option, enter the corresponding information in the search box, and then click **Search**.

You can select multiple filter options to narrow down search results.

| Parameter | Description |
|------------|--|
| ENI Name | The ENI name used to search for the ENI. |
| ENI ID | The ENI ID used to search for the ENI. |
| vSwitch ID | The vSwitch ID used to search for the ENIs that are associated with the vSwitch. |

1.7.8.3. Delete an ENI

You can delete secondary elastic network interfaces (ENIs) that are no longer needed.

Background information

Only secondary ENIs can be deleted. Primary ENIs share the same lifecycle as instances and cannot be deleted.

Procedure

1. [Log on to the VMware Cloud on Alibaba Cloud console.](#)
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, move the pointer over **Region** and select the region where VMware Cloud on Alibaba Cloud is deployed.
4. Find the instance whose secondary ENI is to be deleted and click the **ENIs** tab.
5. Find the secondary ENI and click **Delete** in the **Actions** column.
6. Click **OK**.

1.8. Enterprise

1.8.1. Organizations

1.8.1.1. Create an organization

You can create organizations to store resource sets and their resources.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.

3. In the left-side navigation pane of the Enterprise page, click **Organizations**.
4. In the organization navigation tree, click a parent organization. In the Current Organization section, click **Add Organization**.
5. In the Add Organization dialog box, enter an organization name and click **OK**.
6. (Optional) If the primary node has enabled the multi-cloud management feature, the **Synchronously Create Multi-cloud Organization** check box is displayed after you click **OK**.
Select **Synchronously Create Multi-cloud Organization**, select clouds where the organization is synchronously created, and then click **OK**.

1.8.1.2. Query an organization

You can query an organization by name to view its resource sets, users, and user groups.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Organizations**.
4. In the search box below **Organizations**, enter an organization name.
You can view the information about the corresponding organization.

1.8.1.3. View organization information

You can view information about an organization on the Organizations page.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Organizations**.
4. On the **Organizations** page, click an organization in the organization list.
5. In the right-side area, view the organization information.
 - In the **Resource Sets** section, you can view information such as the name, creation time, and creator of each resource set in the organization. Click the name of a resource set to view its details.
 - In the **Users** section, you can view information such as the name, status, and role of each user in the organization. Click a username to view the user details.
 - In the **User Groups** section, you can view the name, organization, role, members, and creation time of each user group in the organization.

1.8.1.4. Modify the name of an organization

Users that have operation permissions on an organization can modify the name of the organization.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#).
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Organizations**.
4. In the organization navigation tree, click an organization name.
5. In the Current Organization section, click **Edit Organization**.
6. In the Edit Organization dialog box, modify the organization name.

If you want to synchronously modify the name of a multi-cloud organization in other clouds, select **Synchronously Modify Multi-cloud Organization**.

7. Click **OK**.

1.8.1.5. Change organization ownership

Users that have operation permissions on organizations can change the ownership of organizations.

Prerequisites

- Make sure that each organization under the organization that you want to change the ownership has a unique name.
- The ownership of an organization cannot be changed cross level-1 organizations.

Context

Users can change the ownership of an organization cross parent organizations. This way, the ownership of subordinate organizations, users, and resources are also changed in a cascading manner.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Resource Sets**.
4. On the **Change Ownership** page, select an organization and click **Change Organization** in the upper-right corner.
5. In the **Change Organization** dialog box, select the destination organization and click **OK** to change the ownership of the organization along with that of its resources sets and users.

1.8.1.6. Obtain the AccessKey pair of an organization

An AccessKey pair consists of an AccessKey ID and an AccessKey secret. The AccessKey pair is used to implement symmetric encryption to verify the identity of the requester. The AccessKey ID is used to identify a user. The AccessKey secret is used to encrypt the signature string. This topic describes how to obtain the AccessKey pair of an organization.

Prerequisites

Only operations administrators and level-1 organization administrators can obtain the AccessKey pair of an organization.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#).
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, click **Organizations**.
4. In the organization navigation tree, click an organization name.
5. In the Current Organization section, click **Obtain AccessKey Pair**.
6. In the AccessKey message, view the AccessKey pair of the organization.

1.8.1.7. Delete an organization

Administrators can delete organizations that are no longer needed.

Prerequisites

Before you delete an organization, make sure that the organization does not contain users, resource sets, or subordinate organizations. Otherwise, the organization cannot be deleted.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Organizations**.
4. In the organization navigation tree, click an organization name. In the **Current Organization** section, click **Delete Organization**.
5. In the Confirm message, click **OK**.

1.8.2. Resource sets

1.8.2.1. Create a resource set

You must create a resource set before you apply for resources.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Resource Sets**.
4. In the upper-left corner of the **Resource Sets** page, click **Create Resource Set**.
5. In the **Create Resource Set** dialog box, set **Name** and **Organization**.
6. Click **OK**.

1.8.2.2. View the details of a resource set

When you want to use a cloud resource in your organization, you can view the details of the resource set that contains the resource, including all resource instances and users of the resource set.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Resource Sets**.
4. Select an **organization** from the drop-down list, or enter a **resource set** name in the search bar, and then click **Search**.
5. Click the name of the target **resource set**.
6. On the **Resource Set Details** page, click the **Resources** and **Members** tabs to view information about all resource instances and users of the resource set.
7. On the **Resources** tab, click the number of a service to go to the instance list page of the service. The list is automatically filtered and displayed based on the organization and resource set.

1.8.2.3. Modify the name of a resource set

An administrator can modify the name of a resource set to keep it up-to-date.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as an administrator.

2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Resource Sets**.
4. Click **More** in the **Actions** column corresponding to a resource set, and choose **Edit Name** from the shortcut menu.
5. In the dialog box that appears, enter the new name.
6. Click **OK**.

1.8.2.4. Add a member to a resource set

You can add a member to a resource set so that the member can use the resources in the resource set.

Prerequisites

Before adding a member, make sure that the following prerequisites are met:

- A resource set is created. For more information, see [Create a resource set](#).
- A user is created. For more information, see [Create a user](#).

Context

Members of a resource set have the permissions to use resources in the resource set.

Deleting resources from a resource set does not affect the members of the resource set. Similarly, deleting members from a resource set does not affect the resources in the resource set.

You can delete a member that is no longer in use in a resource set. After the member is deleted, it will no longer be able to access the resource set.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Resource Sets**.
4. Click **More** in the **Actions** column corresponding to a resource set, and choose **Add Member** from the shortcut menu.
5. In the dialog box that appears, select a username.
6. Click **OK**.

1.8.2.5. Add or remove a user group of a resource set

You can add or remove a user group of a resource set to manage user group access to resources in the resource set.

Prerequisites

- A resource set is created. For more information, see [Create a resource set](#).
- A user group is created. For more information, see [Create a user group](#).

Context

User groups in a resource set have the permissions to use resources in the resource set.

Deleting resources from a resource set does not affect user groups of the resource set. Similarly, deleting user groups from a resource set does not affect the resources in the resource set.

You can delete a user group that is no longer in use in a resource set. After the user group is deleted, it will no longer be able to access the resource set.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Resource Sets**.
4. Click **More** in the **Actions** column corresponding to the target resource set.
5. Add or remove a user group.
 - Select **Add User Group**. In the dialog box that appears, select a user group. Click **OK** to add the user group.
 - Select **Delete User Group**. In the dialog box that appears, select a user group. Click **OK** to remove the user group.

1.8.2.6. Delete a resource set

You can delete resource sets that are not needed as an administrator.

Prerequisites

Ensure that the resource set to be deleted does not contain resources, users, or user groups.

 **Notice** A resource set cannot be deleted if it contains resources, users, or user groups.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Resource Sets**.
4. Click **More** in the **Actions** column corresponding to the target resource set, and select **Delete**.
5. In the message that appears, click **OK**.

1.8.3. Roles

1.8.3.1. Create a custom role

You can create custom roles in the Apsara Uni-manager Management Console to more efficiently grant permissions to users so that different personnel can work with different features.

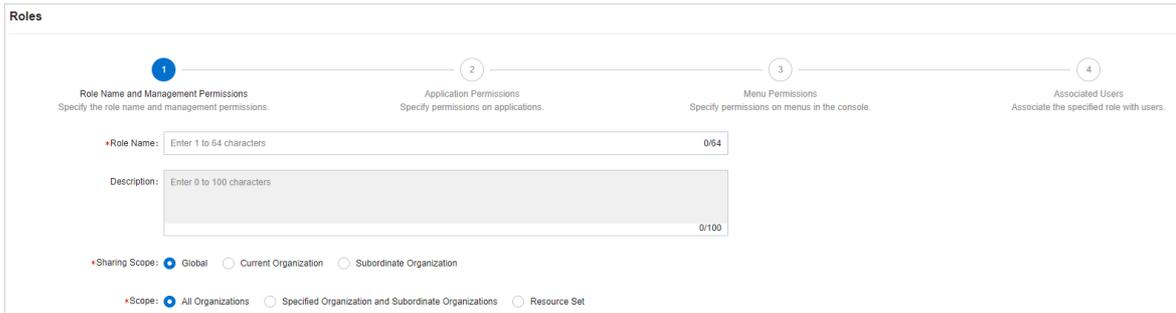
Context

A role is a set of access permissions. Each role has a range of permissions. A user can have multiple roles, which means that the user is granted all of the permissions defined for each role. A role can be used to grant the same set of permissions to a group of users.

The total number of custom and default roles cannot exceed 20.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. In the upper-right corner of the page, click **Create Custom Role**.
5. On the **Roles** page, set the role name and management permissions.



The following table describes the role parameters.

Role parameters

| Parameter | Description |
|----------------------|---|
| Role Name | The name of the RAM role. The name can be up to 15 characters in length and can contain only letters and digits. |
| Description | Optional. The description of the role. The description can be up to 100 characters in length and can contain letters, digits, commas (,), semicolons (;), and underscores (_). |
| Sharing Scope | <ul style="list-style-type: none"> ◦ Global The role is visible and valid to all organizations involved. The default value is Global. ◦ Current Organization The role is visible and valid to the organization to which the user belongs. ◦ Subordinate Organization The role is visible and valid to the organization to which the user belongs and its subordinate organizations. |
| Scope | <ul style="list-style-type: none"> ◦ All Organizations The permissions apply to all organizations involved. ◦ Specified Organization and Subordinate Organizations The permissions apply to the organization to which the user belongs and its subordinate organizations. ◦ Resource Sets The permissions apply to the resource sets that are assigned to the user. |

6. Select the operation permissions that this role has and click **Next**.
7. In the **Application Permissions** step, select the operation permissions that this role has on the cloud services, and click **Next**.
8. In the **Menu Permissions** step, select the operation permissions that this role has on the menus and the homepage template corresponding to the role, and click **Create Role**.
9. In the **Associated Users** step, select the users associated with the role from the drop-down list. The associated users are granted the permissions of the role.

1.8.3.2. View the details of a role

If you are uncertain about the specific permissions of a role, you can go to the **Roles** page to view the role permissions.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. Click the name of the role that you want to view. On the **Roles** page, view the information of the role.

1.8.3.3. Modify custom role information

You can modify the name and permissions of a custom role as an administrator.

Context

Information about preset roles cannot be modified.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. In the role name list, click **More** in the **Actions** column corresponding to the target custom role, and select **Modify**.
5. On the **Roles** page, modify the custom role name, permissions, and associated users or user groups.
 - Modify role name: Enter a new role name in the **Role Name** field.
 - Modify permissions: Click the **Management Permissions**, **Application Permissions**, or **Menu Permissions** tab, select or clear related permissions from the corresponding tab, and then click **Update**.
 - Bind a user to a role: Click the **Associated Users** tab and select a user from the **Select one or more users** drop-down list to add the user. To unbind the user from the role, click **Remove** in the **Actions** column.
 - Manage user groups: Click the **User Groups** tab, click **Add User Group**, select a user group from the drop-down list, and then click **OK** to bind the user group. To unbind the user group from the role, click **Remove** in the **Actions** column.

1.8.3.4. Copy a role

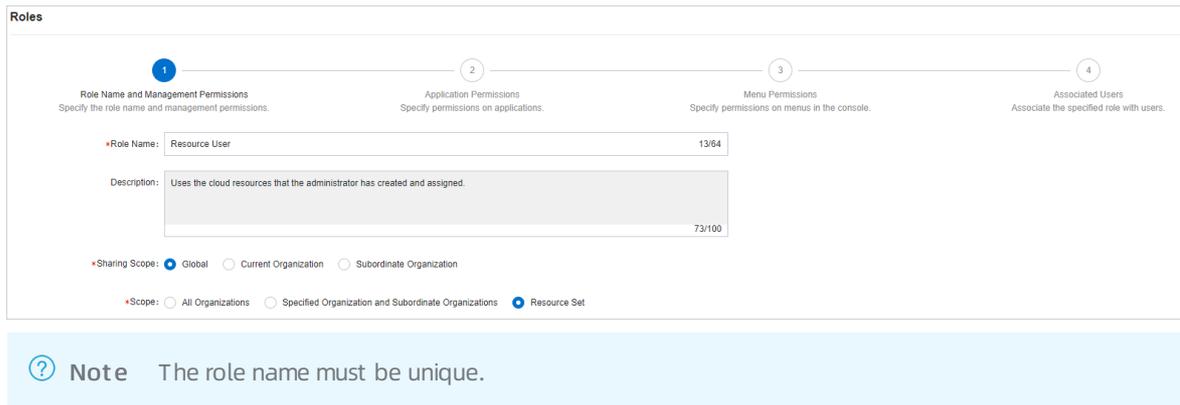
You can copy a preset role or a custom role to create a role that has the same permissions.

Context

Operations on the **Roles** page are the same as those for creating a custom role. You can add, modify, and remove the role permissions in the copied role. By default, if you do not modify the role permissions, the sharing scope, management permissions, application permissions, menu permissions, and associated users of the copied role are all the same as those of the source role.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. In the role list, choose **More > Copy** in the **Actions** column corresponding to a role.
5. On the **Roles** page, set the new role name, sharing scope, and management permissions.



Roles

1 Role Name and Management Permissions
Specify the role name and management permissions.

2 Application Permissions
Specify permissions on applications.

3 Menu Permissions
Specify permissions on menus in the console.

4 Associated Users
Associate the specified role with users.

*Role Name: Resource User 13/64

Description: Uses the cloud resources that the administrator has created and assigned. 73/100

*Sharing Scope: Global Current Organization Subordinate Organization

*Scope: All Organizations Specified Organization and Subordinate Organizations Resource Set

Note The role name must be unique.

6. Select the operation permissions that this role has and click **Next**.
7. In the **Application Permissions** step, select the operation permissions that this role has on the cloud services and click **Next**.
8. In the **Menu Permissions** step, select the operation permissions that this role has on the menus and click **Create Role**.
9. In the **Associated Users** step, select the users that are associated with the role from the drop-down list. The associated users are granted the permissions of the role.

1.8.3.5. Disable a role

When you disable a role, the permissions of the role are disabled.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. In the role list, click **More** in the **Actions** column corresponding to a role and choose **Disable** from the shortcut menu.

1.8.3.6. Enable a role

When you enable a disabled role, the permissions of the role are restored.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. In the role list, click **More** in the **Actions** column corresponding to a disabled role and choose **Enable** from the shortcut menu.

1.8.3.7. Delete a custom role

You can delete a custom role that is no longer needed.

Prerequisites

- Default or preset roles cannot be deleted.
- To delete a role, you must unbind all user groups from the role.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. Choose **More > Delete** in the **Actions** column corresponding to a role.
5. In the Confirm message, click **OK**.

1.8.4. Users

1.8.4.1. System users

1.8.4.1.1. Create a user

You can create a user and assign the user different roles as an administrator to meet different requirements for system access control.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. Use one of the following methods to open the Create User window:
 - In the left-side navigation pane of the **Enterprise** page, click **Organizations**. In the **Users** section of the **Organizations** page, click **Create User**.
 - In the left-side navigation pane of the **Enterprise** page, click **Users**. On the **System Users** tab of the **Users** page, click **Create**.
4. In the Create User dialog box, configure the parameters.

| Parameter | Description |
|--------------|---|
| Username | The Apsara Stack account name of the user. The name must be 1 to 64 characters in length and can contain letters, digits, hyphens (-), underscores (_), periods (.), and at signs (@). It must start with a letter or digit. |
| Display Name | The display name of the user. The name must be 1 to 128 characters in length and can contain letters, digits, hyphens (-), underscores (_), periods (.), and at signs (@). |
| Roles | The role to be assigned to the user. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 5px;"> ? Note You can enter role names in the field. Fuzzy match is supported. </div> |
| Organization | The organization to which the user belongs. |
| Logon Policy | The logon policy that restricts the logon time and IP addresses of the user. The default policy is automatically bound to new users. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 5px;"> ? Note The default policy does not restrict the time period and IP addresses for users to log on. To restrict the allowed logon time and IP addresses of a user, you can modify the logon policy of the user or create a logon policy for the user. For more information, see Create a logon policy. </div> |

| Parameter | Description |
|-------------------------|---|
| Mobile Number | <p>The mobile phone number of the user. This mobile number is used by the system to notify users of resource application and usage. Make sure that this mobile number is valid.</p> <p> Note If the user mobile number changes, update it on the system in a timely manner.</p> |
| Landline Number | <p>Optional. The landline number of the user. The landline number must be 4 to 20 characters in length and can contain only digits and hyphens (-).</p> |
| Email | <p>The email address of the user. Emails about the usage and requests for resources are sent to this email address. Make sure that this email address is valid.</p> <p> Note If the user email address changes, update it on the system in a timely manner.</p> |
| DingTalk Key | <p>The key of the chatbot for the DingTalk group where the user is a member. For more information about how to configure the key, see DingTalk development documentation.</p> |
| Notify User by Email | <p>After this option is selected, the Apsara Uni-manager Management Console informs the user configured as the alert contact by email whenever an alert is generated.</p> <p> Note You must configure an email server to receive an email each time an alert is triggered. For more information, contact on-site O&M engineers.</p> |
| Notify User by DingTalk | <p>After this option is selected, the Apsara Uni-manager Management Console informs the user configured as the alert contact by DingTalk whenever an alert is generated.</p> |

5. Click **OK**.

1.8.4.1.2. Query a user

You can view user information such as name, organization, mobile number, email address, role, logon time, and initial password.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Users**.
4. Click the **System Users** tab.
5. Set **Username**, **Organization**, or **Role**, and then click **Search**.
6. Click **More** in the **Actions** column corresponding to a user, and choose **User Information** from the shortcut menu to view basic information about the user.

1.8.4.1.3. Modify user information

You can modify user information such as display name, mobile number, and email address to keep it up to date.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Users**.
4. Click the **System Users** tab.
5. Find the user that you want to modify and choose **More > Edit** in the **Actions** column.
6. In the **Modify User Information** dialog box, enter the relevant information and click **OK**.

1.8.4.1.4. Change user roles

You can add, change, and delete roles for a user.

Change user roles by using user management

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Users**.
4. Click the **System Users** tab.
5. Find the user that you want to modify and choose **More > Authorize** in the **Actions** column.
6. In the **Role** field, add, delete, or change user roles.
You can enter role names in the field. Fuzzy match is supported.
7. Click **OK**.

Change user roles by changing ownership

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Change Ownership**.
4. Click the  icon to the left of an organization and click **Users**.
5. In the **Users** section on the right, set **Logon Policy** and **Role** or **Username**, and click **Search** to query the user that you want to modify.
6. Find the user and click **Change** in the **Actions** column.
7. In the **Organization to Change** dialog box, select the destination or original organization and select the role to be added or removed from the **Assigned Roles** drop-down list.

Note

- If you change only roles without changing the organization, select the original organization.
- Blue role names are the roles that are selected, and black role names are the roles that are not selected.

8. Click **OK**.

1.8.4.1.5. Modify the information of a user group

On the **Users** page, you can view the user group information and modify the ownership of users in user groups.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.

2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, click **Users**.
4. Click the **System Users** tab, select the user that you want to modify, and then click **More** in the **Actions** column.
 - Select **Add to User Group**. In the dialog box that appears, select a user group and click **OK** to add the user to the user group.
 - Select **Remove from User Group**. In the dialog box that appears, select a user group and click **OK** to remove the user from the user group.

1.8.4.1.6. Modify a user logon policy

An administrator can modify a user logon policy to restrict the permitted logon time and IP addresses of the user.

Prerequisites

A new logon policy is created. For more information about how to create a logon policy, see [Create a logon policy](#).

Modify the logon policy of a single user

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Users**.
4. Click the **System Users** tab.
5. Find the user that you want to modify and choose **More > Logon Policy** in the Actions column.
6. In the **Assign Logon Policy** dialog box, select a logon policy and click **OK**.

Modify multiple user logon policies at a time

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Users**.
4. Click the **System Users** tab.
5. Select users that you want to modify.
6. In the upper-right corner of the page, click **Change Logon Policy**.
7. In the **Batch Assign Logon Policy** dialog box, select a logon policy and click **OK**.

1.8.4.1.7. View the initial password of a user

After a user is created, the system generates an initial password for the user.

Context

Organization administrators can view the initial passwords of all users in the organizations they manage.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. Use one of the following methods to view the initial password of a user on the **Enterprise** page:
 - In the left-side navigation pane, click **Users**. On the **System Users** tab of the **Users** page, select a username.
 - Click **View Initial Password** in the upper-right corner of the **Users** page to view the initial password.

- Choose **More > User Information** in the **Actions** column corresponding to the user. On the user information page, click **View Password** to view the initial password.
- In the left-side navigation pane, click **Organizations**. In the organization navigation tree on the **Organizations** page, click an organization name. In the **Users** section, click a username. On the user information page, click **View Password** to view the initial password.

1.8.4.1.8. Reset the password of a user

If users forget their logon passwords, the system administrator can reset the logon passwords for them.

Prerequisites

Only organization administrators can reset the password of a user.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. Use one of the following methods to go to the **User Information** page:
 - In the left-side navigation pane of the **Enterprise** page, click **Users**. On the **System Users** tab of the **Users** page, click a username.
 - In the left-side navigation pane of the **Enterprise** page, click **Organizations**. On the **Organizations** page, click a username in the **Users** section.
4. Click **Reset Password**.

After the password is reset, a message is displayed, which indicates that the password has been reset. If you want to view the initial password after password reset, click **View Password**.

1.8.4.1.9. Disable or enable a user account

You can disable a user account to prevent the user account from logging on to the Apsara Uni-manager Management Console. User accounts that are disabled must be re-enabled before they can be used to log on to the Apsara Uni-manager Management Console again.

Context

By default, user accounts are enabled when they are created.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Users**.
4. Click the **System Users** tab.
5. Perform the following operations on the current tab:
 - Select a user account whose **Status** is **Enabled**, choose **More > Disable** in the **Actions** column to disable the user account.
 - Select a user whose **Status** is **Disabled**, choose **More > Enable** in the **Actions** column to enable the user account.

1.8.4.1.10. Delete a user

You can delete a specific user as an administrator.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. On the Enterprise page, use one of the following methods to delete a user:
 - In the left-side navigation pane of the **Enterprise** page, click **Users**. On the page that appears, click the **System Users** tab. Click **More** in the **Actions** column corresponding to the target user, and select **Delete**.
 - In the left-side navigation pane of the **Enterprise** page, click **Organizations**. On the page that appears, find the **Users** section. Find the target user, click **More** in the **Actions** column, and then select **Delete**.
4. Click **OK**.

1.8.4.2. Historical users

1.8.4.2.1. Query historical users

You can check whether a user has been deleted and restore a user that has been deleted.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Users**.
4. Click the **Historical Users** tab.
5. Enter the username that you want to query in the **Username** search box.

 **Note** You can search for usernames by fuzzy match.

6. Click **Search**.

1.8.4.2.2. Restore historical users

An administrator can restore a deleted user account from the **Historical Users** tab.

Context

The basic information such as logon password of a restored user is the same as it was before the user was deleted, except for the organization and role.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Users**.
4. Click the **Historical Users** tab.
5. Find the user that you want to restore and click **Restore** in the **Actions** column.
6. In the **Restore User** dialog box, select an organization and a role.
7. Click **OK**.

1.8.5. Logon policies

1.8.5.1. Create a logon policy

To improve the security of the Apsara Uni-manager Management Console, you can create a logon policy as an administrator to control logon access based on the logon time and user IP address.

Context

Logon policies are used to control the time period and IP addresses for users to log on. After a user is bound to a logon policy, user logons are restricted based on the logon time and IP addresses specified in the policy.

A default policy without limits on logon time and IP addresses is automatically generated in the Apsara Uni-manager Management Console. The default policy cannot be deleted.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Logon Policies**.
4. In the upper-right corner of the page, click **Create**.
5. In the **Create Logon Policy** dialog box, set Name, Sharing Scope, Policy Properties, Time Period, and IP Address.

Parameters for creating a logon policy

| Parameter | Description |
|--------------------------|---|
| Name | The name of the logon policy. The name must be 2 to 50 characters in length and can contain only letters and digits. The name must be unique in the system. |
| Description | The description of the logon policy. |
| Sharing Scope | The scope in which the role is visible. <ul style="list-style-type: none"> ◦ Global: The role is globally visible. The default value is Global. ◦ Current Organization: The role is visible only in the current organization and is invisible in subordinate organizations. ◦ Subordinate Organization: The role is visible in the current organization and all its subordinate organizations. |
| Policy Properties | The authentication method of the logon policy. <ul style="list-style-type: none"> ◦ Whitelist: Logon is allowed if the parameter settings are met. ◦ Blacklist: Logon is denied if the parameter settings are met. |
| Time Period | The permitted logon time period. When this policy is configured, users can log on to the Apsara Uni-manager Management Console only during the configured period. Specify the time in minutes in a 24-hour clock. Example: <code>16:32</code> . <div style="background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note When the Policy Properties parameter is set to Whitelist, you can select No Time Limit.</p> </div> |
| IP Address | The permitted CIDR block. <ul style="list-style-type: none"> ◦ If the Policy Properties parameter is set to Whitelist, IP addresses within this CIDR block are allowed to log on to the Apsara Uni-manager Management Console. ◦ If the Policy Properties parameter is set to Blacklist, IP addresses within this CIDR block are not allowed to log on to the Apsara Uni-manager Management Console. <div style="background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note When the Policy Properties parameter is set to Whitelist, you can select No CIDR Block Limit.</p> </div> |

1.8.5.2. Query a logon policy

You can query the detailed information of a logon policy in the Apsara Uni-manager Management Console.

Context

When the Apsara Uni-manager Management Console provides services, it automatically generates a default policy without limits on the logon time and IP addresses.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Logon Policies**.
4. Enter the name of the policy that you want to view and click **Search**.
5. View the logon policy, including the permitted logon time and IP addresses.

1.8.5.3. Modify a logon policy

You can modify the policy name, policy properties, permitted logon time period, and IP addresses of a logon policy.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Logon Policies**.
4. Find the logon policy that you want to modify and choose **More > Modify** in the **Actions** column.
5. In the **Modify Logon Policy** dialog box, modify the logon policy information.
6. Click **OK**.

1.8.5.4. Disable a logon policy

You can disable logon policies that are no longer needed.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#).
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Logon Policies**.
4. Find the logon policy that you want to disable and choose **More > Disable** in the **Actions** column.

1.8.5.5. Enable a logon policy

You can re-enable disabled logon policies.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#).
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Logon Policies**.
4. Click **More** in the **Actions** column corresponding to a policy, and choose **Enable** from the short cut menu.

1.8.5.6. Delete a logon policy

You can delete logon policies that are no longer needed.

Prerequisites

The logon policy to be deleted is not bound to any users. If a logon policy is bound to a user, the logon policy cannot be deleted.

Context

 **Note** The default policy cannot be deleted.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Logon Policies**.

4. Click **More** in the **Actions** column corresponding to a policy, and choose **Delete** from the short cut menu.
5. In the message that appears, click **OK**.

1.8.6. User groups

1.8.6.1. Create a user group

You can create a user group in a selected organization and grant batch authorizations to users in the group.

Prerequisites

Before creating a user group, you must create an organization. For more information, see [Create an organization](#).

Context

Relationship between user groups and users:

- A user group can contain zero or more users.
- You can add users to user groups as needed.
- You can add a user to multiple user groups.

Relationship between user groups and organizations:

- A user group can only belong to a single organization.
- You can create multiple user groups in an organization.

Relationship between user groups and roles:

- A user group can only be bound to a single role.
- A role can be associated with multiple user groups.
- When a role is associated with a user group, the role permissions are automatically granted to users in the user group.

Relationship between user groups and resource sets:

- You can add zero or more user groups to a resource set.
- A user group can be added to multiple resource sets.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **User Groups**.
4. In the upper-right corner of the page, click **Create User Group**.
5. In the dialog box that appears, set **User Group Name** and **Organization**.

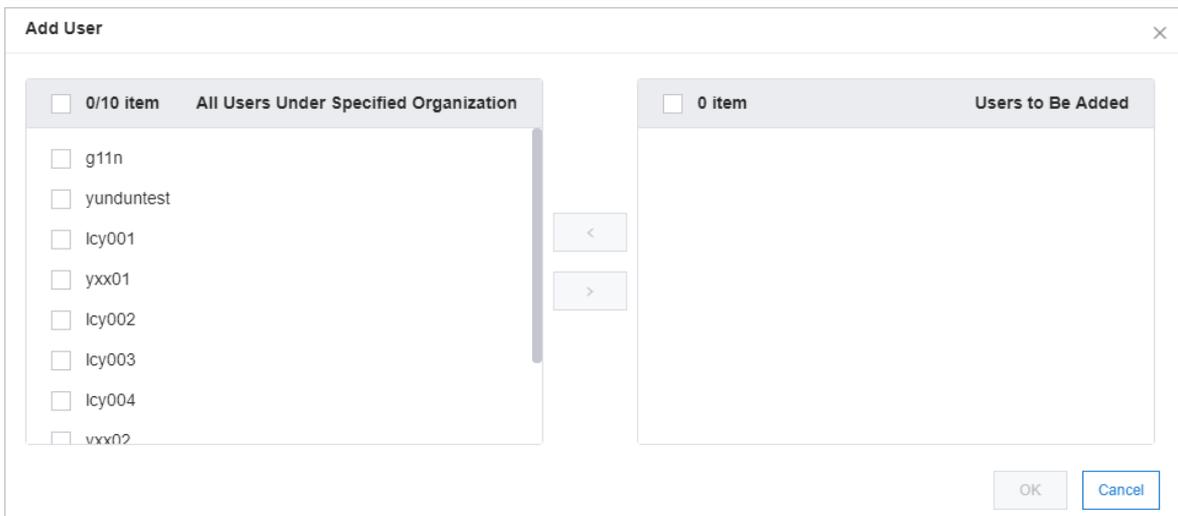
6. Click **OK**.

1.8.6.2. Add users to a user group

You can add users to a user group.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **User Groups**.
4. Click **Add User** in the **Actions** column corresponding to a user group.
5. Select the names of users to be added from the left list, and click the right arrow to move them to the right list.



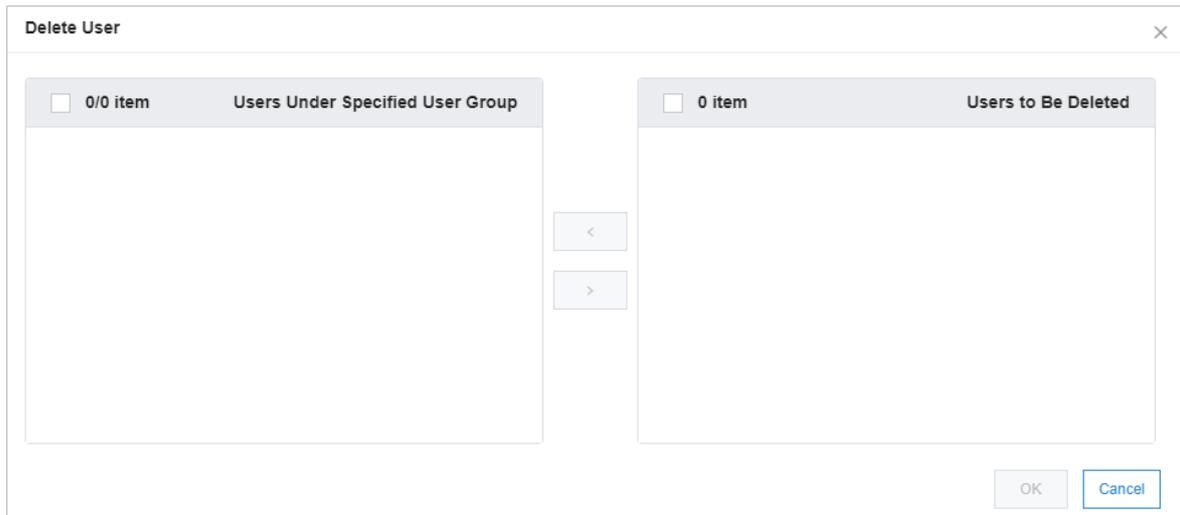
6. Click **OK**.

1.8.6.3. Delete users from a user group

You can delete users from a user group.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **User Groups**.
4. Click **Delete User** in the **Actions** column corresponding to a user group.
5. Select the names of users to be deleted from the **Users Under Specified User Group** list, and click the right arrow to move them to the **Users to Be Deleted** list.



6. Click **OK**.

1.8.6.4. Add a role

You can add a role to a user group and assign the role to all users in the group.

Context

 **Note** You can add only one role to a user group.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **User Groups**.
4. Click **Add Role** in the **Actions** column corresponding to a user group.
5. In the dialog box that appears, select a role.
6. Click **OK**.

1.8.6.5. Delete a role

You can delete existing roles.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **User Groups**.
4. Find the user group from which you want to delete a role and click **Delete Role** in the **Actions** column.
5. In the **Confirm** message, click **OK**.

1.8.6.6. Modify the name of a user group

You can modify the names of user groups.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **User Groups**.
4. Click **Edit User Group** in the **Actions** column corresponding to a user group.
5. In the dialog box that appears, enter the new name.
6. Click **OK**.

1.8.6.7. Delete a user group

You can delete user groups that are no longer needed.

Prerequisites

The user group to be deleted is unbound from all roles. If a user group is bound to a role, the user group cannot be deleted.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **User Groups**.
4. Find the user group that you want to delete and click **Delete User Group** in the **Actions** column.
5. In the **Confirm** message, click **OK**.

1.8.7. Resource pools

1.8.7.1. Update associations

You can deploy the Apsara Uni-manager Management Console in multiple regions. You can update the associations between organizations and regions.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Resource Pools**.
4. In the left-side organization navigation tree, click the name of the organization that you want to update.
5. In the corresponding region list, select the names of regions to be associated.
6. Click **Update Association**.

1.8.8. Change the ownership of an instance

You can change the ownership of an instance from one resource set to another.

Change the ownership of an instance

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Change Ownership**.
4. Click the  icon to the left of an organization and click a resource set.
5. In the resource list on the right side of the page, set a service type and a resource type, enter an instance ID,

and then click **Search** to query the instance.

6. Click **Change Ownership** in the **Actions** column corresponding to the instance to change the ownership of the instance to another resource set.
7. Click **Change Sharing Scope** in the **Actions** column corresponding to the instance to change the sharing scope of the instance.
 - **Current Organization and Subordinate Organizations:** The instance can be shared by the organization that contains the resource set to which the instance belongs and by subordinate organizations.
 - **Current Resource Set:** The instance can be shared by the resource set to which the instance belongs.
 - **Current Organization:** The instance can be shared by the organization that contains the resource set to which the instance belongs.
8. In the **Change Resource Set** dialog box, select a resource set and click **OK**.

1.8.9. Cloud instances

1.8.9.1. Manage Apsara Stack cloud instances

1.8.9.1.1. Export data of the current cloud

You can export the data of secondary Apsara Stack nodes to a configuration file. This can be used by the primary node to manage nodes in a centralized manner.

Procedure:

1. [Log on to the Apsara Uni-manager Management Console.](#)
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the Enterprise page, click **Cloud Instances**.
4. Click the **Apsara Stack Management** tab.
5. Click **Collect Data of Current Cloud** to collect the deployment information of the current cloud.
6. Click **Export** to export the information in the JSON format.

1.8.9.1.2. Add a secondary Apsara Stack node

You can add the configuration information of secondary Apsara Stack nodes to the multi-cloud configuration of the primary Apsara Stack node for centralized management.

Procedure:

1. [Log on to the Apsara Uni-manager Management Console.](#)
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the Enterprise page, click **Cloud Instances**.
4. Click the **Apsara Stack Management** tab.
5. Click **Import**.

- In the **Create Apsara Stack Secondary Node** dialog box, enter the configuration information of a secondary node and click **OK**.

| Parameter | Description |
|----------------------------|--|
| Cloud Instance Information | The configuration file of the secondary node. For more information, see Export data of the current cloud . |
| Secondary Node Name | The name of the secondary node. |
| Username | The username of the operations administrator that manages the secondary node. |
| Password | The password of the operations administrator that manages the secondary node. |
| Description | The description of the secondary node. |

| Parameter | Description |
|------------------|---|
| AccessKey ID | The AccessKey ID of the operations administrator that manages the secondary node. For more information, see View the AccessKey pair of your Apsara Stack tenant account . |
| AccessKey Secret | The AccessKey secret of the operations administrator that manages the secondary node. For more information, see View the AccessKey pair of your Apsara Stack tenant account . |

Notice

You must create an operations administrator account in the secondary node. This account is for dedicated use by the primary node and cannot be the default operations administrator account.

1.8.9.1.3. View managed cloud instances

You can use the multi-cloud management feature to view the details of all managed cloud instances.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#).
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the Enterprise page, click **Cloud Instances**.
4. Click the **Apsara Stack Management** tab.

You can view the name, description, cloud type, cloud role, and address of all managed cloud instances.

5. Enter a cloud instance name in the search box and click **Search** to search for the cloud instance.
6. Click **View Details** in the **Actions** column corresponding to the cloud instance.

In the Manage Cloud Instance message, you can view the version, Apsara Stack API (ASAPI) address, and region of the cloud.

1.8.9.1.4. Modify a cloud instance

If you want to change the information of a cloud instance for more efficient management, you can modify it in the Apsara Uni-manager Management Console.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#).
2. In the top navigation bar, click **Enterprise**. In the left-side navigation pane of the Enterprise page, click **Cloud Instances**.
3. Click the **Apsara Stack Management** tab.
4. Enter the name of a cloud instance that you want to modify in the search box and click **Search** to search for the cloud instance.
5. Click **Edit** in the **Actions** column corresponding to the cloud instance.

6. In the **Edit Cloud Instance** dialog box, set **Cloud Name**, **Username**, **Password**, **Description**, **AccessKey ID**, **AccessKey Secret**, **Longitude**, and **Latitude**, and click **OK**.

1.8.9.1.5. Manage cloud instances

You can manage Apsara Stack cloud instances to check whether they can be connected.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#).
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the Enterprise page, click **Cloud Instances**.
4. Click the **Apsara Stack Management** tab.
5. Enter a cloud instance name in the search box and click **Search** to search for the cloud instance.
6. Click **Manage** in the **Actions** column corresponding to the cloud instance.
7. In the **Manage Cloud Instance** dialog box, click **Test Connectivity**.

1.8.9.2. Manage VMware nodes

1.8.9.2.1. Add a VMware node

You can add the configuration information of VMware nodes to the Apsara Stack VMware management configuration for centralized management.

Prerequisites

- The configuration file of a VMware node is obtained from the deployment personnel.
- The VMware node is configured.

Procedure:

1. [Log on to the Apsara Uni-manager Management Console](#).
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the Enterprise page, click **Cloud Instances**.
4. Click the **VMware Management** tab.
5. Click **Create VMware Node**.

6. In the **Create VMware Node** dialog box, enter the configuration information of a VMware node and click **OK**.

| Parameter | Description |
|----------------------------|--|
| Cloud Instance Information | The configuration file of the VMware node. |
| Cloud Name | The name of the VMware node. |
| Cloud Description | The description of the VMware node. |
| AccessKey ID | The AccessKey ID in the configuration file of the VMware node. |
| AccessKey Secret | The AccessKey secret in the configuration file of the VMware node. |

1.8.9.2.2. Modify a VMware node

If you want to change the information of a VMware node for more efficient management, you can modify it in the Apsara Uni-manager Management Console.

Procedure

1. [Log on to the Apsara Uni-manager Management Console.](#)
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the Enterprise page, click **Cloud Instances**.
4. Click the **VMware Management** tab.

5. Enter the name of the VMware node that you want to modify in the search box and click **Search** to search for the VMware node.
6. Click **Edit** in the **Actions** column corresponding to the VMware node.
7. In the **Edit Cloud Instance** dialog box, set **Cloud Name**, **Cloud Description**, **AccessKey ID**, and **AccessKey Secret**, and click **OK**.

1.8.9.2.3. Test VMware node connectivity

You can manage VMware nodes to check whether they can be connected.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#).
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the Enterprise page, click **Cloud Instances**.
4. Click the **VMware Management** tab.
5. Enter a VMware node name in the search box and click **Search** to search for the VMware node.
6. Click **Manage** in the **Actions** column corresponding to the VMware node.
7. In the **Manage Cloud Instance** dialog box, click **Test Connectivity**.

1.8.10. Data permissions

1.8.10.1. Overview

Data permission management allows you to specify which users can access instances of a specific service, grant data access permissions to the users, and view and modify the data permissions in all the RAM policies attached to specified users.

Apsara Stack controls users and permissions by managing their visibility and operability in the Apsara Uni-manager Management Console. Many Apsara Stack cloud services are directly used by calling their API operations or SDKs instead of in the console. In this case, data access permissions must be controlled by RAM permission verification provided by the cloud services.

RAM policies are configured for such cloud service instances for access control. Automatic judgment is used when personnel are added to or removed from resource sets. However, this judgement method can affect performance and has a high error rate in complex scenarios. To solve this problem, the authorization of cloud services that require data access permissions is separately managed. Organization administrators can configure the data permissions granted to related personnel on the data authorization page.

1.8.10.2. Set the data permissions of resource instances

Organization administrators can set the data permissions of resource instances to allow or prohibit access to and operations on cloud services in the Apsara Uni-manager Management Console.

Prerequisites

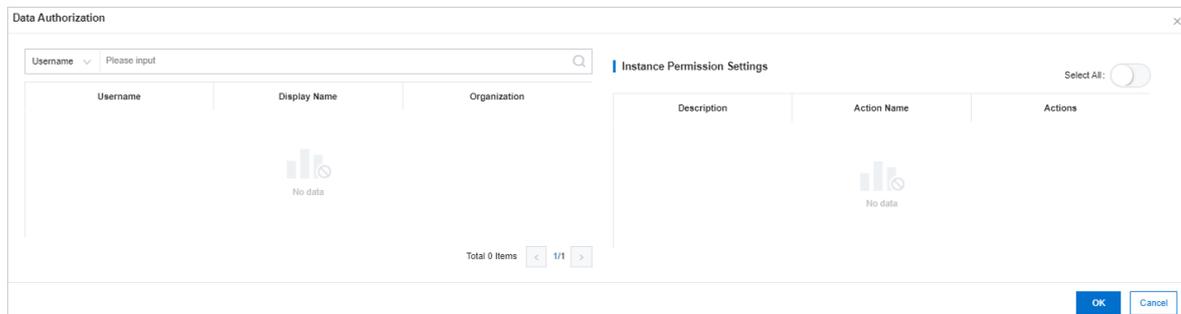
The cloud services that support data authorization include Message Queue (MQ), Object Storage Service (OSS), Log Service, DataHub, and Container Service.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#).
2. In the top navigation bar, click **Enterprise**.

- In the left-side navigation pane of the Enterprise page, click **Data Permissions**.
- Click a resource set and click a product type on the right side of the page.
- Click **Authorize** in the **Actions** column corresponding to the instance that you want to manage.
- In the Data Authorization dialog box, select a user on the left side.
- Turn on or off the data permission switches in the Actions column on the right side.

You can also turn on or off the Select All switch to manage permissions in batches.



- Click **OK**.

1.8.10.3. Edit user permissions

You can use JSON statements to edit user permissions.

Procedure

- [Log on to the Apsara Uni-manager Management Console](#).
- In the top navigation bar, click **Enterprise**.
- In the left-side navigation pane of the Enterprise page, click **Data Permissions**.
- In the organization navigation tree, click the ▶ icon to the left of the organization that contains the user you want to manage.
- Click **Users**.
- Enter the username in the search box and click **Search**.
- Click **Edit Permissions** in the **Actions** column corresponding to the user.
- In the Edit Permissions dialog box, select a data permission on the left side and click **OK**.

If no permissions are available, specify a policy in the text editor. For more information about the syntax and structure of a policy, see [Permission policy structure and syntax](#).

1.8.10.4. View the permissions of a user

You can view the existing policies of a user.

Procedure

- [Log on to the Apsara Uni-manager Management Console](#).
- In the top navigation bar, click **Enterprise**.
- In the left-side navigation pane of the Enterprise page, click **Data Permissions**.
- In the organization navigation tree, find the organization that contains the user you want to manage and click the ▶ icon.
- Click **Users**.

6. Enter the username in the search box and click **Search**.
7. Click **View Permissions** in the **Actions** column corresponding to the user.

1.9. Configurations

1.9.1. Security policies

1.9.1.1. Configure password policies

You can configure password policies for user logons.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as a platform administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane of the **Configurations** page, click **Security Policies**.
4. Click the **Password Policy** tab.
5. On the **Password Policy** tab, set the password policy parameters.

The screenshot displays the 'Password Policy' configuration interface. It includes the following settings:

- Password Length:** 10 to 32 Characters (Minimum: 8)
- Required Character Types in Password:** Lowercase Letters, Uppercase Letters, Digits, Special Characters
- Logon Disabled After Password Expires:** Yes, No
- Password Validity Period (Days):** 90 (The value must be 0 to 1095. The value 0 specifies that the password will not expire.)
- Password Attempts:** Failed logon attempts within an hour cannot exceed 0 password attempts within an hour. (The value must be 0 to 32. The value 0 specifies that the password history check is disabled.)
- Password History Check:** disables the first 0 passwords. (The value must be 0 to 24. The value 0 specifies that the password history check is disabled.)

Buttons for **Save** and **Reset** are located at the bottom of the form.

To restore to the default password policy, click **Reset**.

1.9.1.2. Configure logon control

You can limit the access from multiple clients of users.

1. Log on to the [Apsara Uni-manager Management Console](#) as a platform administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, click **Security Policies**.
4. Click the **Logon Control** tab.
5. Perform the following operations:
 - o Select **Single Session**. A single session means that a user is allowed to log on only by using a single client at the same time.
 - o Select **Multi-session**. A multi-session means that a user is allowed to log on by using multiple clients at the same time.

1.9.2. Menus

1.9.2.1. Create a menu

You can create a menu and add its URL to the Apsara Uni-manager Management Console for quick access.

Procedure

1. Log on to the **Apsara Uni-manager Management Console** as a platform administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane of the **Configurations** page, click **Menu Settings**.
4. On the **Main Menu** page, click **Create** in the upper-right corner.
5. In the **Create** dialog box, configure the menu parameters.

Menu parameters

| Parameter | Description |
|---------------------|---|
| Title | The display name of the menu. |
| URL | The URL of the menu. |
| Console Type | Different console types correspond to different domain names. <ul style="list-style-type: none"> ○ oneconsole: You need only to enter the path in the URL field. The domain name is automatically matched. ○ asconsole: You need only to enter the path in the URL field. The domain name is automatically matched. ○ other: You must enter the domain name in the URL field. |
| Icon | The icon displayed in the left-side navigation pane. The icon cannot be changed. |

| Parameter | Description |
|--------------|--|
| Identifier | The unique identifier of the menu in the system. This identifier can be used to indicate whether the menu is selected in the navigation bar. The identifier cannot be changed. |
| Order | The display order among the same-level menus. The larger the value, the lower the display order. Leave the Order field empty. |
| Parent Level | The displayed tree structure. |
| Open With | Specifies whether to open the menu in the current window or in a new window. |
| Description | The description of the menu. |

1.9.2.2. Modify a menu

You can modify an existing menu, including the menu name, URL, icon, and menu order.

Prerequisites

Default menus cannot be modified.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as a platform administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane of the **Configurations** page, click **Menu Settings**.
4. Click **Edit** in the **Actions** column corresponding to a menu.
5. In the **Edit** dialog box, modify the relevant information of the menu.

6. Click **OK**.

1.9.2.3. Delete a menu

You can delete menus that are no longer needed.

Prerequisites

Default menus cannot be deleted.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as a platform administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane of the **Configurations** page, click **Menu Settings**.
4. Find the menu that you want to delete and click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

1.9.2.4. Show or hide menus

You can perform the following operations to show or hide menus:

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as a platform administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane of the **Configurations** page, click **Menu Settings**.

4. Select or clear the check box in the **Displayed** column corresponding to a menu.

1.9.3. Specifications

1.9.3.1. Specification parameters

This topic describes the specification parameters of each resource type.

NAT Gateway

| Parameter | Description |
|-----------------------------------|---|
| Specifications | The specifications that can be configured for NAT Gateway. |
| Specifications Description | The description of the specifications that can be configured for NAT Gateway. |

AnalyticDB for PostgreSQL

| Parameter | Description |
|----------------------------|---|
| Specifications | The specifications that can be configured for AnalyticDB for PostgreSQL. |
| Specifications Name | The name of the instance type that can be configured for AnalyticDB for PostgreSQL. |
| CPU | The total number of CPU cores that can be configured for AnalyticDB for PostgreSQL. |
| Memory | The memory size that can be configured for AnalyticDB for PostgreSQL. |
| Storage Space | The total storage size that can be configured for AnalyticDB for PostgreSQL. |
| Version | The version number of AnalyticDB for PostgreSQL. |
| Node | The number of nodes that can be configured for AnalyticDB for PostgreSQL. |

SLB

| Parameter | Description |
|----------------------------|--|
| Specifications | The instance type that can be configured for Server Load Balancer (SLB). |
| Specifications Name | The name of the instance type that can be configured for SLB. |
| Maximum Connections | The maximum number of connections that can be configured for SLB. |
| New Connections | The number of new connections that can be configured for SLB. |

| Parameter | Description |
|-------------|---|
| QPS | The queries per second (QPS) that can be configured for SLB. |
| Description | The description of the specifications that can be configured for SLB. |

ApsaraDB RDS

| Parameter | Description |
|----------------------|--|
| Engine Type | The engine type that can be configured for ApsaraDB RDS. |
| Minimum Storage (GB) | The minimum amount of storage space that can be configured for ApsaraDB RDS. |
| Maximum Storage (GB) | The maximum amount of storage space that can be configured for ApsaraDB RDS. |
| Specifications Name | The name of the instance type that can be configured for ApsaraDB RDS. |
| Version | The version number of ApsaraDB RDS. |
| CPUs | The number of CPU cores that can be configured for ApsaraDB RDS. |
| Maximum Connections | The maximum number of connections that can be configured for ApsaraDB RDS. |
| Memory (GB) | The memory size that can be configured for ApsaraDB RDS. |
| Share Type | The share type that can be configured for ApsaraDB RDS. |

PolarDB-X

| Parameter | Description |
|---------------------|--|
| Instance Type | The instance type that can be configured for PolarDB-X. |
| Instance Type Name | The name of the instance type that can be configured for PolarDB-X. |
| Specifications | The specifications that can be configured for PolarDB-X. |
| Specifications Name | The name of the specifications that can be configured for PolarDB-X. |

ECS

| Parameter | Description |
|-----------------|---|
| Instance Family | The instance family that is divided into different instance types based on the scenarios for which they are suitable. |

| Parameter | Description |
|--------------------------------|--|
| Specifications Level | The level of the instance type that can be configured for Elastic Compute Service (ECS). |
| vCPUs | The maximum number of vCPUs that can be configured for ECS. |
| Memory (GB) | The memory size that can be configured for ECS. |
| Instance Specifications | The instance type that can be configured for ECS. |
| GPU Type | The GPU type that can be configured for ECS. |
| GPUs | The number of GPUs that can be configured for ECS. |
| Supported ENIs | The number of elastic network interface (ENIs) that can be configured for ECS. |
| Number Of Private IP Addresses | The number of private IP addresses that can be configured for ECS. |

KVStore for Redis

| Parameter | Description |
|-------------------------|---|
| Specifications Name | The name of the specifications that can be configured for KVStore for Redis. |
| Instance Specifications | The instance type that can be configured for KVStore for Redis. |
| Maximum Connections | The maximum number of connections that can be configured for KVStore for Redis. |
| Maximum Bandwidth | The maximum bandwidth that can be configured for KVStore for Redis. |
| CPUs | The number of CPU cores that can be configured for KVStore for Redis. |
| Version | The version number of KVStore for Redis. |
| Architecture | The architecture of KVStore for Redis. |
| Node Type | The node type of KVStore for Redis. |
| Service Plan | The service plan that can be configured for KVStore for Redis. |

ApsaraDB for MongoDB

| Parameter | Description |
|-------------------------|---|
| Specifications | The instance type that can be configured for ApsaraDB for MongoDB. |
| Instance Specifications | The name of the specifications that can be configured for ApsaraDB for MongoDB. |

| Parameter | Description |
|----------------------|--|
| Engine Type | The engine type that can be configured for ApsaraDB for MongoDB. |
| Version | The version number of ApsaraDB for MongoDB. |
| Serial Number | The serial number of ApsaraDB for MongoDB. |
| Sequence Description | The description of the serial number of ApsaraDB for MongoDB. |
| Maximum Connections | The maximum number of connections that can be configured for ApsaraDB for MongoDB. |
| IOPS | The IOPS of ApsaraDB for MongoDB. |
| Storage Space | The amount of storage space that can be configured for ApsaraDB for MongoDB. |
| Minimum Storage | The minimum amount of storage space that can be configured for ApsaraDB for MongoDB. |
| Maximum Storage | The maximum amount of storage space that can be configured for ApsaraDB for MongoDB. |

1.9.3.2. Create specifications

You can customize specifications for each resource type.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as a platform administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane of the **Configurations** page, click **Specifications**.
4. Select the resource type for which you want to create specifications and click the **Resource Specifications** tab.
5. On the **Resource Specifications** tab, click **Create Specifications** in the upper-right corner.
6. In the dialog box that appears, configure the specifications parameters.
For more information about specification parameters, see [Specification parameters](#).
7. Click **OK**.

1.9.3.3. View specifications

You can view the specifications of each resource type.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as a platform administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane of the **Configurations** page, click **Specifications**.
4. Click the resource type for which you want to view specifications.
5. On the **Resource Specifications** tab, set a **region**, **column**, and **value**. The corresponding information is displayed in the specifications list.

6. Click the **Existing Specifications** tab and view the existing specifications and their quantity.

1.9.3.4. Disable specifications

By default, newly created specifications are in the Enabled state.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as a platform administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane of the **Configurations** page, click **Specifications**.
4. Select the resource type for which you want to disable specifications.
5. Click **Disable** in the **Actions** column corresponding to the specifications that you want to disable.
6. In the message that appears, click **OK**.

1.9.3.5. Export specifications

You can export specifications that you want to view and share.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as a platform administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane of the **Configurations** page, click **Specifications**.
4. Click the resource type for which you want to export specifications.
5. Click the **Resource Specifications** tab and click **Export** in the upper-right corner to export the file to your PC.

1.9.3.6. View specifications of each resource type in previous versions

You can view specifications of each resource type in previous versions.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane of the **Configurations** page, click **Specifications**.
4. On the **Specifications** page, click the resource type for which you want to view specifications.
5. Click the **Specifications History** tab. View the detailed information in the specifications list.

1.9.4. Message center

1.9.4.1. View internal messages

You can view all internal messages, including read and unread messages.

Context

When changes are made to or alerts are generated for an instance in a resource, all users that have read and operation permissions on this resource receive corresponding messages.

Procedure

1. [Log on to the Apsara Uni-manager Management Console.](#)
2. In the top navigation bar, move the pointer over the  icon and click **More**.
3. In the left-side navigation pane of the **Message Center** page, click the target message scope.
 - Choose **Internal Messages > All Messages** to view all messages, including unread and read messages.
 - Choose **Internal Messages > Unread Messages** to view unread messages.
 - Choose **Internal Messages > Read Messages** to view read messages.

1.9.4.2. Mark messages as read

You can mark unread messages as read messages to facilitate message management.

Procedure

1. [Log on to the Apsara Uni-manager Management Console.](#)
2. In the top navigation bar, move the pointer over the  icon and click **More**.
3. In the left-side navigation pane of the **Message Center** page, choose **Internal Messages > Unread Messages**.
In the upper part of the **Unread Messages** page, click different message types to filter messages.
4. On the **Unread Messages** page, find the message that you want to mark as read and click **Mark as Read** in the **Actions** column.
You can also select the check boxes to the left of messages and click **Batch Read** in the upper-left corner of the page.
5. In the **Mark as Read** message, click **OK**.

1.9.4.3. Delete a message

You can delete messages that are no longer needed.

Procedure

1. [Log on to the Apsara Uni-manager Management Console.](#)
2. In the top navigation bar, move the pointer over the  icon and click **More**.
3. In the left-side navigation pane of the **Message Center** page, choose **Internal Messages > All Messages**.
4. Find the message that you want to delete on the **All Messages** tab or other tabs and click **Delete**.
You can also select the check boxes to the left of messages and click **Batch Delete** in the upper-left corner of the page.

1.9.5. Resource pool management

You can modify the maximum usage of each resource.

Prerequisites

- If the physical inventory is unlimited, the logical inventory cannot be less than the used inventory.
- If the physical inventory is limited, the logical inventory cannot be less than the used inventory or greater than the physical inventory.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane of the **Configurations** page, click **Resource Pool Management**.
4. On the **Resource Pools** page, click the  icon in the module that you want to modify and modify the number of resources.

Resource Pool Configuration

Region: 

| ECS | | | | VPC | | | | OSS | | | |
|-----------------------|-------------------|--------------------|------|-----------|-------------------|--------------------|------|----------------|-------------------|--------------------|------|
| Item | Logical Inventory | Physical Inventory | Used | Item | Logical Inventory | Physical Inventory | Used | Item | Logical Inventory | Physical Inventory | Used |
| CPU Quota | 20,000 | Unknown | 31 | VPC Quota | 10,000 | Unlimited | 4 | OSS Quota (GB) | Not Set | Unknown | 0 |
| Memory Quota (GB) | 60,000 | Unknown | 219 | | | | | | | | |
| GPU Quota | 60,000 | Unknown | 0 | | | | | | | | |
| SSD Quota (GB) | 600,000 | Unknown | 80 | | | | | | | | |
| Ultra Disk Quota (GB) | 6,000,000 | Unknown | 520 | | | | | | | | |

| RDS-MysQL | | | | RDS-SQLServer | | | | RDS-postgreSQL | | | |
|-------------------|-------------------|--------------------|------|-------------------|-------------------|--------------------|------|-------------------|-------------------|--------------------|------|
| Item | Logical Inventory | Physical Inventory | Used | Item | Logical Inventory | Physical Inventory | Used | Item | Logical Inventory | Physical Inventory | Used |
| CPU Quota | Not Set | Unknown | 0 | CPU Quota | Not Set | Unknown | 0 | CPU Quota | Not Set | Unknown | 0 |
| Memory Quota (GB) | Not Set | Unknown | 0 | Memory Quota (GB) | Not Set | Unknown | 0 | Memory Quota (GB) | Not Set | Unknown | 0 |
| Disk Quota (GB) | Not Set | Unknown | 0 | Disk Quota (GB) | Not Set | Unknown | 0 | Disk Quota (GB) | Not Set | Unknown | 0 |

| SLB | | | | EIP | | | | ODPS | | | |
|-------------------------|-------------------|--------------------|------|-----------|-------------------|--------------------|------|-----------------|-------------------|--------------------|------|
| Item | Logical Inventory | Physical Inventory | Used | Item | Logical Inventory | Physical Inventory | Used | Item | Logical Inventory | Physical Inventory | Used |
| Virtual IP Quota | 2,304 | 2,304 | 0 | EIP Quota | 10,000 | Unlimited | 1 | CU Quota | Not Set | Unknown | 0 |
| Public Virtual IP Quota | 512 | 512 | 0 | | | | | Disk Quota (GB) | Not Set | Unknown | 0 |

5. Click the  icon to complete modification.

1.9.6. Custom configurations

1.9.6.1. Configure brands

You can customize the icon, platform name, and logon page of the Apsara Uni-manager Management Console logon interface.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as a platform administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane, click **Custom Configurations**.
4. In the **Brand Settings** section, click a language tab of the page that you want to modify.
5. Configure the following parameters and click OK.

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|------------------|---|
| Icon | <p>You can click Upload Image, select an image file, and then click Open.</p> <p>The recommended resolution of an image is 160 × 36 pixels, and the recommended formats are PNG and JPG.</p> <p>After the image is uploaded, you can click the  icon to preview it.</p> <p>Before you replace the image, you must click the  icon to delete the current image.</p> |
| Platform Name | <p>You can customize a platform name.</p> <p>The name can be up to 8 characters in length.</p> |
| Logon Page | <p>You can click + Add Screen or the  icon after screens to increase or decrease the number of scrolling screens.</p> <p>The number of scrolling screens on the logon page must be in the range of one to three.</p> <p>After you click the tab of a screen, you can perform the following operations:</p> <ul style="list-style-type: none"> ◦ Background Image: Click Upload Image, select a background image file, and then click Open. <p>The recommended resolution of an image is 1,880 × 1,600 pixels, and the recommended formats are PNG and JPG.</p> <p>After the image is uploaded, you can click the  icon to preview it.</p> <p>Before you replace the image, you must click the  icon to delete the current image.</p> <ul style="list-style-type: none"> ◦ Image Copywriting: Customize the image copywriting. <p>The image copywriting can be up to 40 characters in length.</p> |
| Copyright Notice | <p>You can customize the information of the copyright notice.</p> |

1.10. Operations

1.10.1. Quotas

1.10.1.1. Quota parameters

This topic describes the quota parameters of each service.

An organization administrator can set resource quotas and create resources within the allowed quotas for the organization. When the quotas for the organization are used up, the system does not allow the organization administrator to create more resources for the organization. To create more resources, you must first increase the quotas for the organization.

If no quotas are set, you can create an unlimited amount of resources.

ECS

| Parameter | Description |
|-----------------------|--|
| CPU Quota (Cores) | The total number of CPU cores that you can configure for Elastic Compute Service (ECS) and the number of used cores. |
| Memory Quota (GB) | The total memory size that you can configure for ECS. |
| GPU Quota (Cores) | The total number of GPU cores that you can configure for ECS. |
| SSD Quota (GB) | The total SSD capacity that you can configure for ECS. |
| Ultra Disk Quota (GB) | The total number of disks that you can configure for an ECS instance. |

VPC

| Parameter | Description |
|-----------|---|
| VPC Quota | The maximum number of virtual private clouds (VPCs) that you can configure. |

OSS

| Parameter | Description |
|----------------|--|
| OSS Quota (GB) | The maximum capacity that you can allocate for Object Storage Service (OSS). |

ApsaraDB RDS for MySQL

| Parameter | Description |
|-------------------|---|
| CPU Quota (Cores) | The total number of CPU cores that you can configure for ApsaraDB RDS for MySQL and the number of used cores. |
| Memory Quota (GB) | The total memory size that you can configure for ApsaraDB RDS for MySQL. |
| Disk Quota (GB) | The total storage size that you can configure for ApsaraDB RDS for MySQL. |

PolarDB

| Parameter | Description |
|-------------------|--|
| CPU Quota (Cores) | The total number of CPU cores that you can configure for PolarDB and the number of used cores. |
| Memory Quota (GB) | The total memory size that you can configure for PolarDB. |
| Disk Quota (GB) | The total storage size that you can configure for PolarDB. |

ApsaraDB RDS for SQL Server

| Parameter | Description |
|--------------------------|--|
| CPU Quota (Cores) | The total number of CPU cores that you can configure for ApsaraDB RDS for SQL Server and the number of used cores. |
| Memory Quota (GB) | The total memory size that you can configure for ApsaraDB RDS for SQL Server. |
| Disk Quota (GB) | The total storage size that you can configure for ApsaraDB RDS for SQL Server. |

ApsaraDB RDS for PostgreSQL

| Parameter | Description |
|--------------------------|--|
| CPU Quota (Cores) | The total number of CPU cores that you can configure for ApsaraDB RDS for PostgreSQL and the number of used cores. |
| Memory Quota (GB) | The total memory size that you can configure for ApsaraDB RDS for PostgreSQL. |
| Disk Quota (GB) | The total storage size that you can configure for ApsaraDB RDS for PostgreSQL. |

SLB

| Parameter | Description |
|--------------------------------|--|
| Virtual IP Quota (a) | The maximum number of internal IP addresses that you can configure for Server Load Balancer (SLB). |
| Public Virtual IP Quota | The maximum number of public IP addresses that you can configure for SLB. |

EIP

| Parameter | Description |
|------------------|---|
| EIP Quota | The maximum number of elastic IP addresses (EIPs) that you can configure. |

MaxCompute

| Parameter | Description |
|------------------------|---|
| CU Quota (a) | The total number of capacity units (CUs) that you can configure for MaxCompute. |
| Disk Quota (GB) | The total storage size that you can configure for MaxCompute. |

KVStore for Redis

| Parameter | Description |
|--------------------------|---|
| Memory Quota (GB) | The total memory size that you can configure for KVStore for Redis. |

PolarDB-X

| Parameter | Description |
|-------------------|--|
| CPU Quota (Cores) | The total number of CPUs that you can configure for PolarDB-X. |

AnalyticDB for PostgreSQL

| Parameter | Description |
|-------------------|--|
| CPU Quota (Cores) | The total number of CPU cores that you can configure for AnalyticDB for PostgreSQL and the number of used cores. |
| Memory Quota (GB) | The total memory size that you can configure for AnalyticDB for PostgreSQL. |
| Disk Quota (GB) | The total storage size that you can configure for AnalyticDB for PostgreSQL. |

ApsaraDB for MongoDB

| Parameter | Description |
|-------------------|---|
| CPU Quota (Cores) | The total number of CPU cores that you can configure for ApsaraDB for MongoDB and the number of used cores. |
| Memory Quota (GB) | The total memory size that you can configure for ApsaraDB for MongoDB. |
| Disk Quota (GB) | The total storage size that you can configure for ApsaraDB for MongoDB. |

1.10.1.2. Set quotas for a cloud service

The Apsara Uni-manager Management Console allows you to set quotas to properly allocate resources among organizations.

Prerequisites

You must set quotas for a parent organization before you can set quotas for its subordinate organizations.

Context

If the parent organization has quotas (except when the parent organization is a level-1 organization), the available quotas for a subordinate organization are equal to the quotas for the parent organization minus the quotas for other subordinate organizations.

This topic describes how to modify quotas for Elastic Compute Service (ECS). You can modify quotas for other cloud resources in a similar manner.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an organization administrator.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane of the **Operations** page, click **Quotas**.
4. In the left-side navigation tree, click the name of the organization for which you want to create cloud resources.

5. Select the cloud service for which you want to set quotas. In this example, ECS is selected.
6. In the upper-right corner of the quota section, click **Set Quota**.
7. Set the total quotas and click **Save**.

For more information about quota parameters, see [Quota parameters](#).

1.10.1.3. Modify quotas

Administrators can adjust quotas for cloud resources based on organizational requirements.

Context

This topic describes how to modify quotas for ECS. You can modify quotas for other cloud resources in a similar manner.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane of the **Operations** page, click **Quotas**.
4. In the left-side navigation tree, click the name of the organization for which you want to create cloud resources.
5. Select the Apsara Stack service for which you want to modify quotas. For this example, ECS is selected.
6. In the upper-right corner of the quota area, click **Modify**.
7. Set the total quotas and click **Save**.

For more information about quota parameters, see [Quota parameters](#).

1.10.1.4. Reset quotas

Administrators can reset quotas as needed.

Prerequisites

Before deleting a quota for an organization, make sure that no subordinate organizations have any quotas.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane of the **Operations** page, click **Quotas**.
4. In the left-side organization navigation tree, click the name of the target organization.
5. Select the cloud service for which you want to reset quotas. For this example, ECS is selected.
6. In the upper-right corner of the quota section, click **Reset**.
7. In the message that appears, click **OK**.

1.10.2. Usage statistics

1.10.2.1. View the usage statistics of cloud resources

The Apsara Uni-manager Management Console shows the number of resource instances that run in the Apsara Stack environment by time, organization, resource set, and region. You can also export statistical reports from the Apsara Uni-manager Management Console.

Context

The cloud resources that can be metered include Elastic Compute Service (ECS), Virtual Private Cloud (VPC), Server Load Balancer (SLB), Object Storage Service (OSS), ApsaraDB RDS for MySQL, Elastic IP Address (EIP), Tablestore, PolarDB-X, KVStore for Redis, AnalyticDB for PostgreSQL, ApsaraDB for MongoDB, ApsaraDB RDS for PostgreSQL, ApsaraDB RDS for SQL Server, Log Service, ECS disks, ECS snapshots, scaling group rules, API gateways, and Key Management Service (KMS).

This topic describes how to create quotas for ECS. You can set quotas for other cloud resources in a similar manner.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#).
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane of the **Operations** page, click **Usage Statistics**.
4. In the **Resource Type** section, click **Elastic Compute Service ECS**.
5. In the **Search Conditions** section, set **Time Period**, **Organization**, **Resource Set**, **Region**, and **Instance ID** to filter resources.

You can view the statistics in the console or click **Export** in the upper-right corner to export the statistics to your PC as an XLS file.

 **Note** In the console, you can view or export up to 1,000 statistical records to an Excel file. Use the statistics query API to obtain more statistical data.

The exported file is named *<Resource type name>.xls*. Find the downloaded file from the download path of the browser.

1.10.3. Statistical analysis

1.10.3.1. View reports of current data

You can use reports to view the most recent data of each service.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#).
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane of the **Operations** page, choose **Statistical Analysis > Reports**.
4. Click the tab that you want to view.

You can click the **Resource Reports**, **Quota Reports**, or **CloudMonitor Reports** tab.

5. (Optional) Set **Organizations and Resource Sets** and **Region** and click **Search**.
6. Click the tab of the service that you want to view.

View the most recent data of the service.

1.10.3.2. Export reports of current data

You can batch export data that you want to view by cloud service.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#).
2. In the top navigation bar, click **Operations**.

3. In the left-side navigation pane of the Operations page, choose **Statistical Analysis > Reports**.
4. Click the tab that you want to view.
You can click the **Resource Reports**, **Quota Reports**, or **CloudMonitor Reports** tab.
5. (Optional) Set **Organizations and Resource Sets** and **Region** for which you want to view data and click **Search**.
Click **Reset** to clear all filter properties.
6. Use one of the following methods to export data:
 - o Export data by service.
 - a. Click **Export Reports** on the right side of the page.
 - b. In the **Select Products to Export** dialog box, select the service that you want to view and click **OK**.
You can also select **Select All** in the lower-left corner and click **OK**.
 - o Export data by instance.
 - a. Click the tab of the service that you want to view.
 - b. Select the instance that you want to view and click **Export Selected Reports** in the lower-left corner of the page.

1.10.3.3. Download reports of historical data

You can download data reports of cloud services within the specified period of time, resource set, and region by creating download tasks.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#).
2. In the top navigation bar, click **Operations**.
3. In the left-side navigation pane of the Operations page, choose **Statistical Analysis > Download Center**.
4. In the upper-right corner of the page, click **Create Download Task**.

5. Enter the information of the download task.

Create Download Task
✕

*Report Name:

*Report Type:

Select an option
▼

*Product:

Select an option
▼

*Start Time and End Time:

Select a date
📅

Select a date
📅

*Organizations and Resource Sets:

Select one or more organizations or resource sets
▼

*Region:

Select an option
▼

OK
Cancel

| Parameter | Description |
|---------------------------------|--|
| Report Name | The name of the report. |
| Report Type | The type of the report. Valid values: <ul style="list-style-type: none"> ○ Resource Reports ○ CloudMonitor Reports |
| Product | The cloud service for which you want to download reports. You can select multiple cloud services. |
| Start Time and End Time | The start and end time of the data. |
| Organizations and Resource Sets | The organization to which the data belongs. You can select multiple organizations. |
| Region | The region of the data. You can select multiple regions. |

6. Click OK.

7. After the **Created** message appears, the Download Center page appears. Enter the information of the created report in the search box and click **Search** to search for the created download task.

- After **In Progress** changes to **Completed** in the **Status** column, click **Download Report** in the **Actions** column.

1.11. Security

1.11.1. View operation logs

You can view operation logs to obtain up-to-date information for various resources and functional modules in the Apsara Uni-manager Management Console. You can also export operation logs to your PC.

Procedure

- [Log on to the Apsara Uni-manager Management Console](#) as a security administrator.
- In the top navigation bar, click **Security**.
- You can filter logs by username, resource type, resource ID, product, organization, keyword, source IP address, start time, and end time.

The following table describes the fields in the query result.

Fields in the query result

| Log field | Description |
|----------------------|--|
| Organization | The organization of the object on which operations are performed. |
| Resource Set | The resource set of the object on which operations are performed. |
| Resource Type | The resource type of the object on which operations are performed. |
| Resource ID | The ID of the object on which operations are performed. |
| Status | The current status of the operation. |
| User | The name of the operator. |
| Event Type | The operation performed on an Apsara Stack service. The operations include creating, modifying, deleting, querying, updating, binding, unbinding, enabling and disabling service instances, applying for and releasing service instances, and changing the ownership of service instances. |
| Source IP | The IP address of the operator. |
| Details | A brief introduction of the operation. |
| Start Time | The time when the operation started. |
| End Time | The time when the operation ended. |
| Region | The region of the object on which operations are performed. |
| Level | The level of logs. |

 **Note** Click the  icon. In the **Show/Hide Columns** dialog box, select fields that you want to show in the query result.

- (Optional) Click **Download** to download the logs displayed on the current page to your PC as an XLS file.
The exported log file is named *log.xls* and stored in the *C:\Users\Username\Downloads* directory.

1.12. RAM

1.12.1. RAM introduction

Resource Access Management (RAM) is a resource access control service provided by Apsara Stack.

You can use RAM to manage users and control which resources are accessible to employees, systems, and applications.

RAM provides the following features:

- RAM role
To authorize a cloud service in a level-1 organization to use other resources in the organization, you must create a RAM role. This role specifies the operations that the cloud service can perform on resources.
Only system administrators and level-1 organization administrators can create RAM roles.
- User group
You can create multiple users within an organization and grant them different operation permissions on cloud resources.
You can create RAM user groups to classify and authorize RAM users within your Apsara Stack tenant account. This simplifies the management of RAM users and their permissions.
You can create RAM permission policies to grant different operation permissions to different user groups.

1.12.2. Permission policy structure and syntax

This topic describes the structure and syntax used to create or update permission policies in Resource Access Management (RAM).

Policy characters and usage rules

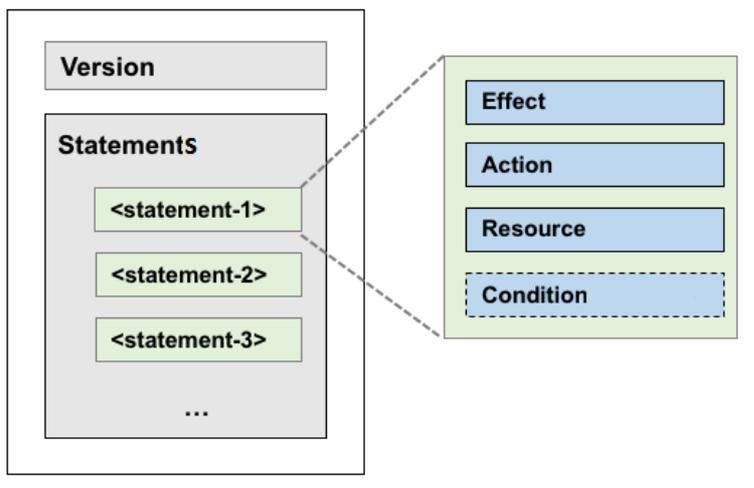
- Characters in a policy
 - The following characters are JSON tokens and are included in policies: `{ } [] " , : .`
 - The following characters are special characters in the syntax and are not included in policies: `= < > () |`.
- Use of characters
 - If an element can have more than one value, you can perform the following operations:
 - Separate multiple values by using commas (,) as delimiters between each value and use an ellipsis (...) to describe the remaining values. Example: `[<action_string>, <action_string>, ...]`.
 - Include only one value. Examples: `"Action": [<action_string>]` and `"Action": <action_string>`.
 - A question mark (?) following an element indicates that the element is optional. Example: `<condition_block?>`.
 - A vertical bar (|) between elements indicates multiple options. Example: `("Allow" | "Deny")`.
 - Elements that must be text strings are enclosed in double quotation marks ("). Example: `<version_block> = "Version": ("1")`.

Policy structure

The policy structure includes the following components:

- The version number.
- A list of statements. Each statement contains the following elements: Effect, Action, Resource, and Condition.

The Condition element is optional.



Policy syntax

```

policy = {
  <version_block>,
  <statement_block>
}
<version_block> = "Version" : ("1")
<statement_block> = "Statement" : [ <statement>, <statement>, ... ]
<statement> = {
  <effect_block>,
  <action_block>,
  <resource_block>,
  <condition_block? >
}
<effect_block> = "Effect" : ("Allow" | "Deny")
<action_block> = ("Action" | "NotAction") :
  ("*" | [<action_string>, <action_string>, ...])
<resource_block> = ("Resource" | "NotResource") :
  ("*" | [<resource_string>, <resource_string>, ...])
<condition_block> = "Condition" : <condition_map>
<condition_map> = {
  <condition_type_string> : {
    <condition_key_string> : <condition_value_list>,
    <condition_key_string> : <condition_value_list>,
    ...
  },
  <condition_type_string> : {
    <condition_key_string> : <condition_value_list>,
    <condition_key_string> : <condition_value_list>,
    ...
  }, ...
}
<condition_value_list> = [<condition_value>, <condition_value>, ...]
<condition_value> = ("String" | "Number" | "Boolean")
    
```

Description:

- The current policy version is 1.
- The policy can have multiple statements.

- The effect of each statement can be either `Allow` or `Deny`.

 **Note** In a statement, both the Action and Resource elements can have multiple values.

- Each statement can have its own conditions.

 **Note** A condition block can contain multiple conditions with different operators and logical combinations of these conditions.

- You can attach multiple policies to a RAM user. If policies that apply to a request include an `Allow` statement and a `Deny` statement, the Deny statement overrides the Allow statement.
- Element value:
 - If an element value is a number or Boolean value, it must be enclosed in double quotation marks (") in the same way as strings.
 - If an element value is a string, characters such as the asterisk (*) and question mark (?) can be used for fuzzy matching.
 - The asterisk (*) indicates any number (including zero) of allowed characters. For example, `ecs:Describe*` indicates all ECS API operations that start with `Describe`.
 - The question mark (?) indicates an allowed character.

Policy format check

Policies are stored in RAM as JSON documents. When you create or update a policy, RAM first checks whether the JSON format is valid.

- For more information about JSON syntax standards, see [RFC 7159](#).
- We recommend that you use tools such as JSON validators and editors to check whether the policies meet JSON syntax standards.

1.12.3. RAM roles

1.12.3.1. View basic information about a RAM role

You can view basic information about a RAM role, including its user groups and existing permission policies.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. On the Roles page, click the name of the target RAM role.
5. In the basic information section, click the **User Groups** and **Permissions** tabs to view relevant information.

1.12.3.2. Create a RAM role

To authorize a cloud service in a level-1 organization to use other resources in the organization, you must create a RAM role. This role contains the operations that the cloud service can perform on resources.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.

- In the upper-right corner of the page, click **Create RAM Role**.
- On the **Roles - Create RAM Role** page, set **Role Name**, **Description**, and **Sharing Scope**.
Valid values of the **Sharing Scope** parameter:
 - Global**
The role is visible and valid to all organizations involved. The default value is Global.
 - Current Organization**
The role is visible and valid to the organization to which the user belongs.
 - Subordinate Organization**
The role is visible and valid to the organization to which the user belongs and its subordinate organizations.
- Click **Create**.

1.12.3.3. Create a policy

To use a cloud service to access other cloud resources, you must create a policy and attach it to a user group.

Procedure

- [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
- In the top navigation bar, click **Enterprise**.
- In the left-side navigation pane of the **Enterprise** page, click **Roles**.
- In the role name list, find the RAM role that you want to modify and choose **More > Modify** in the **Actions** column to go to the **Roles** page.
- Click the **Permissions** tab.
- Click **Add Permission Policy**.
- In the Add Permission Policy dialog box, enter information of the policy.

Add Permission Policy

*Policy Name:
Enter a policy name 0/15

Description:
Enter 0 to 100 characters 0/100

*Policy Details:
1 | The details of the specified policy must be 2,048 characters in length, and follow the JSON format

OK Cancel

For more information about how to enter the policy content, see [Permission policy structure and syntax](#).

1.12.3.4. Modify the content of a RAM policy

You can modify the content of a RAM policy.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. In the role list, find the RAM role that you want to modify and choose **More > Modify** in the **Actions** column to go to the **Roles** page.
5. Click the **Permissions** tab.
6. Click the name of a policy in the **Permission Policy Name** column.
7. In the **Modify Permission Policy** dialog box, modify the relevant information and click **OK**.

For more information about how to modify the policy content, see [Permission policy structure and syntax](#).

1.12.3.5. Modify the name of a RAM policy

You can modify the name of a RAM policy.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. In the role name list, find the RAM role that you want to modify and choose **More > Modify** in the **Actions** column to go to the **Roles** page.
5. Click the **Permissions** tab. Click the name of that policy that you want to modify in the **Permission Policy Name** column.
6. In the **Modify Permission Policy** dialog box, modify the policy name.

1.12.3.6. Add a RAM role to a user group

You can bind RAM roles to user groups.

Prerequisites

You must create a user group before RAM roles can be added. For more information, see [Add a role](#).

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. In the role name list, find the RAM role that you want to modify and choose **More > Modify** in the **Actions** column to go to the **Roles** page.
5. Click the **User Groups** tab.
6. Click **Add User Group**. In the **Add User Group** dialog box, select a user group.
7. Click **OK**.

1.12.3.7. Grant permissions to a RAM role

When you grant permissions to a RAM role, all users in the user groups that are assigned this role share the granted permissions.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. In the role list, find the RAM role that you want to modify and choose **More > Modify** in the **Actions** column to go to the **Roles** page.
5. Click the **Permissions** tab.
6. Click **Select Existing Permission Policy**.
7. In the dialog box that appears, select a RAM policy and click **OK**.
If no RAM policies are available, see [Add a permission policy](#).

1.12.3.8. Remove permissions from a RAM role

You can remove permissions that are no longer needed from RAM roles.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. In the role list, find the RAM role that you want to modify and choose **More > Modify** in the **Actions** column to go to the **Roles** page.
5. Click the **Permissions** tab.
6. Find the policy that you want to remove and click **Remove** in the **Actions** column.

1.12.3.9. Modify a RAM role name

Administrators can modify the names of RAM roles.

Context

The name of a preset role cannot be modified.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. In the role name list, find the RAM role that you want to modify and choose **More > Modify** in the **Actions** column to go to the **Roles** page.
5. Move the pointer over the role name and click the  icon to enter a new role name.

1.12.3.10. Delete a RAM role

This topic describes how to delete a RAM user.

Prerequisites

Before you delete a RAM role, make sure that no policies are attached to the RAM role.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. In the role name list, click **More** in the **Actions** column corresponding to a RAM role, and choose **Delete** from the shortcut menu.
5. In the message that appears, click **OK**.

1.12.4. RAM authorization policies

1.12.4.1. Create a service-linked role

You can create authorization policies and grant them to organizations.

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane of the **Configurations** page, click **Service-linked Roles**.
4. In the upper-right corner of the page, click **Create RAM Role**.
5. On the **Create RAM Role** page, set **Organization Name** and **Service Name**.
6. Click **OK**.

1.12.4.2. View the details of a service-linked role

You can view the details of a Resource Access Management (RAM) role, including its role name, creation time, description, and Alibaba Cloud Resource Name (ARN).

Procedure

1. Log on to the [Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane of the **Configurations** page, click **Service-linked Roles**.
4. On the **Service-linked Roles** page, set **Role Name**, **Service Name**, or **Organization Name**, and click **Search** in the upper-right corner.
To perform another search, click **Clear**.
5. Find the service-linked role that you want to view and click **Details** in the **Actions** column.

1.12.4.3. View RAM authorization policies

You can view the details of a Resource Access Management (RAM) authorization policy, including its policy name, policy type, default version, description, association time, and policy content.

Prerequisites

A RAM authorization policy is created. For more information, see [Create a RAM role](#).

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Configurations**.
3. In the left-side navigation pane of the **Configurations** page, click **Service-linked Roles**.
4. On the **RAM Roles** page, set **Role Name** or **Service Name** and click **Search** in the upper-right corner. To perform another search, click **Clear**.
5. Find the service-linked role that you want to view and click **Details** in the **Actions** column.
6. Click the **Role Policy** tab to view the information of the role authorization policy. Click **Details** in the **Actions** column to view the policy details.

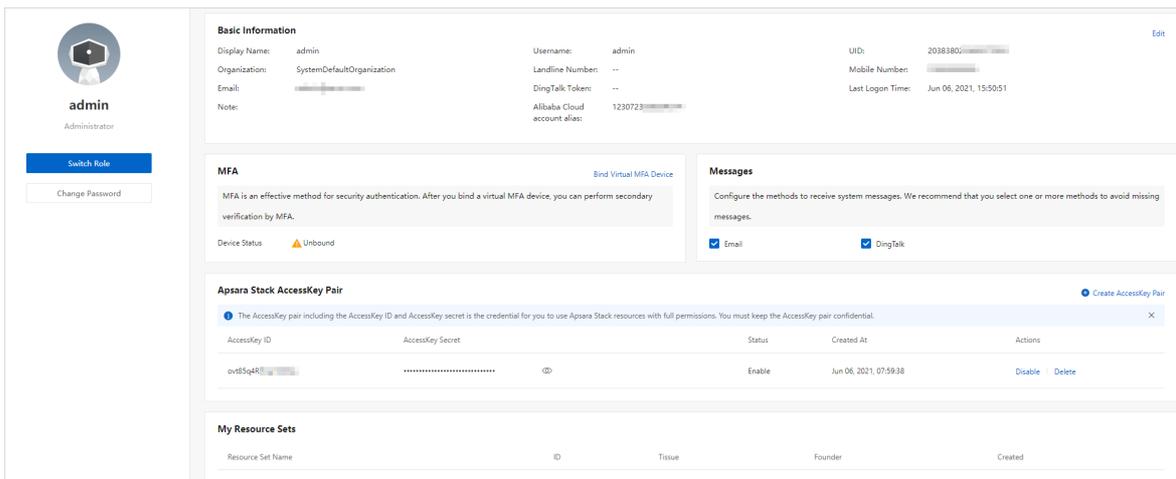
1.13. Personal information management

1.13.1. Modify personal information

You can modify your personal information to keep it up to date.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#).
2. In the upper-right corner of the homepage, move the pointer over the profile picture and choose **User Information** from the shortcut menu.



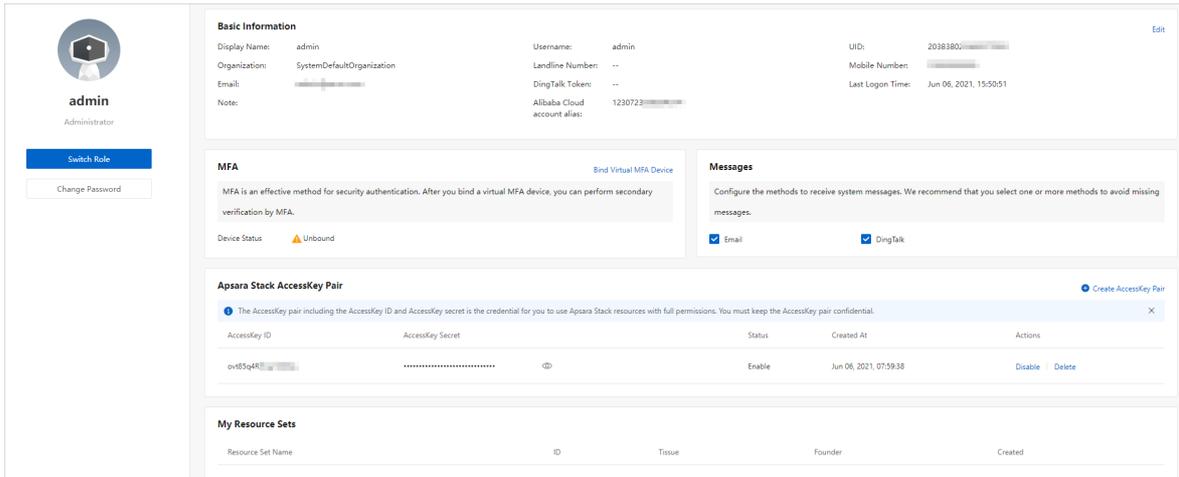
3. In the upper-right corner of the **Basic Information** section, click **Edit**.
4. In the **Modify User Information** dialog box, modify the personal information.
5. Click **OK**.

1.13.2. Change the logon password

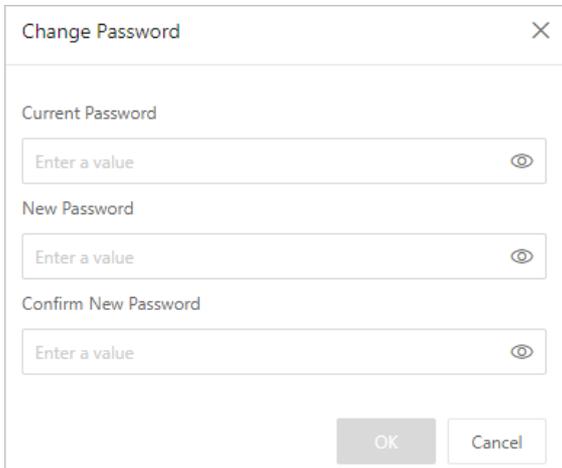
To improve security, you must change your logon password in a timely manner. This topic describes how to change your logon password.

Procedure

1. [Log on to the Apsara Uni-manager Management Console](#).
2. In the upper-right corner of the homepage, move the pointer over the profile picture and choose **User Information** from the shortcut menu.



3. Click **Change Password**.
4. In the **Change Password** dialog box, set **Current Password**, **New Password**, and **Confirm New Password**.



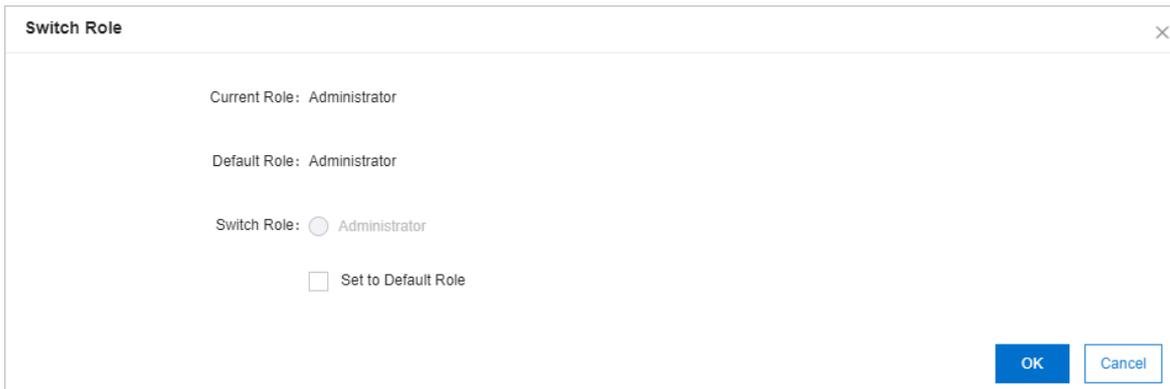
5. Click **OK**.

1.13.3. Switch the current role

You can switch the scope of your current role.

Procedure

1. **Log on to the Apsara Uni-manager Management Console** as an administrator.
2. In the upper-right corner of the homepage, move the pointer over the user profile picture and choose **User Information** from the short cut menu.
3. Click **Switch Role**.
4. In the **Switch Role** dialog box that appears, select the role that you want to switch to.



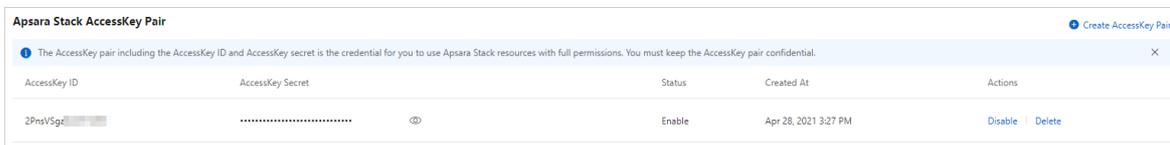
You can also switch back to the default role.

1.13.4. View the AccessKey pair of your Apsara Stack tenant account

To ensure the security of cloud resources, the system must verify the identity of visitors and ensure that they have the relevant permissions. You must obtain the AccessKey ID and AccessKey secret of your Apsara Stack tenant to access cloud resources.

Procedure

1. [Log on to the Apsara Uni-manager Management Console.](#)
2. In the upper-right corner of the homepage, move the pointer over the profile picture and choose **User Information** from the short cut menu.
3. In the **Apsara Stack AccessKey Pair** section, view your AccessKey pair.



Note The AccessKey pair consists of an AccessKey ID and an AccessKey secret. These credentials provide you with full permissions on Apsara Stack resources. You must keep the AccessKey pair confidential.

1.13.5. Create an AccessKey pair

You can delete your old AccessKey pairs and create new ones to implement the rotation of your AccessKey pairs.

Procedure

1. [Log on to the Apsara Uni-manager Management Console.](#)
2. In the upper-right corner of the homepage, move the pointer over the profile picture and choose **User Information** from the short cut menu.
3. In the upper-right corner of the **Apsara Stack AccessKey Pair** section, click **Create AccessKey Pair**.

Note Each user can have up to two AccessKey pairs.

1.13.6. Delete an AccessKey pair

You can delete AccessKey pairs that are no longer needed.

Prerequisites

- Each user can delete only its own AccessKey pairs. The administrator cannot delete the AccessKey pairs of users.
- Each user retains at least one AccessKey pair.

Procedure

1. [Log on to the Apsara Uni-manager Management Console.](#)
2. In the upper-right corner of the homepage, move the pointer over the profile picture and choose **User Information** from the short cut menu.
3. In the **Apsara Stack AccessKey Pair** section, find the AccessKey pair that you want to delete and click **Delete** in the **Actions** column.
4. In the message that appears, click **OK**.

1.13.7. Disable an AccessKey pair

You can disable AccessKey pairs that are no longer needed. Newly created AccessKey pairs are in the Enable state by default.

Procedure

1. [Log on to the Apsara Uni-manager Management Console.](#)
2. In the upper-right corner of the homepage, move the pointer over the profile picture and choose **User Information** from the short cut menu.
3. In the **Apsara Stack AccessKey Pair** section, find the AccessKey pair that you want to disable and click **Disable** in the **Actions** column.

Each user can disable only its own AccessKey pairs. The administrator cannot disable the AccessKey pairs of users.

4. In the message that appears, click **OK**.

Note

Make sure that at least one AccessKey pair is in the Enable state.

1.13.8. Enable an AccessKey pair

Disabled AccessKey pairs must be re-enabled before you can continue to use them. Newly created AccessKey pairs are in the Enable state by default.

Procedure

1. [Log on to the Apsara Uni-manager Management Console.](#)
2. In the upper-right corner of the homepage, move the pointer over the profile picture and choose **User Information** from the short cut menu.
3. In the **Apsara Stack AccessKey Pair** section, find the AccessKey pair that you want to enable and click **Enable** in the **Actions** column.
4. In the message that appears, click **OK**.

 Note

Make sure that at least one AccessKey pair is in the Enable state.

1.13.9. MFA

1.13.9.1. Overview

Multi-factor authentication (MFA) is an identity authentication method in computer systems. It requires users to provide two or more verification factors to gain access to a resource. This topic introduces the principles and use scenarios of MFA.

Introduction to MFA

When MFA is enabled, you must enter your username and password (first security factor) and then a variable verification code (second security factor) from an MFA device when you log on to the Apsara Uni-manager Management Console. Two-factor authentication enhances security for your account.

MFA devices use the Time-based One-time Password (TOTP) algorithm to generate time-dependent 6-digit dynamic verification codes. The Apsara Uni-manager Management Console supports software-based virtual MFA devices. You can install software such as the Alibaba Cloud app that supports MFA on your mobile device such as your mobile phone to act as a virtual MFA device.

Limits

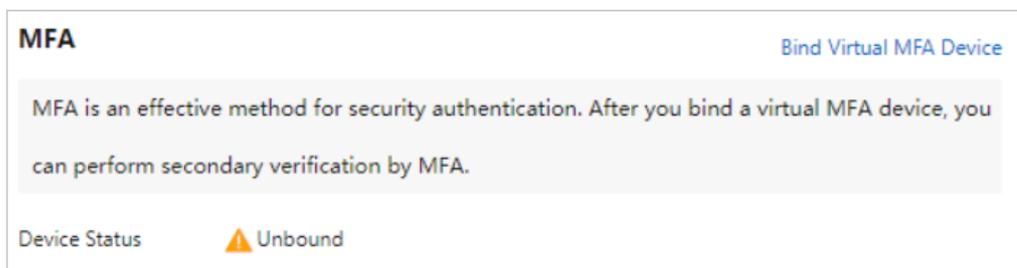
The TOTP algorithm requires that the time of the system clock of the Apsara Uni-manager Management Console remain consistent with the standard time on the Internet. Otherwise, discrepancies in time can lead to inconsistent MFA verification codes and leave you unable to log on to the Apsara Uni-manager Management Console.

1.13.9.2. Bind a virtual MFA device to enable MFA

Multi-factor authentication (MFA) is automatically enabled after you bind an MFA device. This topic describes how to bind a virtual MFA device.

1. [Log on to the Apsara Uni-manager Management Console.](#)
2. In the upper-right corner of the homepage, move the pointer over the profile picture and choose **User Information** from the shortcut menu.
3. In the upper-right corner of the **MFA** section, click **Bind Virtual MFA Device**.
4. On the **Bind Virtual MFA Device** page, follow the instructions to bind an MFA device.

The following figure shows the MFA section after the MFA device is bound.



Results

After a virtual MFA device is bound, you must enter a 6-digit MFA verification code in addition to your username and password before you can log on to the Apsara Uni-manager Management Console.

1.13.9.3. Unbind a virtual MFA device to disable MFA

To disable multi-factor authentication (MFA), you must disable the MFA device that is bound. This topic describes how to unbind a virtual MFA device to disable MFA.

1. [Log on to the Apsara Uni-manager Management Console](#).
2. In the upper-right corner of the homepage, move the pointer over the profile picture and choose **User Information** from the short cut menu.
3. In the upper-right corner of the **MFA** section, click **Disable Virtual MFA Device**.
4. In the dialog box that appears, click **Disable Virtual MFA Device**.

After you disable MFA, you need only to enter your username and password the next time you log on to the Apsara Uni-manager Management Console.

1.13.9.4. Forcibly enable MFA

Administrators including the platform administrator, operations administrator, and organization administrator can check whether their users have multi-factor authentication (MFA) enabled. If MFA is disabled for the users, the administrators can forcibly enable MFA for the users.

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Users > Users**.
4. Click the **System Users** tab.
5. Find the user for whom you want to forcibly enable MFA and choose **More > View MFA Status** in the **Actions** column.

You can search for the user by username, organization, or role.

6. In the MFA Status Prompt message, click **OK**.

Note

- After MFA is forcibly enabled for a user, the user must go to the Bind Virtual MFA Device page to bind a virtual MFA device before the user can log on to the Apsara Uni-manager Management Console.
- MFA can be forcibly enabled for users, but cannot be forcibly disabled. Before MFA is enabled, the MFA status of the user is **Not Enabled**.
- After MFA is enabled, the MFA status of the user is **Enabled but Not Bound**.

1.13.9.5. Reset MFA

If you lose your multi-factor authentication (MFA) key, you must have the administrator reset the key. The MFA key is automatically reset while the password is reset.

1. [Log on to the Apsara Uni-manager Management Console](#) as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane, choose **Users > Users**.
4. Click the **System Users** tab.
5. Find the user for whom you want to reset MFA and click the username.

You can search for the user by username, organization, or role.

6. In the user details panel, click **Reset password**.

2. Elastic Compute Service (ECS)

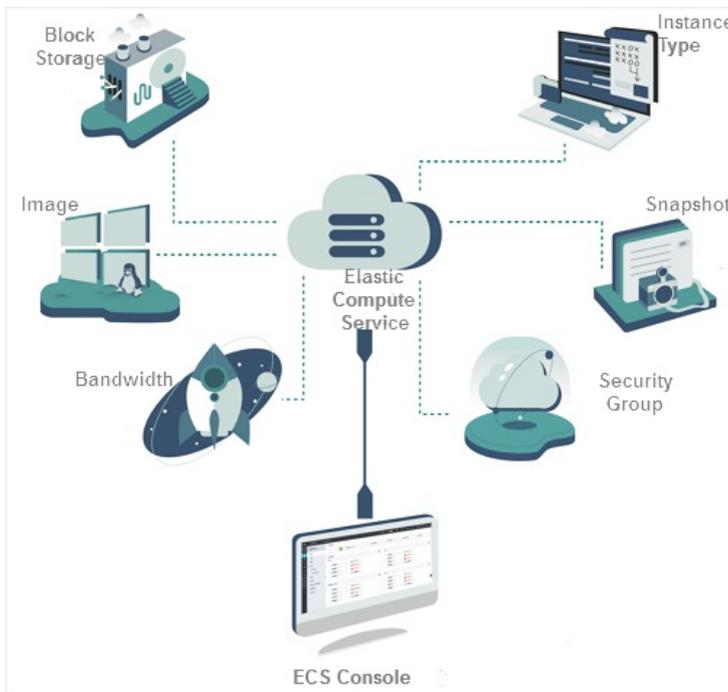
2.1. What is ECS?

2.1.1. Overview

Elastic Compute Service (ECS) is a type of computing service that features elastic processing capabilities. Compared with physical servers, ECS can be more efficiently managed and is more user-friendly. You can create instances, resize disks, and add or release any number of ECS instances at any time based on your business needs.

An ECS instance is a virtual computing environment that contains the most basic components of computers such as the CPU, memory, and storage. Users perform operations on ECS instances. Instances are core components of ECS, and operations can be performed on instances by using the ECS console. Other resources such as block storage, images, and snapshots can only be used after they are integrated into ECS instances. For more information, see [ECS components](#).

ECS components



2.1.2. Instance lifecycle

The lifecycle of an ECS instance begins when the instance is created and ends when the instance is released. This topic describes the instance states in the ECS console, state attributes, and corresponding instance states in API responses.

The following table describes the states that an ECS instances may go through during its lifecycle.

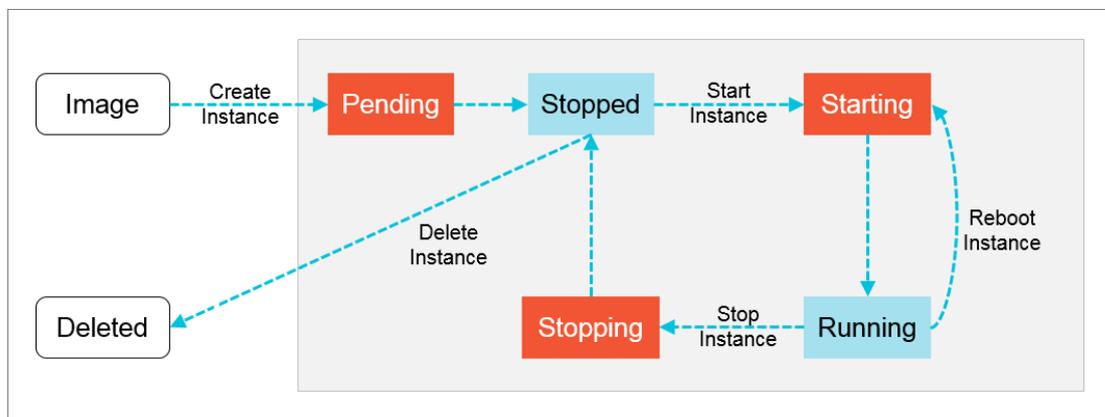
Instance states

| State | State attribute | Description | State in an API response |
|------------------------|-----------------|--|--------------------------|
| Instance being created | Intermediate | The instance is being created and waiting to be started. If an instance remains in the Instance being created state for an extended period of time, an exception has occurred. | Pending |

| State | State attribute | Description | State in an API response |
|----------------------|-----------------|--|--------------------------|
| Starting | Intermediate | When you start or restart an instance by using the ECS console or calling an API operation, the instance enters this state before it enters the Running state. If an instance remains in the Starting state for an extended period of time, an exception has occurred. | Starting |
| Running | Stable | While an instance is in the Running state, the instance can function normally and can accommodate your business needs. | Running |
| Stopping | Intermediate | When you stop an instance by using the ECS console or calling an API operation, the instance enters this state before it enters the Stopped state. If an instance remains in the Stopping state for an extended period of time, an exception has occurred. | Stopping |
| Stopped | Stable | An instance enters this state when it is stopped. An instance in the Stopped state cannot provide external services. | Stopped |
| Reinitializing | Intermediate | When you re-initialize the system disk or a data disk of an instance by using the ECS console or calling an API operation, the instance enters this state before it enters the Running state. If an instance remains in the Reinitializing state for an extended period of time, an exception has occurred. | Stopped |
| Changing system disk | Intermediate | When you replace the system disk of an instance by using the ECS console or calling an API operation, the instance enters the Changing system disk state before it enters the Running state. If an instance remains in the Changing system disk state for an extended period of time, an exception has occurred. | Stopped |

Instance states describes the relationships between instance states in the ECS console and instance states in API responses. The following figure shows the transitions between instance states in API responses.

Transitions between instance states in API responses



2.2. Instructions

2.2.1. Restrictions

Learn about restrictions before performing operations on ECS instances.

- Do not upgrade the kernel or operating system version of an ECS instance.
- Do not start SELinux for Linux systems except CentOS and RedHat.
- Do not detach PVDriver.
- Do not arbitrarily modify the MAC address of the network interface.

2.2.2. Suggestions

Consider the following suggestions to make more efficient use of ECS:

- ECS instances with 4 GiB or higher memory must use a 64-bit operating system. 32-bit operating systems have a maximum of 4 GiB of memory addressing.
- A 32-bit Windows operating system supports a maximum of 4 CPU cores.
- To ensure service continuity and avoid failover-induced service unavailability, we recommend that you configure service applications to boot automatically at system startup.

2.2.3. Limits

Before using ECS instances, you must be familiar with the limits of instance families.

General limits

- Windows operating systems support a maximum of 64 vCPUs in instance specifications.
- ECS instances do not support the installation of virtualization software and secondary virtualization.
- Sound card applications are not supported. Only GPU instances support virtual sound cards. External hardware devices, such as hardware dongles, USB flash drives, external hard disks, and bank U keys, cannot be directly connected to ECS instances.
- ECS does not support multicast protocols. If multicasting services are required, we recommend that you use unicast instead.

Instance family ga1

To create a ga1 instance, you must use one of the following images pre-installed with drivers:

- Ubuntu 16.04 with an AMD GPU driver pre-installed
- Windows Server 2016 English version with an AMD GPU driver pre-installed
- Windows Server 2008 R2 English version with an AMD GPU driver pre-installed

Note:

- A ga1 instance uses an optimized driver provided by Alibaba Cloud and AMD. The driver is installed in images provided by Alibaba Cloud and is currently unavailable for download.
- If the GPU driver malfunctions due to improper removal of related components, you need to replace the system disk to restore GPU related functions.

 **Note** This operation causes data loss.

- If the driver malfunctions because an improper image is selected, you need to replace the system disk to reselect an image with an AMD GPU driver pre-installed.
- For Windows Server 2008 or earlier, you cannot connect to the VNC after the GPU driver takes effect. The VNC is irresponsive with a black screen or stuck at the splash screen. You can use other methods such as Remote

Desktop Protocol (RDP) to access the system.

- RDP does not support DirectX, OpenGL, or other related applications. You need to install the VNC and a client, or use other supported protocols such as PCoIP and XenDesktop HDX 3D.

Instance families gn4, gn5i, and gn5

- **Bandwidth:** If you use an image of Windows Server 2008 R2 for a gn4 instance, you cannot use the Connect to VNC function in the ECS console to connect to the instance after the installed GPU driver takes effect. You need to set the bandwidth to a non-zero value or attach an Elastic IP address to the created instance.
- **Image:** If an NVIDIA GPU driver is not required, you can select any image, and then [Install the CUDA and GPU drivers for a Linux instance](#) or [Install the CUDA and GPU drivers for a Windows instance](#).

2.2.4. Notice for Windows users

Before using Windows-based ECS instances, you must consider the following points:

- Data loss may occur if a local disk is used as the data disk of an instance. We recommend that you use a cloud disk to create your instance if you are not sure about the reliability of the data architecture.
- Do not close the built-in shutdownmon.exe process. Otherwise, the server may require a longer time to restart.
- Do not rename, delete, or disable Administrator accounts or it may affect the use of the server.
- We do not recommend that you use virtual memory.
- When you modify your computer name, you must synchronize the following key values in the registry. Otherwise, the computer name cannot be modified, causing failure when installing certain third-party programs. The following key values must be modified in the registry:

```
HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Control\ComputerName\ActiveComputerName
```

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\ComputerName\ComputerName
```

2.2.5. Notice for Linux users

Before using Linux-based ECS instances, you must consider the following points:

- Do not modify content of the default /etc/issue files under a Linux instance. Otherwise, the custom image created from the instance cannot be recognized, and instances created based on the image cannot start as expected.
- Do not arbitrarily modify the permissions of each directory in the partition where the root directory is located, especially permissions of /etc, /sbin, /bin, /boot, /dev, /usr, and /lib directories. Improper modification of permissions can cause errors.
- Do not rename, delete, or disable Linux root accounts.
- Do not compile or perform any other operations on the Linux kernel.
- We do not recommend the use of Swap for partitioning.
- Do not enable the NetworkManager service. This service conflicts with the internal network service of the system, causing network errors.

2.2.6. Notice on defense against DDoS attacks

You need to purchase Anti-DDoS Pro to defend against DDoS attacks. For more information, see *Apsara Stack Security Product Introduction*.

2.3. Quick start

2.3.1. Overview

This topic describes how to quickly create and connect to an ECS instance.

Perform the following procedure:

1. **Create a security group**

A security group is a virtual firewall used to control traffic to and from ECS instances. Each ECS instance must be added to at least one security group. Before creating an instance, you must select a security group to control traffic to and from the instance.

2. **Create an instance**

An ECS instance is a virtual machine that contains basic computing components such as CPU, memory, operating system, network, and disks. After a security group is created, you can select an instance type based on your business requirements. For more information, see [Instance types](#).

3. **Connect to an instance**

Select a remote connection method based on the network configuration and operating system of the ECS instance and your local operating system. After you log on to the instance, you can perform other operations on it, such as installing applications.

2.3.2. Log on to the ECS console

This topic describes how to log on to the ECS console.

Prerequisites

- Before logging on to the ASCM console, make sure that you have obtained the IP address or domain name of the ASCM console from the deployment personnel. The URL used to access the ASCM console is in the following format: `https://[IP address or domain name of the ASCM console]`.
- We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL used to access the ASCM console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password for logging on to the console from the operations administrator.

 **Note** When you log on to the ASCM console for the first time, you must change the password of your username as prompted. Due to security concerns, your password must meet the minimum complexity requirements: The password must be 8 to 20 characters in length and must contain at least two of the following character types: uppercase letters, lowercase letters, digits, and special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%).

3. Click **Login** to go to the ASCM console homepage.
4. In the top navigation bar, choose **Products > Elastic Computing > Elastic Compute Service**.

2.3.3. Create a security group

Security groups are an important means for network security isolation. They are used to set network access control for one or more ECS instances.

Prerequisites

A Virtual Private Cloud (VPC) has been created. For more information, see *VPC User Guide*.

Context

Instances that belong to the same account and are in the same region and in the same security group can communicate with each other over the internal network. If instances that belong to the same account in the same

region are in different security groups, you can implement internal network communication by authorizing mutual access between two security groups.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > Security Groups**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Click **New Security Group**.
5. Configure the parameters of the security group.

| Type | Parameter | Required | Description |
|----------------|---------------------|----------|--|
| Region | Organization | Yes | The organization to which the security group belongs. Make sure that the security group and the VPC belong to the same organization. |
| | Resource Set | Yes | The resource set to which the security group belongs. Make sure that the security group and the VPC belong to the same resource set. |
| | Region | Yes | The region to which the security group belongs. Make sure that the security group and the VPC belong to the same region. |
| | Zone | Yes | The ID of the zone where the security group resides. |
| Basic Settings | VPC | Yes | The VPC to which the security group belongs. |
| | Security Group Name | No | The name must be 2 to 128 characters in length and start with a letter. It can contain letters, digits, periods (.), underscores (_), hyphens (-), colons (:), and commas (,). It cannot start with http:// or https://. |
| | | | |

| Type | Parameter | Required | Description |
|------|-------------|----------|--|
| | Description | No | The description of the security group. We recommend that you provide an informational description to simplify future management operations. The name must be 2 to 256 characters in length and start with a letter. It can contain letters, digits, periods (.), underscores (_), hyphens (-), and commas (,). It cannot start with http:// or https://. |

6. Click **Submit**.

2.3.4. Create an instance

An ECS instance is a virtual machine that contains the basic computing components of a server, such as CPU, memory, operating system, network, and disks.

Prerequisites

- A VPC and a VSwitch are created. For more information, see *Create a VPC and Create a VSwitch* in Apsara Stack VPC User Guide.
- If you want to assign an IPv6 address to the instance that you want to create, make sure that the VPC and VSwitch are associated with IPv6 CIDR blocks. For more information, see *Enable an IPv6 CIDR block for a VPC and Enable an IPv6 CIDR block for a VSwitch* in Apsara Stack VPC User Guide.
- A security group is available. If no security group is available, create a security group. For more information, see [Create a security group](#).

Context

Some limits apply when you create GPU-accelerated instances. For more information, see [Limits](#).

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Instances**.
3. Click **Create Instance**.
4. Configure parameters listed in the following tables to create an instance.
 - i. Configure the basic settings of the instance.

| Parameter | Required | Description |
|--------------|----------|---|
| Organization | Yes | Select an organization in which to create the instance. |
| Resource Set | Yes | Select a resource set in which to create the instance. |

ii. Select a region and zone for the instance.

| Parameter | Required | Description |
|-----------|----------|--|
| Region | Yes | Select a region in which to create the instance. |
| Zone | Yes | Select a zone in which to create the instance. Zones are the physical zones with separate power supplies and networks in the same region. The internal networks of zones are interconnected, and faults in one zone are isolated from the other zones. To increase the availability of your applications, we recommend that you create instances in different zones. |

iii. Configure the network of the instance.

| Parameter | Required | Description |
|--------------|----------|---|
| Network Type | Yes | Select the type of the network in which to create the instance. VPC is available. |
| VPC | Yes | Select a VPC in which to create the instance. |
| VSwitch | Yes | Select a VSwitch to be connected to the instance. |
| Private IP | No | Specify a private IP address for the instance. The private IP address must be within the CIDR block of the VSwitch. If you do not specify a private IP address, the system will automatically allocate a private IP address to the instance. |

iv. (Optional) Specify whether to assign an IPv6 address to the instance.

v. Select a security group in which to create the instance.

vi. Select an instance family and instance type for the instance.

| Parameter | Required | Description |
|-----------------|----------|--|
| Instance Family | Yes | Select an instance family for the instance. After you select an instance family, you must select an instance type. |
| Instance Type | Yes | Select an instance type. Windows Server images require specific CPU and memory combinations. For more information, see Limits in <i>ECS Product Introduction</i> . |

vii. Configure the image to be used by the instance.

| Parameter | Required | Description |
|--------------|---------------------------|--|
| Image Type | Yes | Select an image type. Valid values: Public Image and Custom Image . |
| Public Image | Subject to the image type | <p>Select a public image for the instance. Public images provided by Alibaba Cloud are licensed, secure, and stable. Public images include Windows Server images and major Linux images.</p> <p>This parameter must be specified when you set Image Type to Public Image.</p> <p>When you use an image that supports DHCPv6 to create an instance, an IPv6 address is automatically assigned to the instance. The created instance can use this IPv6 address to communicate over the internal network. When you use an image that does not support DHCPv6 to create an instance, you must manually assign an IPv6 address to the instance. The following images support DHCPv6:</p> <ul style="list-style-type: none"> ■ Linux images: <ul style="list-style-type: none"> ■ CentOS 7.6 IPv6 64-bit ■ CentOS 6.10 64-bit ■ SUSE Linux Enterprise Server 12 SP4 64-bit ■ Windows Server images <div style="background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p>Note To use an IPv6 address to communicate over the Internet, you must also enable public bandwidth for the IPv6 address. For more information, see Enable Internet bandwidth for an IPv6 address in <i>Apsara Stack VPC User Guide</i>.</p> </div> |
| Custom Image | Subject to the image type | <p>Select a custom image for the instance. Custom images are created from instances or snapshots, or imported from your local device.</p> <p>This parameter must be specified when you set Image Type to Custom Image.</p> |

viii. Configure the storage settings for the instance.

| Parameter | Required | Description |
|-------------|----------|--|
| System Disk | Yes | Specify the disk category and capacity of the system disk. Valid disk categories: Ultra Disk and SSD Disk . The system disk size must range from 20 GiB to 500 GiB. |
| Data Disk | No | You can click Data Disk to add data disks. Specify the disk category and capacity of each data disk. Valid disk categories: Ultra Disk and SSD Disk . A maximum of 16 data disks can be added to an instance. The maximum capacity of each data disk is 32 TiB. You can select or clear Release with Instance and Encrypt for each data disk. To encrypt a data disk, set Encryption Algorithm to AES256 or SM4 and set Encryption Key to a key created in Key Management Service (KMS) . You can also add data disks after the instance is created. For more information, see Create a disk . |

ix. Configure the logon password settings for the instance.

| Parameter | Required | Description |
|------------------|----------|--|
| Set Password | Yes | Specify when to set the password. Valid values: Now and Later . If Later is selected, you can use the password reset feature to set a password at a later time. For more information, see Change the logon password of an instance . <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px;">Note The password is used to log on to the instance, not to the VNC.</div> |
| Logon Password | No | Set the password to be used to log on to the instance. The password must be 8 to 30 characters in length and must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. The supported special characters include () ` ~ ! @ # \$ % ^ & * - _ + = { } [] ; ' < > , . ? / |
| Confirm Password | No | Re-enter the password. |

x. (Optional) Select a deployment set in which to create the instance.

xi. (Optional) Enter a name for the instance.

The name must be 2 to 128 characters in length and start with a letter. It can contain periods (.), underscores (_), colons (:), and hyphens (-).

If you do not specify a name, the system will assign an instance name at random.

xii. (Optional) In the User Data field, enter the user data to be automatically run upon instance start up.

Windows supports Batch and PowerShell scripts. Before you perform Base64 encoding, make sure that the content to be encoded includes `[bat]` or `[powershell]` as the first line. Linux supports shell scripts.

xiii. Enter the number of instances that you want to create.

The number must be an integer ranging from 1 to 100.

5. Click **Submit**.

Result

The instance appears in the instance list. When the instance is being created, it is in the **Preparing** state. After the instance is created, it enters the **Running** state.

2.3.5. Connect to an instance

2.3.5.1. Instance connecting overview

After an instance is created, you can connect to the instance to perform operations such as installing applications.

You can use one of the following methods to connect to an instance:

- Use remote connection tools to connect to instances that have public IP addresses. For more information about the procedure, see the following topics:
 - [Connect to a Linux-based instance by using SSH commands in Linux or Mac OS X](#)
 - [Connect to a Linux-based instance by using remote connection tools in Windows](#)
 - [Connect to a Windows-based instance by using RDP](#)
- Use the VNC feature in the ECS console. For more information, see [Connect to an instance by using a VNC management terminal](#).

The username of a Windows instance is Administrator, and that of a Linux instance is root.

2.3.5.2. Connect to a Linux instance by using SSH commands in Linux or Mac OS X

This topic describes how to use SSH commands to connect to a Linux instance.

Prerequisites

- The instance and the security group are created.
- The instance is in the **Running** state.
- A logon password is set for the instance.
- An Elastic IP address (EIP) is bound with the instance.
- An inbound security group rule is added to the security group to allow the SSH port.

| Rule direction | Authorization policy | Protocol type | Port range | Priority | Authorization type | Authorization object |
|----------------|----------------------|---------------|------------|----------|--------------------|----------------------|
| Inbound | Accept | TCP | 22/22 | 1 | IPv4 CIDR block | 0.0.0.0/0 |

Procedure

1. Enter the following command and press the Enter key.

```
ssh root@instance IP
```

2. Enter the instance password of the root user and press the Enter key.

2.3.5.3. Connect to a Linux-based instance by using remote connection tools in Windows

This topic describes how to connect to an instance by using the PuTTY tool.

Prerequisites

Remote connection tools are designed with similar logics. In this example, PuTTY is used to connect to an instance. Download PuTTY at the following URL: .

Procedure

1. Download and install PuTTY for Windows.
 2. Start the PuTTY client and complete the following settings:
 - Host Name (or IP Address): Enter the EIP of the instance to be connected.
 - Port: Select the default port 22.
 - Connection Type: Select SSH.
 - Saved Session: Enter the name of the session. Click **Save**. After the settings are saved, PuTTY remembers the name and IP address of the instance. This eliminates the need to enter them every time you connect to the instance.
 3. Click **Open** to connect to the instance.

When you connect to the instance for the first time, PuTTY displays security alerts. Click **Yes** to proceed.
 4. Enter the username root and press Enter.
 5. Enter the password for the instance and press Enter.
- If a message similar to the following one appears, a connection to the instance is established.

```
Welcome to aliyun Elastic Compute Server!
```

2.3.5.4. Connect to a Windows instance by using RDP

This topic describes how to connect to a Windows instance by using Remote Desktop Protocol (RDP).

Prerequisites

- The instance and the security group are created.
- The instance is in the **Running** state.
- A logon password is set for the instance.
- An Elastic IP address (EIP) is bound with the instance.
- An inbound security group rule is added to the security group to allow the RDP port.

| Rule direction | Authorization policy | Protocol type | Port range | Priority | Authorization type | Authorization object |
|----------------|----------------------|---------------|------------|----------|--------------------|----------------------|
| Inbound | Accept | TCP | 3389/3389 | 1 | IPv4 CIDR block | 0.0.0.0/0 |

- CredSSP-related security updates are installed on the operating system of the instance.

Procedure

1. Activate the Remote Desktop Connection feature by using any of the following methods:

- Click **Start**, enter `mstsc` in the search box, and click `mstsc` in the search result.
 - Press Windows Key + R. In the **Run** dialog box that appears, enter `mstsc` and click **OK**.
2. In the **Remote Desktop Connection** dialog box, enter the EIP of the instance and click **Show Options**.
 3. Enter the username.
The default username is `administrator`.
 4. (Optional) If you do not want to enter the password upon subsequent logons, select **Allow me to save credentials**.
 5. Click **Connect**.
 6. In the **Windows Security** dialog box that appears, enter the password for the account and click **OK**.

Result

After you log on to the instance, the Windows desktop appears.

If authentication errors occur or the required function is not supported, install security updates.

1. [Connect to an ECS instance by using the VNC](#) before proceeding.
2. Choose **Start > Control Panel**.
3. Click **System and Security**.
4. Click **Check for updates** in the **Windows Updates** pane.
5. If updates are available, click **Install updates**.
6. Restart the instance.

2.3.5.5. Connect to an ECS instance by using the VNC

You can access your instance by using the VNC in the ECS console when other SSH clients such as PuTTY, Xshell, and SecureCRT do not work properly.

Prerequisites

- The instance is in the **Running** state.
- The root certificate is imported to your web browser. For more information, see [Install a certificate](#).
- If you log on to an instance for the first time after the instance is created, make sure that you set a new VNC password. For more information, see [Change the VNC password](#).

Context

The VNC password is used to log on to the VNC of the ECS console, and the instance password is used to log on to an instance.

You can use the VNC to connect to an instance to solve issues shown in the following table.

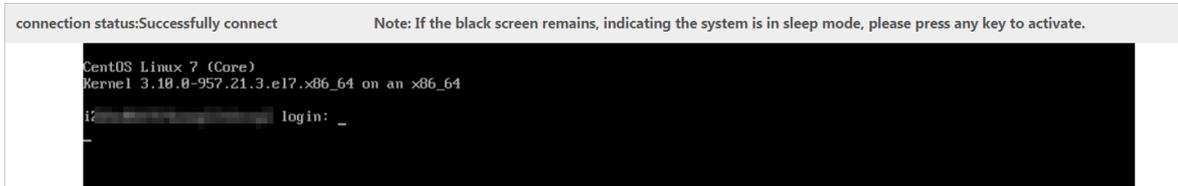
| Scenario | Resolution |
|---|--|
| The instance startup is slowly due to self-check upon startup. | Check the progress of the self check. |
| The firewall of the operating system is enabled by mistake. | Disable the firewall. |
| Abnormal processes appear, which consume large amounts of CPU or bandwidth resources. | Troubleshoot and terminate the abnormal processes. |

Procedure

1. [Log on to the ECS console](#).

2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the ECS instance, and click **Remote Connection** in the **Actions** column.
5. Enter the VNC password, and then click **OK**.

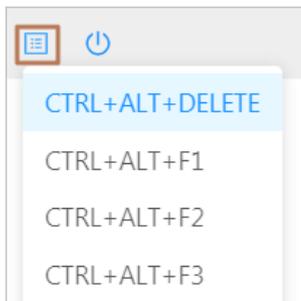
After the connection is successful, the logon page is displayed, as shown in the following figure.



6. Enter the username and password.
 - o For Linux instances: Enter the username `root` and the logon password.

Note Passwords in Linux are not displayed as you type. Press the Enter key after you enter the password.

- o For Windows instances: To use key combinations such as Ctrl + Alt + Delete, click the corresponding key combination in the upper-right corner of the VNC page.



Enter the username and password as prompted, and click the Log In icon .

2.4. Instances

2.4.1. Create an instance

An Elastic Compute Service (ECS) instance is a virtual machine that contains the basic computing components of a server, such as CPU, memory, operating system, network settings, and disks.

Prerequisites

- A virtual private cloud (VPC) and a vSwitch are created. For more information, see the "Create a VPC" and "Create a vSwitch" topics in *Apsara Stack VPC User Guide*.
- If you want to assign an IPv6 address to the instance that you want to create, make sure that the VPC and vSwitch are associated with IPv6 CIDR blocks. For more information, see the "Enable an IPv6 CIDR block for a VPC" and "Enable an IPv6 CIDR block for a vSwitch" topics in *Apsara Stack VPC User Guide*.
- One or more security groups are available. If no security groups are available, create one. For more information, see [Create a security group](#).

Context

Some limits apply when you create GPU-accelerated instances. For more information, see [Limits](#).

Procedure

1. [Log on to the ECS console.](#)
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Click **Create Instance**.
5. Configure the parameters listed in the following tables.
 - i. Configure the basic settings of the instance.

| Parameter | Required | Description |
|--------------|----------|-------------------------|
| Organization | Yes | Select an organization. |
| Resource Set | Yes | Select a resource set. |

- ii. Select a region and zone for the instance.

| Parameter | Required | Description |
|-----------|----------|---|
| Region | Yes | Select a region. |
| Zone | Yes | Select a zone. Zones are the physical zones with separate power supplies and networks within the same region. The internal networks of zones are connected, and faults in one zone are isolated from the other zones. To increase the availability of your applications, we recommend that you create instances in different zones. |

- iii. Configure network settings for the instance.

| Parameter | Required | Description |
|--------------------|----------|---|
| Network Type | Yes | Select a network type. Only VPC is supported. |
| VPC | Yes | Select a VPC. |
| vSwitch | Yes | Select a vSwitch. |
| Private IP Address | No | Specify a private IPv4 address for the instance. The private IPv4 address must be within the CIDR block of the selected vSwitch. If you do not specify a private IP address, the system allocates one to the instance. |
| IPv6 | No | Specify whether to assign an IPv6 address to the instance. |

- iv. Configure instance settings such as security group, instance family, and instance type and specify the number of instances that you want to create.

| Parameter | Required | Description |
|-----------------|----------|--|
| Security Group | Yes | Select a security group. |
| Deployment Set | No | Select a deployment set. You can use deployment sets to disperse or aggregate the instances involved in your business. |
| User Data | No | <p>In the User Data field, enter the user data to be automatically run on instance startup.</p> <ul style="list-style-type: none">Windows supports batch and PowerShell scripts. Before you perform Base64 encoding of user data, make sure that the data to be encoded includes <code>[bat]</code> or <code>[powershell]</code> as the first line.Linux supports shell scripts. |
| Quantity | Yes | Specify the number of instances that you want to create. The number must be an integer within the range of 1 to 100. |
| Instance Family | Yes | Select an instance family. After you select an instance family, you must select an instance type. |
| Instance Type | Yes | Select an instance type. Instance types that have specific CPU and memory combinations do not support Windows Server images. For more information, see <i>ECS Product Introduction</i> . |

- v. Configure the image to be used by the instance.

| Parameter | Required | Description |
|------------|----------|--|
| Image Type | Yes | Select an image type. Valid values: Public Image , Custom Image , and Shared Custom Image . |

| Parameter | Required | Description |
|---------------------|---------------------------|---|
| Public Image | Subject to the image type | <p>Select a public image. Public images provided by Alibaba Cloud are licensed, secure, and stable. Public images include Windows Server images and major Linux images.</p> <p>This parameter is required when you set Image Type to Public Image.</p> <p>When you use an image that supports Dynamic Host Configuration Protocol version 6 (DHCPv6) to create an instance, an IPv6 address is automatically assigned to the instance. The instance can use this IPv6 address to communicate over the internal network. When you use an image that does not support DHCPv6 to create an instance, you must manually assign an IPv6 address to the instance. The following images support DHCPv6:</p> <ul style="list-style-type: none"> ■ Linux images: <ul style="list-style-type: none"> ■ CentOS 7.6 IPv6 64-bit ■ CentOS 6.10 64-bit ■ SUSE Linux Enterprise Server 12 SP4 64-bit ■ Windows Server images <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note To use an IPv6 address to communicate over the Internet, you must enable public bandwidth for the IPv6 address. For more information, see the "Enable Internet bandwidth for an IPv6 address" topic in <i>Apsara Stack VPC User Guide</i>.</p> </div> |
| Custom Image | Subject to the image type | <p>Select a custom image. Custom images are created from instances or snapshots or are imported from your device.</p> <p>This parameter is required when you set Image Type to Custom Image.</p> |
| Shared Custom Image | Subject to the image type | <p>Select a custom image that is shared by another Apsara Stack tenant.</p> <p>This parameter is required when you set Image Type to Shared Custom Image.</p> |

vi. Configure the storage settings for the instance.

| Parameter | Required | Description |
|-------------|----------|--|
| System Disk | Yes | <p>Select a disk category from the drop-down list and specify the system disk capacity. Valid values for the disk category: Ultra Disk and Standard SSD.</p> <p>The system disk capacity must range from 20 GiB to 500 GiB.</p> |
| Data Disk | No | <p>You can click Data Disk to create and attach data disks. For each data disk, select a disk category from the drop-down list and specify the disk capacity. Valid values for the disk category: Ultra Disk and Standard SSD.</p> <p>A maximum of 16 data disks can be attached to an instance. The maximum capacity of each data disk is 32 TiB. You can optionally select Release with Instance and Encryption for each data disk.</p> <p>To encrypt a data disk, you can select AES256 or SM4 from the Encryption Method drop-down list and then select a key created in Key Management Service (KMS) from the Encryption Key drop-down list.</p> <p>You can also attach data disks after the instance is created. For more information, see Create a disk.</p> |

vii. Configure the logon password settings for the instance.

| Parameter | Required | Description |
|------------------|----------|---|
| Password Setting | Yes | Specify when to set a password. Valid values: Set Now and Set After Purchase . If you select Set After Purchase , you can use the password reset feature to set a password after the instance is created. For more information, see Change the logon password of an instance . Note The password is used to log on to the instance, not to a VNC management terminal. |
| Logon Password | No | Set the password to be used to log on to the instance. The password must be 8 to 30 characters in length and must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. The supported special characters include <code>()`~!@#\$%^&*-_+={}[];'<>,.?/</code> . Note This parameter is required when you set Password to Set Now . |
| Confirm Password | No | Enter the password again. Note This parameter is required when you set Password to Set Now . |

viii. (Optional) Enter a name for the instance.

The name must be 2 to 128 characters in length and can contain `hyphens (-)`, `underscores (_)`, and `colons (:)`. It must start with a letter and cannot start with `http://` or `https://`.

If you do not specify a name, the system assigns one at random.

6. Click **Submit**.

Result

The new instance appears in the instance list. When the instance is being created, it is in the **Preparing** state. When the instance is created, it enters the **Running** state.

2.4.2. Connect to an instance

2.4.2.1. Instance connecting overview

After an instance is created, you can connect to the instance to perform operations such as installing applications.

You can use one of the following methods to connect to an instance:

- Use remote connection tools to connect to instances that have public IP addresses. For more information about the procedure, see the following topics:
 - [Connect to a Linux-based instance by using SSH commands in Linux or Mac OS X](#)
 - [Connect to a Linux-based instance by using remote connection tools in Windows](#)

- [Connect to a Windows-based instance by using RDP](#)
- Use the VNC feature in the ECS console. For more information, see [Connect to an instance by using a VNC management terminal](#).

The username of a Windows instance is Administrator, and that of a Linux instance is root.

2.4.2.2. Connect to a Linux-based instance by using SSH commands in Linux or Mac OS X

This topic describes how to use SSH commands to connect to a Linux-based instance.

Prerequisites

Create a security group and an instance.

Procedure

1. Enter the following command: `ssh root@instance IP`.
2. Enter the password for the `root` user to log on to the instance.

2.4.2.3. Connect to a Linux-based instance by using remote connection tools in Windows

This topic describes how to connect to an instance by using the PuTTY tool.

Prerequisites

Remote connection tools are designed with similar logics. In this example, PuTTY is used to connect to an instance. Download PuTTY at the following URL: .

Procedure

1. Download and install PuTTY for Windows.
2. Start the PuTTY client and complete the following settings:
 - Host Name (or IP Address): Enter the EIP of the instance to be connected.
 - Port: Select the default port 22.
 - Connection Type: Select SSH.
 - Saved Session: Enter the name of the session. Click **Save**. After the settings are saved, PuTTY remembers the name and IP address of the instance. This eliminates the need to enter them every time you connect to the instance.

3. Click **Open** to connect to the instance.

When you connect to the instance for the first time, PuTTY displays security alerts. Click **Yes** to proceed.

4. Enter the username `root` and press Enter.
5. Enter the password for the instance and press Enter.

If a message similar to the following one appears, a connection to the instance is established.

```
Welcome to aliyun Elastic Compute Server!
```

2.4.2.4. Connect to a Windows instance by using RDC

This topic describes how to connect to a Windows instance by using Remote Desktop Connection (RDC).

Prerequisites

- A security group and a Windows instance are created.
- The instance is in the **Running** state.
- A logon password is set for the instance.
- An Elastic IP address is associated with the instance.
- An inbound security group rule is added to the security group to allow traffic on the RDP port.

| Rule direction | Action | Protocol | Port range | Priority | Authorization type | Authorization object |
|----------------|--------|----------|------------|----------|--------------------|----------------------|
| Inbound | Allow | tcp | 3389/3389 | 1 | IPv4 addresses | 0.0.0.0/0 |

Procedure

1. Use one of the following methods to enable RDC:
 - Click **Start**, enter `mstsc` in the search box, and click `mstsc` in the search result.
 - Press the Windows logo key+R. In the **Run** dialog box that appears, enter `mstsc` and click **OK**.
2. In the **Remote Desktop Connection** dialog box, enter the Elastic IP address of the instance and click **Show Options**.
3. Enter the username.
The default username is administrator.
4. (Optional)If you do not want to enter the password upon subsequent logons, select **Allow me to save credentials**.
5. Click **Connect**.
6. In the **Windows Security** dialog box that appears, enter the password corresponding to the username you entered and click **OK**.

Result

If the Windows desktop appears, a connection to the Windows instance is established.

If an error message is returned indicating that an authentication error has occurred and the function requested is not supported, install CredSSP updates and try again. Follow these steps to install the updates:

1. [Connect to an ECS instance by using the VNC](#).
2. Choose **Start > Control Panel**.
3. Click **System and Security**.
4. Click **Check for updates** in the **Windows Updates** section.
5. If updates are available, click **Install updates**.
6. Restart the instance.

2.4.2.5. Install the certificate for VNC in Windows

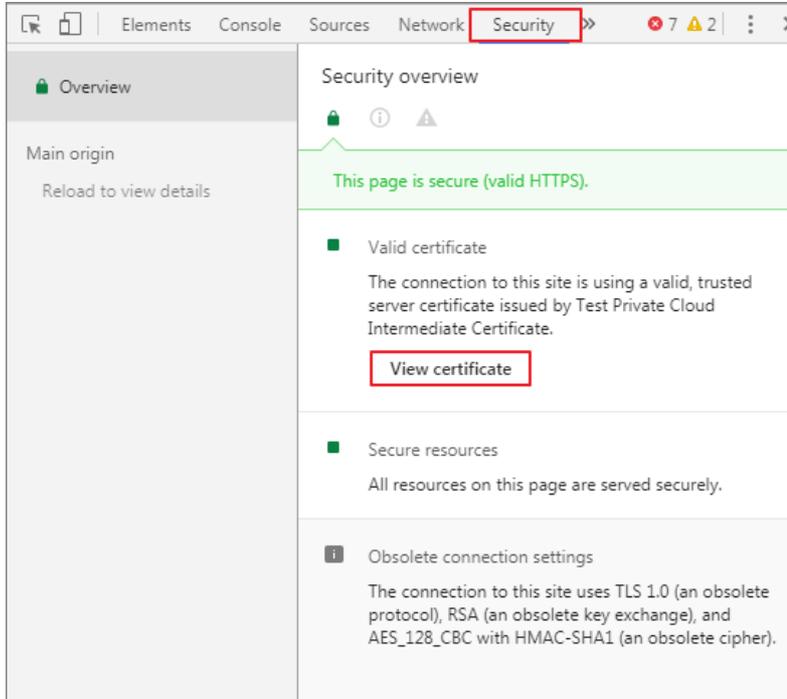
Before you log on to VNC, you must export the relative certificate from the site such as the Apsara Stack Cloud Management (ASCM) console and install the certificate in your local browser.

Context

The VNC feature is provided by the VNC proxy service. The VNC proxy service uses the relative certificate different from that of Apsara Infrastructure Management Framework. The certificate of the VNC proxy service must be manually imported.

Procedure

1. Export the certificate from the ASCM console.
 - i. Log on to the ASCM console. Press the **F12** key or **Fn+F12** to view and select the certificate.
For example, in the Chrome browser, press the **F12** key to open Chrome DevTools.



- ii. In the **Certificate** dialog box, click the **Certificate Path** tab, select the root certificate, and then click **View Certificate**.
 - iii. In the **Certificate** dialog box, click the **Details** tab, and click **Copy to File**.
 - iv. In the **Certificate Export Wizard** dialog box, click **Next**.
 - v. Select **DER encoded binary X.509 (.CER)** as the format and then click **Next**.
 - vi. Click **Browse**, select where to store the certificate, enter a file name, and then click **Save**.
 - vii. Click **Next**.
 - viii. Click **Finish**.
 - ix. Click **OK**.
 2. Install the certificate in your local browser.
 - i. Double-click the certificate.
 - ii. In the **Certificate** dialog box, click **Install Certificate**.
 - iii. In the **Certificate Import Wizard** dialog box, click **Next**.
 - iv. Select **Place all certificates in the following store** and click **Browse**.
 - v. In the **Select Certificate Store** dialog box, select **Trusted Root Certificate Authority** and then click **OK**.
 - vi. In the **Certificate Import Wizard** dialog box, click **Next**.
 - vii. Click **Finish**.
 - viii. If a security warning message is displayed, click **Yes**.
 3. Restart your browser and log on to the ASCM console.
If no security warning message is displayed in the left part of the address bar, the certificate is installed.



2.4.2.6. Connect to an ECS instance by using the VNC

You can access your instance by using the VNC in the ECS console when other SSH clients such as PuTTY, Xshell, and SecureCRT do not work properly.

Prerequisites

- The instance is in the **Running** state.
- The root certificate is imported to your web browser. For more information, see [Install a certificate](#).
- If you log on to an instance for the first time after the instance is created, make sure that you set a new VNC password. For more information, see [Change the VNC password](#).

Context

The VNC password is used to log on to the VNC of the ECS console, and the instance password is used to log on to an instance.

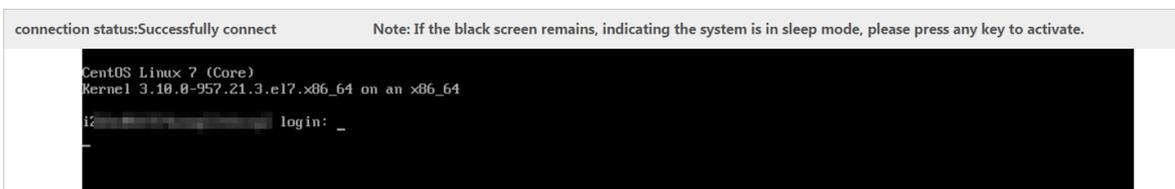
You can use the VNC to connect to an instance to solve issues shown in the following table.

| Scenario | Resolution |
|---|--|
| The instance startup is slowly due to self-check upon startup. | Check the progress of the self check. |
| The firewall of the operating system is enabled by mistake. | Disable the firewall. |
| Abnormal processes appear, which consume large amounts of CPU or bandwidth resources. | Troubleshoot and terminate the abnormal processes. |

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the ECS instance, and click **Remote Connection** in the **Actions** column.
5. Enter the VNC password, and then click **OK**.

After the connection is successful, the logon page is displayed, as shown in the following figure.

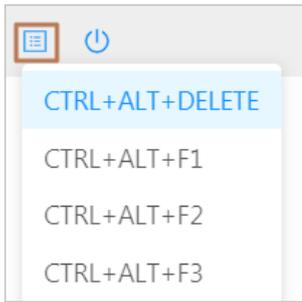


6. Enter the username and password.
 - For Linux instances: Enter the username `root` and the logon password.

Note Passwords in Linux are not displayed as you type. Press the Enter key after you enter the password.

- For Windows instances: To use key combinations such as Ctrl + Alt + Delete, click the corresponding key

combination in the upper-right corner of the VNC page.



Enter the username and password as prompted, and click the Log In icon .

2.4.3. View instances

You can view the list of created instances and the details of individual instances. The details of an instance include basic configurations, disks, snapshots, security groups, and elastic network interfaces (ENIs).

Procedure

1. [Log on to the ECS console.](#)
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region. The created instances that match the specified criteria are displayed.
4. Select a filter option from the drop-down list, enter the relevant information in the search bar, and click **Search**.

You can select multiple filter options to narrow down search results.

| Filter option | Description |
|-------------------|---|
| Instance Name | Enter an instance name to search for the instance. |
| Instance ID | Enter an instance ID to search for the instance. |
| IP Address | Enter the IP address of an instance to search for the instance. |
| VPC ID | Enter a VPC ID to search for the instances that belong to the VPC. |
| Image ID | Enter an image ID to search for the instances that use the image. |
| Status | Select an instance status to search for the instances in that status. Valid values: <ul style="list-style-type: none">◦ Running◦ Stopped◦ Starting◦ Stopping |
| Security Group ID | Enter a security group ID to search for the instances that belong to the security group. |

| Filter option | Description |
|------------------|---|
| Operating System | Enter the name of operating system to search for the instances that use the operating system. |

5. Use one of the following methods to go to the Instance Details page of an instance:
 - In the **Instance ID/Name** column, click the instance ID.
 - Click **Manage** in the **Actions** column corresponding to the instance.
 - In the **Actions** column corresponding to the instance, choose **More > Show Details**.

2.4.4. Modify an instance

You can modify the name, description, and custom data of an existing instance.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the instance and choose **More > Modify** from the **Actions** column.
5. Modify the name, description, and custom data for the instance.

The name must be 2 to 128 characters in length. The description must be 2 to 256 characters in length. The custom data must be 2 to 999 characters in length.
6. Click **OK**.

2.4.5. Stop an instance

You can stop an instance that is not in use. Stopping an instance will interrupt the services that are running on it. Exercise caution when you stop an instance.

Prerequisites

The instance is in the **Running** state.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Use one of the following methods to stop the instance:
 - To stop a single instance, find the instance and choose **More > Instance Status > Stop** in the **Actions** column.
 - To stop one or more instances at a time, select the instances and click **Stop** in the lower-left corner of the **Instances** page.
5. In the message that appears, click **OK**.

Result

In the **Status** column, the instance state changes from **Running** to **Stopping**. After the instance is stopped, its state changes to **Stopped**.

2.4.6. Start an instance

You can start a stopped instance.

Prerequisites

The instance is in the **Stopped** state.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Use one of the following methods to start the instance:
 - To start a single instance, find the instance and choose **More > Instance Status > Start** in the **Actions** column.
 - To start one or more instances at a time, select the instances and click **Start** in the lower-left corner of the Instances page.
5. In the message that appears, click **OK**.

Result

In the **Status** column, the instance state changes from **Stopped** to **Starting**. After the instance is started, its state changes to **Running**.

2.4.7. Restart an instance

You must restart an instance after you change its logon password or install system updates. Restarting an instance will stop the instance and interrupt the services that are running on it for a period of time. Exercise caution when you restart an instance.

Prerequisites

The instance is in the **Running** state.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Use one of the following methods to restart the instance:
 - To restart a single instance, find the instance and choose **More > Instance Status > Restart** in the **Actions** column.
 - To restart one or more instances at a time, select the instances and click **Restart** in the lower-left corner of the Instances page.
5. In the Restart Instance dialog box, select a restart mode.
 - **Restart**: restarts the instance normally.
 - **Force Restart**: forces the instance to restart. This may result in loss of unsaved data.
6. Click **OK**.

2.4.8. Delete an instance

You can delete an instance that is no longer needed to release its resources. Deleted instances cannot be recovered. We recommend that you back up data before you delete an instance. If a data disk is released along with the instance, the data on the disk cannot be recovered.

Prerequisites

The instance is in the **Stopped** state.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Select the instance and click **Delete** in the lower-left corner of the Instances page.
5. In the message that appears, click **OK**.

2.4.9. View the monitoring information of an instance

You can view monitoring charts in the CloudMonitor console to learn about the running conditions of Elastic Compute Service (ECS) instances. This topic describes how to go to the CloudMonitor console to view the monitoring information of an ECS instance.

Context

CloudMonitor provides real-time monitoring, alerting, and notification services for resources to protect your services and business. For more information, see **CloudMonitor overview** in *Apsara Uni-manager Management Console User Guide*.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the ECS instance whose monitoring information you want to view and click **Monitor** in the Monitoring column.
5. On the **Monitoring Charts** page, view the monitoring information of the ECS instance.

2.4.10. Change the instance type of an instance

You can change the instance type of an instance to suit your business needs. This eliminates the need to create a new instance.

Prerequisites

The instance is in the **Stopped** state.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the instance and click **Upgrade/Downgrade** in the **Actions** column.
5. On the Change Specifications page that appears, select a new instance type and click **Submit**.
The instance types available for selection are displayed on the Change Specifications page.
6. Restart the instance by using the console or calling an API operation for the new instance type to take effect.
For more information, see [Start an instance](#) or `StartInstance` in *ECS Developer Guide*.

2.4.11. Change an instance logon password

If you did not set a logon password when creating an instance or forgot the password, you can reset the password.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the instance and select one of the following methods to access the instance details page.
 - In the **Instance ID/Name** column, click the instance ID.
 - Click **Manage** in the **Actions** column.
 - In the **Actions** column, choose **More > Show Details**.
5. Click **Change Password**.
6. Enter and confirm the new password, and then click **OK**.

The password must be 8 to 30 characters in length and must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. The supported special characters include `()'~!@#$%^&*-_+=|{}[]:;<>,.?/`

7. Restart the instance in the console or by calling an API operation to make the new password take effect.
For more information, see [Restart an instance](#) or *the Reboot Instance section* in ECS Developer Guide.

2.4.12. Change the VNC password

If you log on to the VNC for the first time or forget the VNC password, you can reset the password.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the instance and select one of the following methods to access the instance details page.
 - In the **Instance ID/Name** column, click the instance ID.
 - Click **Manage** in the **Actions** column.
 - In the **Actions** column, choose **More > Show Details**.
5. Click **Change VNC Password**.
6. Enter and confirm the new password, and click **OK**.

The password must be 6 characters in length and can contain digits and uppercase and lowercase letters. It does not support special characters.

7. Restart the instance in the console or by calling an API operation to make the new password take effect.
For more information, see [Restart an instance](#) or *the Reboot Instance section* in ECS Developer Guide.

2.4.13. Add an ECS instance to a security group

You can add a created instance to a security group and use security group rules to control network access to the instance.

Context

A security group acts as a virtual firewall and is used to provide security isolation. A security group controls access to ECS instances.

Instances that belong to the same account and are in the same region and in the same security group can communicate with each other over the internal network. If instances that belong to the same account in the same region are in different security groups, you can implement internal network communication by authorizing mutual access between two security groups.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > Security Groups**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the security group and click **Manage Instances** in the **Actions** column.
5. Click **Add Instance**.
6. Select the instance and click **OK**.

An instance can be added to five security groups. After an instance is added, the security group rules automatically apply to the instance.

2.4.14. Customize instance data

ECS allows you to run the instance customization script upon startup and import data into instances.

Context

The instance data customization feature is applicable to both Windows and Linux instances. It allows you to:

- Run the instance customization script upon startup.
- Import data into instances.

Usage instructions

- **Limits**

The instance data customization feature can only be used when an instance meets all the following requirements:

- Network type: VPC
- Image: a system image or a custom image that is inherited from the system image
- Operating system: one type included in [Supported operating systems](#) Supported operating systems

| Windows | Linux |
|--|---|
| <ul style="list-style-type: none"> ▪ Windows Server 2016 64-bit ▪ Windows Server 2012 64-bit ▪ Windows Server 2008 64-bit | <ul style="list-style-type: none"> ▪ CentOS ▪ Ubuntu ▪ SUSE Linux Enterprise ▪ OpenSUSE ▪ Debian ▪ Aliyun Linux |

- When you configure instance data customization scripts, you must enter custom data based on the type of operating system and script.

 **Note** Only English characters are allowed.

- If your data is Base64 encoded, select **Enter Base64 Encoded Information**.

 **Note** The size of the customization script cannot exceed 16 KB before the data is Base64 encoded.

- For Linux instances, the script format must meet the requirements described in [Types of Linux instance customization scripts](#).
- For Windows instances, the script can only start with `[bat]` or `[powershell]`.
- After starting an instance, run a command to view the following information:
 - Execution result of the instance customization script
 - Data imported to instances
- **Console:** You can modify the custom instance data in the console. Whether the modified instance customization script needs to be re-executed depends on the script type. For example, if the `bootcmd` script in Cloud Config is modified for Linux instances, the script is automatically executed each time instances are restarted.
- **OpenAPI:** You can also use OpenAPI to customize instance data. For more information, see [Create Instance](#) and [Modify Instance Attribute](#) in *ECS Developer Guide*.

Linux instance data customization scripts

Linux instance data customization scripts provided by Alibaba Cloud are designed based on the cloud-init architecture. They are used to automatically configure parameters of Linux instances. Customization script types are compatible with the cloud-init.

Description of Linux instance data customization scripts

- Linux instance customization scripts are executed after instances are started and before `/etc/init` is executed.
- Linux instance customization scripts can only be executed with root permissions by default.

Types of Linux instance customization scripts

- **User-Data Script**
 - Description: A script, such as shell script, is used to customize data.
 - Format: The first line must include `#!`, such as `#!/bin/sh`.
 - Limit: The script size (including the first line) cannot exceed 16 KB before the data is Base64 encoded.
 - Frequency: The script is executed only when instances are started for the first time.
 - Example:

```
#!/bin/sh
echo "Hello World. The time is now $(date -R)!" | tee /root/output10.txt
```

- **Cloud Config Data**
 - Description: Predefined data is used to configure services, such as specifying yum sources or importing SSH keys.
 - Format: The first line must be `#cloud-config`.
 - Limit: The script size (including the first line) cannot exceed 16 KB before the data is Base64 encoded.
 - Frequency: The script execution frequency varies with the specific service.

- Example:

```
#cloud-config
apt:
  primary:
    - arches: [default]
    uri: http://us.archive.ubuntu.com/ubuntu/
```

- **Include**

- Description: The configuration content can be saved in a text file and imported into cloud-init as a URL.
- Format: The first line must be `#include`.
- Limit: The script size (including the first line) cannot exceed 16 KB before the data is Base64 encoded.
- Frequency: The script execution frequency depends on the script type in the text file.
- Example:

```
#include
http://ecs-image-test.oss-cn-hangzhou.aliyuncs.com/userdata/cloudconfig
```

- **GZIP format**

- Description: Cloud-init limits the size of customization scripts to 16 KB. You can compress and import the script file into the customization script if the file size exceeds 16 KB.
- Format: The .gz file is imported into the customization script as a URL in `#include`.
- Frequency: The script execution frequency depends on the script content contained in the GZIP file.
- Example:

```
#include
http://ecs-image-test.oss-cn-hangzhou.aliyuncs.com/userdata/config.gz
```

View the custom data of a Linux instance

Run the following command in the instance:

```
curl http://100.100.100.200/latest/user-data
```

Windows instance customization scripts

Windows instance customization scripts independently developed by Alibaba Cloud can be used to initialize Windows instances.

There are two types of Windows instance customization scripts:

- Batch processing program: starts with `[bat]` and serves as the first line. The script size must be smaller than 16 KB before the data is Base64 encoded.
- PowerShell script: starts with `[powershell]` and serves as the first line. The script size must be smaller than 16 KB before the data is Base64 encoded.

View the custom data of a Windows instance

Run the following PowerShell command in the instance:

```
Invoke-RestMethod http://100.100.100.200/latest/user-data/
```

2.4.15. Modify a private IP address

Each instance is assigned a private NIC and bound with a private IP address. You can modify the private IP address of the instance. The private IP address you use must be within the CIDR block of the VSwitch to which the instance belongs and cannot be used by another instance.

Prerequisites

The instance is in the **Stopped** state.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the instance and choose **More > Change Private IP Address** from the **Actions** column.
5. Enter a new private IP address and click **OK**.

The private IP address you can use must be within the CIDR block of the VSwitch to which the instance belongs and cannot be used by another instance or for a specific purpose.

For example, if the CIDR block of the VSwitch is 192.168.1.0/24, you can use an IP address in the range of 192.168.1.1 to 192.168.1.254. The first address 192.168.1.0 identifies the subnet itself, and the last address 192.168.1.255 identifies the broadcast address. Both addresses are reserved and cannot be used.

2.4.16. Install the CUDA and GPU drivers for a Linux instance

You must install a GPU driver on GPU instances to use the GPU. If the image you use does not contain a pre-installed GPU driver, you must manually install the CUDA and GPU drivers for the instance.

Prerequisites

If your instance cannot connect to the Internet, the installation file cannot be downloaded. You can install an FTP client on the instance to transfer the installation file to the instance.

Context

When installing NVIDIA drivers, you must install the kernel package that contains the kernel header file before you install the CUDA and GPU drivers on the instance.

Procedure

1. Install the kernel package.
 - i. Run the `uname -r` command to view the current kernel version.

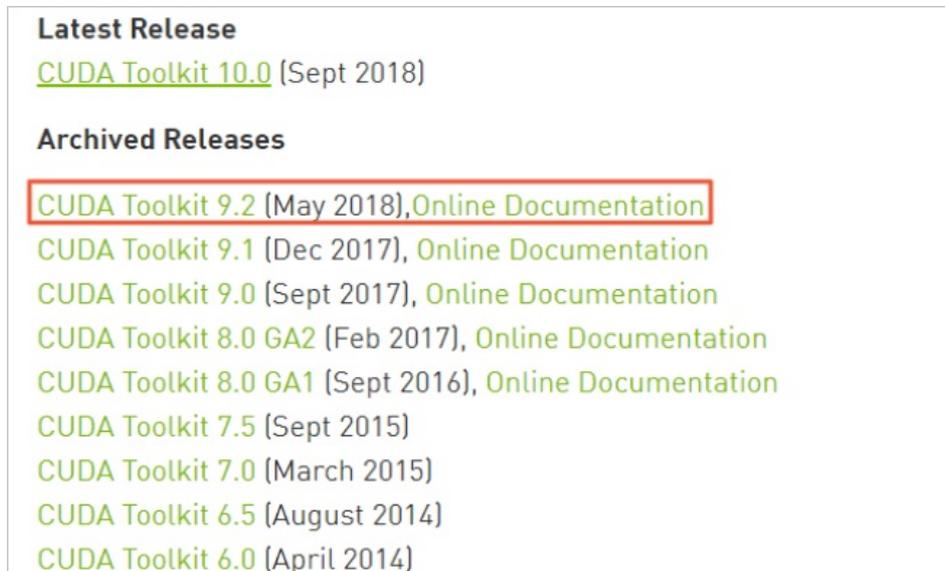
A similar output is displayed:

 - CentOS: 3.10.0-862.14.4.el7.x86_64
 - Ubuntu: 4.4.0-117-generic
 - ii. Copy the kernel package of the corresponding version to the instance and install the package.
 - CentOS: Copy the RPM package of the `kernel-devel` component and run the `rpm -ivh 3.10.0-862.14.4.el7.x86_64.rpm` command to install the package. `3.10.0-862.14.4.el7.x86_64.rpm` is used as an example. Replace it with the actual package name.
 - Ubuntu: Copy the DEB package of the `linux-headers` component and run the `dpkg -i 4.4.0-117-generic.deb` command to install the package. `4.4.0-117-generic.deb` is used as an example. Replace it with the actual package name.
2. Download the CUDA Toolkit.

- i. Access the [official CUDA download page](#). Choose the version based on the GPU application requirements for CUDA.

This example uses [CUDA Toolkit 9.2](#).

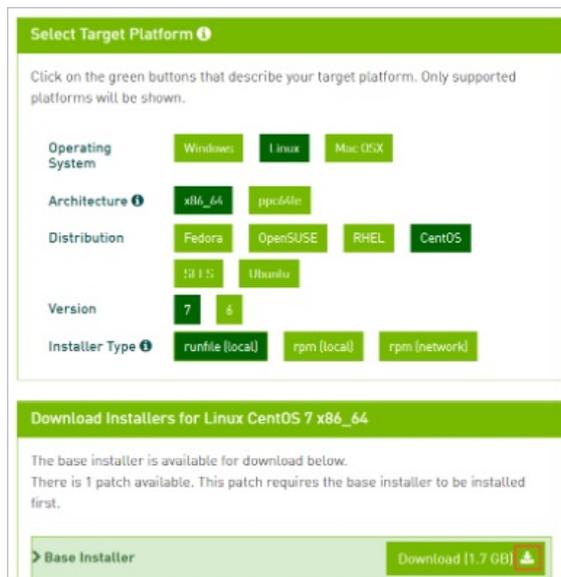
Download the CUDA Toolkit



- ii. Choose a platform based on your operating system. Select **Installer Type** to **runfile (local)** and click **Download**.

NVIDIA drivers are already included in the CUDA Toolkit.

Download the drivers



3. Copy the downloaded `cuda_9.2.148_396.37_linux.run` file to the instance. `cuda_9.2.148_396.37_linux.run` is used as an example. Replace it with the actual file name.
4. Run the `sudo sh ./cuda_9.2.148_396.37_linux.run --silent --verbose --driver --toolkit --samples` command to install the CUDA driver. `cuda_9.2.148_396.37_linux.run` is used as an example. Replace it with the actual file name.

The installation takes about 10 to 20 minutes. When `Driver: Installed` is displayed, the installation is successful.

View the CUDA installation result

```

=====
= Summary =
=====
Driver: Installed
Toolkit: Installed in /usr/local/cuda-9.2
Samples: Installed in /home/lb164654, but missing recommended libraries

Please make sure that
- PATH includes /usr/local/cuda-9.2/bin
- LD_LIBRARY_PATH includes /usr/local/cuda-9.2/lib64, or, add /usr/local/cuda-9.2/lib64 to /etc/ld.so.conf and run ldconfig as root

To uninstall the CUDA Toolkit, run the uninstall script in /usr/local/cuda-9.2/bin
To uninstall the NVIDIA Driver, run nvidia-uninstall

Please see CUDA_Installation_Guide_Linux.pdf in /usr/local/cuda-9.2/doc/pdf for detailed information on setting up CUDA.

Logfile is /tmp/cuda_install_19765.log

```

5. Run the `nvidia-smi` command to view the GPU driver status. If the output displays the details of the GPU driver, the driver is running properly.

View the GPU driver status

```

$ nvidia-smi
Mon Oct 15 19:05:00 2018

+-----+-----+
| NVIDIA-SMI 396.37                Driver Version: 396.37          |
+-----+-----+
| GPU  Name      Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+
|   0   Tesla P4             Off | 00000000:00:08.0 Off |                    0 |
| N/A   28C    P0      23W / 75W |  0MiB / 7611MiB |      0%   Default |
+-----+-----+

+-----+-----+
| Processes:                         GPU Memory          |
|  GPU       PID    Type   Process name           Usage              |
+-----+-----+
| No running processes found         |
+-----+-----+

```

What's next

If you want to run the OpenGL program, you must first purchase the licenses and install the GRID drivers. For information about the installation procedure, see the official NVIDIA documentation.

2.4.17. Install the CUDA and GPU drivers for a Windows instance

You must install a GPU driver on GPU instances to use the GPU. If the image you use does not contain a pre-installed GPU driver, you must manually install the CUDA and GPU drivers for the instance.

Prerequisites

- If your instance cannot connect to the Internet, the installation file cannot be downloaded. You can install an FTP client on the instance to transfer the installation file to the instance.
- To compile CUDA programs, first install a Windows compiling environment, such as Visual Studio 2015. If you do not need to compile CUDA programs, ignore it.

Procedure

1. Download the CUDA Toolkit.

- i. Access the [official CUDA download page](#). Choose the version based on the GPU application requirements for CUDA.
This example uses [CUDA Toolkit 9.2](#).
 - ii. Choose a platform based on your operating system. Set **Installer Type** to **exe (local)** and click **Download**.
NVIDIA drivers are already included in the CUDA Toolkit.
2. Copy the downloaded `cuda_9.2.148_windows.exe` file to the instance. `cuda_9.2.148_windows.exe` is used as an example. Replace it with the actual file name.
 3. Double-click `cuda_9.2.148_windows.exe` and follow the installation wizard to install the CUDA driver. `cuda_9.2.148_windows.exe` is used as an example. Replace it with the actual file name.
The installation takes about 10 to 20 minutes. When `Installed: - Nsight Monitor and HUD Launcher` is displayed, the driver is installed.
 4. Press `Win + R` and enter `devmgmt.msc`.
The NVIDIA device is displayed in **Display Adapter**.
 5. Press `Win + R`, enter `cmd`, and run the `"C:\Program Files\NVIDIA Corporation\NVSMI\nvidia-smi"` command.
If the output displays the details of the GPU driver, the driver is running properly.

What's next

If you want to run the OpenGL and DirectX programs, you must first purchase the licenses and install the GRID drivers. For information about the installation procedure, see the official NVIDIA documentation.

2.5. Disks

2.5.1. Create a disk

You can create an independent data disk and attach it to an ECS instance to increase the storage space of the instance. This topic describes how to create an empty data disk. You cannot create independent system disks.

Context

We recommend that you plan the number and size of data disks before you create them. The following limits apply to data disks:

- A maximum of 16 data disks can be attached to an instance. Disks and Shared Block Storage devices share this quota.
- Each Shared Block Storage device can be attached to up to four ECS instances at the same time.
- Each ultra disk, shared ultra disk, standard SSD, or shared SSD can have a maximum capacity of 32 TiB.
- Disks cannot be combined in ECS. They are independent of each other. You cannot combine multiple disks into one by formatting them.

We recommend that you do not use Logical Volume Manager (LVM) to create logical volumes across multiple disks, because a snapshot can only back up data of a single disk. If you create a logical volume across several disks, data discrepancies will occur when you restore these disks.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Storage > Disks**.
3. Click **Create Disk**.
4. Configure parameters listed in the following table to create a disk.

| Section | Parameter | Required | Description |
|----------------|----------------------|----------|--|
| Region | Organization | Yes | Select an organization in which to create the disk. |
| | Resource Set | Yes | Select a resource set in which to create the disk. |
| | Region | Yes | Select a region in which to create the disk. |
| | Zone | Yes | Select a zone in which to create the disk. |
| Basic Settings | Name | Yes | Enter a name for the disk. The name must be 2 to 128 characters in length. It must start with a letter and cannot start with http:// or https://. It can contain letters, digits, periods (.), underscores (_), hyphens (-), colons (:), and commas (,). |
| | Specifications | Yes | Select a disk category and specify the disk size. Valid disk categories: <ul style="list-style-type: none"> ◦ SSD Disk ◦ Ultra Disk ◦ Shared SSD: shared SSD ◦ Shared Ultra Disk: shared ultra disk The disk size must range from 20 GiB to 32768 GiB. |
| | Encrypt | No | Specify whether to encrypt the disk. |
| | Encryption Algorithm | No | Select an encryption algorithm. This parameter is required when you set Encrypt to Yes . Valid values: <ul style="list-style-type: none"> ◦ AES256: AES256 encryption algorithm. ◦ SM4: Chinese encryption algorithm SM4. |

| Section | Parameter | Required | Description |
|---------|----------------|----------|--|
| | Encryption Key | No | Select an encryption key. This parameter is required when you set Encrypt to Yes .  Note If no key is available, create a key in KMS. |
| | Use Snapshot | No | Specify whether to create the disk from a snapshot. If you select Yes , you must specify a snapshot. The size of the created disk depends on the size of the specified snapshot. <ul style="list-style-type: none"> ◦ If the disk size that you specify is greater than the snapshot size, the disk will be created with the size you specified. ◦ If the disk size that you specify is smaller than the snapshot size, the disk will be created with the snapshot size. |

5. Click **Submit**.

Result

The created disk is displayed in the disk list and in the **Pending** state.

What's next

After the disk is created, you must attach the disk to an instance and partition and format the disk. For more information, see the following topics:

- [Attach a disk](#)
- [Format a data disk for a Linux instance](#)
- [Format a data disk of a Windows instance](#)

2.5.2. Attach a disk

You can attach a disk that is created separately to an ECS instance as a data disk. The disk and the instance must be in the same region and the same zone.

Prerequisites

The disk is in the **Pending** state.

Context

- You do not need to attach data disks that are created at the same time as an instance.

- A disk can only be attached to an instance that is in the same zone and region as the disk.
- You can attach a disk to a single ECS instance at a time.
- A Shared Block Storage device can be attached to a maximum of four ECS instances at the same time.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Storage > Disks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the disk and choose **More > Attach** from the **Actions** column.
5. Specify the destination instance and configure the release mode as needed.
 - If you select **Release Disk with Instance**, the disk is released together when the instance it is attached to is deleted.
 - If you do not select **Release Disk with Instance**, the disk changes to the **Pending** state when the instance it is attached to is deleted.
6. Click **OK**.

2.5.3. Partition and format disks

2.5.3.1. Format a data disk for a Linux instance

Data disks created separately are not partitioned or formatted. This topic describes how to partition and format a data disk of a Linux instance.

Prerequisites

The disk has been attached to the instance.

Procedure

1. [Connect to the instance](#).
2. Run the `fdisk -l` command to view all data disks attached to the ECS instance.

If `/dev/vdb` is not displayed in the command output, the ECS instance does not have a data disk. Check whether the data disk is attached to the instance.

```
[root@iZ*****eZ ~]# fdisk -l
Disk /dev/vda: 42.9 GB, 42949672960 bytes
255 heads, 63 sectors/track, 5221 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00078f9c

Device Boot      Start         End      Blocks   Id  System
/dev/vda1 *        1      5222    41940992   83  Linux
Disk /dev/vdb: 21.5 GB, 21474836480 bytes
16 heads, 63 sectors/track, 41610 cylinders
Units = cylinders of 1008 * 512 = 516096 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00000000
```

3. Create partitions for the data disk.
 - i. Run the `fdisk /dev/sdb` command.
 - ii. Enter `n` to create a new partition.

- iii. Enter *p* to set the partition as the primary partition.
- iv. Enter a partition number and press the Enter key. In this example, 7 is entered to create Partition 1.
- v. Enter the number of the first available sector. This example uses the default value. This is done by pressing the Enter key. You can also enter a value from 1 to 41610 and press the Enter key.
- vi. Enter the number of the last sector. This example uses the default value. This is done by pressing the Enter key. You can also enter a value from 1 to 11748 and press the Enter key.
- vii. (Optional)Optional. To create multiple partitions, repeat steps b through f until all four primary partitions are created.
- viii. Run the **wq** command to start partitioning.

```
[root@iZ*****eZ ~]# fdisk /dev/vdb
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel
Building a new DOS disklabel with disk identifier 0x01ac58fe.
Changes will remain in memory only, until you decide to write them.
After that, of course, the previous content won't be recoverable.
Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)
WARNING: DOS-compatible mode is deprecated. It's strongly recommended to
switch off the mode (command 'c') and change display units to
sectors (command 'u').
Command (m for help): n
Command action
e extended
p primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-41610, default 1):
Using default value 1
Last cylinder, +cylinders or +size[K,M,G] (1-41610, default 41610):
Using default value 41610
Command (m for help): wq
The partition table has been altered!
```

4. Run the **fdisk -l** command to view the partitions.

If `/dev/vdb1` is displayed in the command output, new partition `vdb1` is created.

```
[root@iZ*****eZ ~]# fdisk -l
Disk /dev/vda: 42.9 GB, 42949672960 bytes
255 heads, 63 sectors/track, 5221 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00078f9c
Device Boot Start End Blocks Id System
/dev/vda1 * 1 5222 41940992 83 Linux
Disk /dev/vdb: 21.5 GB, 21474836480 bytes
16 heads, 63 sectors/track, 41610 cylinders
Units = cylinders of 1008 * 512 = 516096 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x01ac58fe
Device Boot Start End Blocks Id System
/dev/vdb1 1 41610 20971408+ 83 Linux
```

5. Format the new partition. In this example, the new partition is formatted as ext3 after you run the **mkfs.ext3 /dev/vdb1** command.

The time required for formatting depends on the disk size. You can also format the new partition to another file system. For example, you can run the **mkfs.ext4 /dev/vdb1** command to format the partition as ext4.

Compared with ext2, ext3 only adds the log function. Compared with ext3, ext4 improves on some important data structures. ext4 provides better performance and reliability, and more functions.

```
[root@iZ*****eZ ~]# mkfs.ext3 /dev/vdb1
mke2fs 1.41.12 (17-May-2010)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
1310720 inodes, 5242852 blocks
262142 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
160 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
 32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
 4096000
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
This filesystem will be automatically checked every 25 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
```

6. Run the `echo '/dev/vdb1 /mnt ext3 defaults 0 0' >> /etc/fstab` command to write the information of the new partition to the `/etc/fstab` file. You can run the `cat /etc/fstab` command to view the new partition information.

Ubuntu 12.04 does not support barriers. To write the information of the new partition into the `/etc/fstab` file, you must run the `echo '/dev/vdb1 /mnt ext3 barrier=0 0 0' >> /etc/fstab` command.

In this example, the partition information is added to the ext3 file system. You can also modify the ext3 parameter to add the partition information to another type of file system.

To attach the data disk to a specific folder, for example, to store web pages, modify the `/mnt` part of the preceding command.

```
[root@iZ*****eZ ~]# echo '/dev/vdb1 /mnt ext3 defaults 0 0' >> /etc/fstab
[root@iZbp19cdhgdj0aw5r2izleZ ~]# cat /etc/fstab
#
# /etc/fstab
# Created by anaconda on Thu Aug 14 21:16:42 2014
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
UUID=94e4e384-0ace-437f-bc96-057dd64f**** / ext4 defaults,barrier=0 1 1
tmpfs      /dev/shm      tmpfs defaults 0 0
devpts     /dev/pts      devpts gid=5,mode=620 0 0
sysfs      /sys          sysfs defaults 0 0
proc       /proc         proc defaults 0 0
/dev/vdb1 /mnt ext3 defaults 0 0
```

7. Mount the new partitions. Run the `mount -a` command to mount all the partitions listed in `/etc/fstab` and run the `df -h` command to view the result.

If the following information is displayed, the new partitions are mounted and available for use.

```
[root@iZ*****eZ ~]# mount -a
[root@iZ*****eZ ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/vda1       40G  5.6G  32G  15% /
tmpfs           499M  0  499M   0% /dev/shm
/dev/vdb1       20G  173M  19G   1% /mnt
```

2.5.3.2. Format a data disk of a Windows instance

Data disks created separately are not partitioned or formatted. This topic describes how to partition and format a data disk of a Windows instance. This example uses Windows Server 2008.

Prerequisites

The disk has been attached to an instance.

Procedure

1. In the lower-left corner of the screen, click the **Server Manager** icon.
2. In the left-side navigation pane of the Server Manager window, choose **Storage > Disk Management**.
3. Right-click an empty partition and select **New Simple Volume** from the short cut menu.
If the disk status is **Offline**, change it to **Online**.
4. Click **Next**.
5. Set the size of the simple volume, which is the partition size. Then click **Next**.
The default value is the maximum value of the disk space. You can specify the partition size as needed.
6. Specify the drive letter and then click **Next**.
7. Specify the formatting options and then click **Next**.
We recommend that you format the partition with the default settings provided by the wizard.
8. When the wizard prompts that the partition has been completed, click **Finish** to close the wizard.

2.5.4. View disks

You can view the list of created disks and the details of individual disks.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Storage > Disks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
The created disks that match the specified criteria are displayed.
4. Select a filter option from the drop-down list, enter the relevant information in the search box, and click **search**.

You can select multiple filter options to narrow down search results.

| Filter option | Description |
|---------------|---|
| Disk Name | Enter a disk name to search for the disk. |
| Disk ID | Enter a disk ID to search for the disk. |
| Instance ID | Enter an instance ID to search for the disks that are attached to the instance. |

| Filter option | Description |
|-------------------|--|
| Disk Status | <p>Select a disk status to search for disks in that status. Valid values:</p> <ul style="list-style-type: none"> ◦ Running ◦ Pending ◦ Attaching ◦ Detaching ◦ Creating ◦ Deleting ◦ Deleted <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin: 10px 0;"> <p> Note Deleted disks are no longer displayed in the disk list.</p> </div> <ul style="list-style-type: none"> ◦ Initializing ◦ All Statuses |
| Disk Properties | <p>Select a disk type to search for disks of that type. Valid values:</p> <ul style="list-style-type: none"> ◦ All ◦ System Disk ◦ Data Disk |
| Policy ID | Enter the ID of an automatic snapshot policy to search for the disks that use the policy. |
| Encryption Key ID | Enter the ID of an encryption key to search for the disks that are encrypted with the key. |

5. In the **Disk ID/Name** column, click a disk ID to go to the Disk Details page of the disk. The properties and mount information of the disk are displayed on the Disk Details page.

2.5.5. Roll back a disk

If you have created snapshots for a disk, you can use a snapshot to roll back the disk to the state it was when the snapshot was taken. Rolling back a disk is irreversible. Once the disk is rolled back, the disk data before the rollback time cannot be restored. Exercise caution when you perform this operation.

Prerequisites

- Snapshots have been created for the disk.
- The disk is not released.
- The instance where the target disk resides must be in the **Stopped** state.

Procedure

1. [Log on to the ECS console.](#)
2. In the left-side navigation pane, choose **Snapshots and Images > Snapshots**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Select a filtering option from the drop-down list, enter the relevant information in the search bar, and click **Search**.

You can select multiple filtering options to narrow down the search results.

| Filtering option | Description |
|------------------|---|
| Snapshot Name | Enter a snapshot name to search for the snapshot. |
| Snapshot ID | Enter a snapshot ID to search for the snapshot. |
| Instance ID | Enter an instance ID to search for the snapshots related to the instance. |
| Disk ID | Enter a disk ID to search for the snapshots related to the disk. |
| Snapshot Type | Select a snapshot type to search for the snapshots of that type. Options include: <ul style="list-style-type: none"> ◦ All ◦ User Snapshots: manual snapshots. ◦ Automatic snapshots: automatic snapshots. |
| Creation Time | Enter a creation time to search for the snapshots that were created at that time. |

5. Find the snapshot and click **Restore** in the **Actions** column.
6. Click **OK**.

2.5.6. Modify the disk properties

You can modify the properties of a created disk, such as changing the settings of the Release Disk with Instance and Release Automatic Snapshots with Disk options.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Storage > Disks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the disk and choose **More > Modify Disk Properties** from the **Actions** column.
5. Modify the release mode.
 - **Release Disk with Instance**: When this option is selected, the disk is released together when the instance it is attached to is deleted. When this option is not selected, the disk changes to the **Pending** state when the instance it is attached to is deleted.
 - **Release Automatic Snapshots with Disk**: When this option is selected, the automatic snapshots created for the disk is released together when the disk is deleted. When this option is not selected, the automatic snapshots are retained.
6. Click **OK**.

2.5.7. Modify the disk description

You can modify the name and description of a created disk.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Storage > Disks**.

3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the disk and choose **More > Modify Disk Description** from the **Actions** column.
5. Modifies the name and description of the disk.
 - The name must be 2 to 128 characters in length and start with a letter. It can contain periods (.), underscores (_), colons (:), and hyphens (-).
 - The description must be 2 to 256 characters in length and cannot start with http:// and https://.
6. Click **OK**.

2.5.8. Expand a disk

You can expand system or data disks online. After a disk is expanded, you do not need to restart the instance to which the disk is attached for the new disk capacity to take effect.

Prerequisites

- To avoid data loss, we recommend that you create a snapshot to back up disk data before you expand a disk. For more information, see [Create a snapshot](#).
- No snapshot is being created for the disk to be expanded.
- The disk or the instance to which the disk is attached meets the following requirements:
 - If the disk is a system disk, the instance is in the **Running** state.
 - If the disk is a data disk, one of the following requirements is met:
 - The disk is in the **Pending** state.
 - If the disk is attached to an instance, the instance is in the **Running** state.
 - If the disk is a Shared Block Storage device, it is in the **Pending** state.

Context

The following limits apply when you expand a disk.

| Limit | Description |
|-------------------|--|
| Disk category | <ul style="list-style-type: none"> • Ultra disks and standard SSDs can be expanded. • Shared SSDs and shared ultra disks can be expanded. |
| Operating system | The system disks of Windows Server 2003 instances cannot be expanded. |
| Partitioning mode | You cannot expand a data disk that uses the MBR partitioning scheme to more than 2 TiB. To expand this kind of disk to more than 2 TiB, we recommend that you create and attach a new data disk with the desired size. Use the GPT partitioning scheme to partition the new data disk and then copy data from the original data disk to the new data disk. |
| File system | For Windows instances, you can expand only disks that use NTFS file systems. |
| Maximum capacity | <ul style="list-style-type: none"> • Ultra disk and standard SSD: 32,768 GiB • Shared SSD and shared ultra disk: 32,768 GiB |
| Operations | <ul style="list-style-type: none"> • When you expand disks, only the capacity of the disks is expanded. The sizes of partitions and file systems do not change. You must manually re-allocate storage space on a disk after the disk is expanded. • You cannot shrink an expanded disk by any means, such as by rolling it back. |

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Storage > Disks**.
3. Find the target disk and choose **More > Expand Disk** in the **Actions** column.
4. In the Expand Disk dialog box that appears, specify a new capacity for the disk.
The new capacity must be greater than the current capacity.
5. Click **OK**.

Result

When you expand disks, only the capacity of the disks is expanded. The sizes of partitions and file systems do not change. You must manually re-allocate storage space on a disk after the disk is expanded.

2.5.9. Encrypt a disk

2.5.9.1. Encrypt a system disk

In the scenarios that require data security and regulatory compliance, you can encrypt disks to secure your data stored on the disks. To encrypt system disks, you can encrypt custom images and then use the encrypted custom images to create instances. The system disks of the created instances are automatically encrypted.

Context

You can encrypt system disks only by encrypting custom images.

Step 1: Create a custom image from an instance

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Snapshots and Images > Snapshots**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the snapshot from which you want to create a custom image and click **Create Custom Image** in the **Actions** column.
5. Configure the parameters listed in the following table to create a custom image.

| Parameter | Description |
|--------------------------|--|
| Custom Image Name | Enter a name for the custom image. The name must be 2 to 128 characters in length and can contain letters, digits, periods (.), underscores (_), hyphens (-), and colons (:). It must start with a letter. |
| Sharing Scope | Select the scope in which to share the custom image. <ul style="list-style-type: none"> ◦ Current Organization and Subordinate Organizations ◦ Current Resource Set ◦ Current Organization |
| Description | Enter a description for the custom image. The description must be 2 to 256 characters in length and cannot start with http:// or https://. |

6. Click **OK**.

Step 2: Encrypt the custom image

1. [Log on to the ECS console.](#)
2. In the left-side navigation pane, choose **Snapshots and Images > Images**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Click the **Custom Images** tab.
5. Find the custom image that you want to encrypt and click **Encrypt Image** in the **Actions** column.
6. In the **Encrypt Image** dialog box, configure the parameters listed in the following table.

| Parameter | Description |
|--------------------------|---|
| Image ID | The system automatically obtains the ID of the image to be encrypted. You do not need to configure this parameter. |
| Custom Image Name | Enter a name for the new encrypted custom image. The name must be 2 to 128 characters in length and can contain letters, digits, underscores (_), periods (.), and hyphens (-). It must start with a letter. |
| Description | Enter a description for the new encrypted custom image. The description must be 2 to 256 characters in length and cannot start with http:// or https://. |

7. Click **OK**.

Step 3: Use the encrypted custom image to create an instance

1. [Log on to the ECS console.](#)
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Click **Create Instance**.
5. Configure the parameters for the instance.

For more information about how to configure these parameters, see [Create an instance](#)

In the **Image** section, set Image Type to **Custom Image**. Then, select the image that you encrypted from the **Custom Image** drop-down list.

6. Click **Submit**.

Result

After the instance is created, you can click its ID to go to the **Instance Details** page. Then, you can click the **Disks** tab and check the value in the **Disk Encryption** column corresponding to the system disk. If the value is **Yes**, the system disk is encrypted.

2.5.9.2. Encrypt a data disk

In the scenarios that require data security and regulatory compliance, you can encrypt disks to secure your data stored on the disks. After a data disk is created, you cannot change its encryption state. If you want to encrypt a data disk, enable encryption for the disk when you create it.

Context

We recommend that you determine the number and sizes of data disks before you create them. The following limits apply to data disks:

- A maximum of 16 data disks can be attached to an instance. Disks and Shared Block Storage devices share this quota.

- Each Shared Block Storage device can be attached to up to four ECS instances at the same time.
- Each ultra disk, standard SSD, Ultra Shared Block Storage device, or SSD Shared Block Storage device can have a maximum capacity of 32 TiB.
- Disks cannot be combined in ECS. They are independent of each other. You cannot combine multiple disks into one by formatting them.

We recommend that you do not use Logical Volume Manager (LVM) to create logical volumes across multiple disks, because a snapshot can back up data only of a single disk. If you create a logical volume across several disks, data discrepancies may occur when you restore these disks.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Storage > Disks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Click **Create Disk**.
5. On the **Create Disk** page, configure the parameters for the disk.

When you encrypt the disk, take note of the following parameters:

- **Encrypted**: Select **Yes**.
- **Encryption Method**: Select an encryption algorithm.
 - **AES256**: indicates the AES256 encryption algorithm.
- **Encryption Key**: Select an encryption key.

For information about how to configure the other parameters to create a disk, see [Create a disk](#).

6. Click **Submit**.

2.5.10. Reinitialize a disk

You can reinitialize a disk to restore it to its initial state.

Prerequisites

- The disk is in the **Running** state.
- The instance is in the **Stopped** state.
- After a disk is reinitialized, its data is lost and cannot be recovered. Exercise caution when you perform this operation. We recommend that you back up the data of the disk or create snapshots before you reinitialize the disk. For more information, see [Create a snapshot](#).

Context

The result of disk reinitialization depends on the disk type and how the disk is created.

- System disk:
 - The disk is restored to the initial state of the image used by the disk.
 - If the original image is deleted, the disk cannot be reinitialized.
- Data disk:
 - If the disk is empty when created, the disk is restored to an empty disk.
 - If the disk is created from a snapshot, the disk is restored to a disk with only the data of the source snapshot.
 - If the disk is created from a snapshot and the snapshot is deleted, the disk cannot be reinitialized.

Procedure

1. [Log on to the ECS console](#).

2. In the left-side navigation pane, choose **Storage > Disks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the data disk and click **Reinitialize** in the **Actions** column.
5. Perform the following operations based on the disk type.
 - For a system disk, enter and confirm a new logon password, select the **Start Instance After Reinitializing Disk** option as needed, and then click **OK**.

The password must be 8 to 30 characters in length and must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. The supported special characters include `()'~!@#$$%^&*-_+={}|[]:;<>,.?/`
 - For a data disk, click **OK**.

Result

When the disk is being reinitialized, the disk enters the **Initializing** state. After the reinitialization, it changes to the **Running** state.

2.5.11. Detach a data disk

You can detach a data disk, not a system disk.

Prerequisites

- For a Windows instance, you must bring the data disk offline in Disk Management.

 **Note** To guarantee data integrity, we recommend that you stop read/write operations on the data disk when you detach the disk. Otherwise, data may be lost.

- For a Linux instance, you must connect to the instance and unmount the partitions on the disk.

 **Note** If you have configured the `/etc/fstab` file to automatically mount the disk partitions upon instance startup, you must delete the mounting information from the `/etc/fstab` file before you detach the disk. Otherwise, you cannot connect to the instance after the instance is restarted.

- The data disk to be detached is in the **Running** state.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Storage > Disks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the data disk and choose **More > Detach** from the **Actions** column.
5. Click **OK**.

2.5.12. Release a data disk

You can release a data disk that is no longer needed. The released data disk cannot be recovered. Exercise caution when you release a data disk.

Prerequisites

The data disk is in the **Pending** state. If the data disk is attached to an instance, detach the disk from the instance first.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Storage > Disks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the data disk and choose **More > Release** from the **Actions** column.
5. Click **OK**.

2.6. Images

2.6.1. Create a custom image

You can create a custom image and use it to create identical instances or replace the system disks of existing instances. This allows you to have multiple instances with the same operating system and data environment.

Create a custom image from a snapshot

You can create a custom image from a system disk snapshot to fully load the operating system and data environment of the snapshot to the image. Before you perform this operation, make sure that a snapshot of system disks is used. You cannot create a custom image from snapshots of data disks.

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Snapshots and Images > Snapshots**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the snapshot and click **Create Custom Image** in the **Actions** column.
5. Enter the name and description of the image, and then click **OK**.

The name must be 2 to 128 characters in length and can contain periods (.), underscores (_), hyphens (-), and colons (:). It cannot start with a special character or digit.

The description must be 2 to 256 characters in length and cannot start with `http://` and `https://`.

Create a custom image from an instance

You can create a custom image from an instance to completely replicate the data of all disks of the instance, including the system disk and data disks.

 **Note** To avoid data security risks, delete sensitive data before you create a custom image.

When you create a custom image from an instance, a snapshot is generated for each disk in the instance, and all the snapshots constitute a complete custom image.

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Instances**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the instance and choose **More > Create Custom Image** from the **Actions** column.
5. Enter the name and description of the custom image, and then click **OK**.

The name must be 2 to 128 characters in length and can contain periods (.), underscores (_), hyphens (-), and colons (:). It cannot start with a special character or digit.

The description must be 2 to 256 characters in length and cannot start with `http://` and `https://`.

2.6.2. View images

You can view the list of created images.

Procedure

1. [Log on to the ECS console.](#)
2. In the left-side navigation pane, choose **Snapshots and Images > Images**.
3. In the top navigation bar, select an organization, a resource set, and a region. The created images that match the specified criteria are displayed.
4. Select the tab based on the type of images you want to view.
You can select the **Custom Images** or **Public Images** tab.
5. Select a filtering option from the drop-down list, enter the relevant information in the search bar, and click **Search**.

You can select multiple filtering options to narrow down the search results.

| Filtering option | Description |
|------------------|--|
| Image Name | Enter an image name to search for the image. |
| Image ID | Enter an image ID to search for the image. |
| Snapshot ID | Enter a snapshot ID to search for the images associated with the snapshot. This option is not available for public images. |

2.6.3. View instances related to an image

You can view the instances that use the specified image.

Procedure

1. [Log on to the ECS console.](#)
2. In the left-side navigation pane, choose **Snapshots and Images > Images**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Select a tab based on the type of the image.
You can select the **Custom Images** or **Public Images** tab.
5. Find the target image and click **Related Instances** in the **Actions** column.

Result

The Instances page appears, showing the instances that use the image. You can perform operations on these instances, such as updating the image.

2.6.4. Modify the description of a custom image

You can modify the description of a created custom image.

Procedure

1. [Log on to the ECS console.](#)
2. In the left-side navigation pane, choose **Snapshots and Images > Images**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the target custom image and click **Modify Description** in the **Actions** column.
5. In the dialog box that appears, modify the image description in the Basic Settings field.
The description must be 2 to 256 characters in length and cannot start with `http://` or `https://`.
6. Click **OK**.

2.6.5. Share custom images

You can share a custom image that you create to organizations that you manage to create multiple identical ECS instances in a short time.

Context

Only custom images can be shared. Shared images are not counted towards the image quota assigned to the organization.

The organization can use the shared image to create instances or replace system disks of existing instances.

You can delete shared images. After a shared image is deleted, the image is no longer visible to the organization to which the image was shared. The system disk of instances created from the shared image can no longer be reinitialized.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Snapshots and Images > Images**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the image and click **Share Image** in the **Actions** column.
5. Select the organization to which you want to share the image and click **OK**.
An image can be shared only to the organizations that the image owner manages.

2.6.6. Encrypt a custom image

This topic describes how to encrypt a custom image to generate a new identical encrypted custom image.

Prerequisites

The custom image that you want to encrypt is in the Available (Available) state.

Context

To meet the requirements for data security compliance, you can use encrypted custom images to create instances. The system disks of the created instances are automatically encrypted.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Snapshots and Images > Images**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Click the **Custom Images** tab.
5. Find the custom image that you want to encrypt and click **Encrypt Image** in the **Actions** column.
6. In the **Encrypt Image** dialog box, configure the parameters listed in the following table.

| Parameter | Description |
|-----------------|--|
| Image ID | The system automatically obtains the ID of the image to be encrypted. You do not need to configure this parameter. |

| Parameter | Description |
|-------------------|---|
| Custom Image Name | Enter a name for the new encrypted custom image. The name must be 2 to 128 characters in length and can contain letters, digits, underscores (_), periods (.), and hyphens (-). It must start with a letter. |
| Description | Enter a description for the new encrypted custom image. The description must be 2 to 256 characters in length and cannot start with http:// or https://. |

7. Click OK.

Result

After you encrypt the custom image, a new identical encrypted custom image is generated and displayed on the Custom Images tab.

2.6.7. Import custom images

2.6.7.1. Limits on importing custom images

This topic describes the limits on importing images. You must understand the limits to ensure image availability and improve import efficiency.

The following limits apply when you import custom images:

- [Limits on importing custom images in Linux](#)
- [Limits on importing custom images in Windows](#)

Limits on importing custom images in Linux

When you import custom images in Linux, note the following limits:

- Multiple network interfaces are not supported.
- IPv6 addresses are not supported.
- The password must be 8 to 30 characters in length. It must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters.
- The firewall is disabled, and port 22 is opened.
- The Linux system disk size ranges from 40 GiB to 500 GiB.
- DHCP must be enabled in the image.
- SELinux is disabled.
- The Kernel-based Virtual Machine (KVM) driver must be installed.
- We recommend that you install cloud-init to configure the hostname and NTP and yum sources.

Limits

| Item | Standard operating system image | Non-standard operating system image |
|------|---------------------------------|-------------------------------------|
|------|---------------------------------|-------------------------------------|

| Item | Standard operating system image | Non-standard operating system image |
|-------------------------------------|--|---|
| Description | <p>The supported standard 32-bit and 64-bit operating systems include:</p> <ul style="list-style-type: none"> CentOS Ubuntu SUSE openSUSE Red Hat Debian CoreOS Aliyun Linux <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note Support for standard operating systems may be subject to version changes. You can access the ECS console to check the latest supported operating systems.</p> </div> | <p>Non-standard operating systems include:</p> <ul style="list-style-type: none"> Operating systems that are not supported by Alibaba Cloud. Standard operating systems that do not meet the requirements of critical system configuration files, basic system environments, and applications. <p>If you want to use non-standard operating system images, you must select Others Linux when importing images. If you import non-standard operating system images, Alibaba Cloud does not perform any processing on the instances created from these images. After you create an instance, you must connect to the instance by clicking Connect to VPN in the ECS console. You can then configure the IP address, route, and password for the instance.</p> |
| Critical system configuration files | <ul style="list-style-type: none"> Do not modify <code>/etc/issue*</code>. Otherwise, the version of the operating system cannot be identified, which leads to system creation failure. Do not modify <code>/boot/grub/menu.lst</code>. Otherwise, the system may fail to start. Do not modify <code>/etc/fstab</code>. Otherwise, partitions cannot be loaded, which leads to system startup failure. Do not modify <code>/etc/shadow</code> to read-only. Otherwise, the password file cannot be modified, which leads to system creation failure. Do not modify <code>/etc/selinux/config</code> to enable SELinux. Otherwise, the system may fail to start. | <p>Requirements for standard operating systems</p> |

| Item | Standard operating system image | Requirements for standard operating system image |
|---|---|--|
| Requirements for the basic system environment | <ul style="list-style-type: none"> Do not adjust the system disk partitions. Only disks with a single root partition are supported. Make sure that the system disk has sufficient storage space. Do not modify critical system files, such as <code>/sbin</code>, <code>/bin</code>, and <code>/lib*</code>. Before importing an image, confirm the integrity of the file system. Only ext3 and ext4 file systems are supported. | |
| Applications | Do not install <code>qemu-ga</code> on a custom image. Otherwise, some of the services that Alibaba Cloud needs may be unavailable. | |
| Image file formats | Only images in the RAW, VHD, or qcow2 format can be imported. If you want to import images in other formats, use a tool to convert the format before importing the image. We recommend that you import images in the VHD format, which has a smaller transmission footprint. | |
| Image file size | We recommend that you configure the disk size for importing images based on the virtual disk size (not the image file size). The disk size for importing images must be at least 40 GiB. | |

Limits on importing custom images in Windows

When you import custom images in Windows, note the following limits:

- The password must be 8 to 30 characters in length. It must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters.
- Imported Windows images do not provide the Windows activation service.
- The firewall must be disabled. Otherwise, remote logon is unavailable. Port 3389 must be opened.
- The Windows system disk size ranges from 40 GiB to 500 GiB.

Limits

| Item | Description |
|------|-------------|
|------|-------------|

| Item | Description |
|---|--|
| Operating system versions | <p>Alibaba Cloud supports importing the following versions of 32-bit and 64-bit operating system images:</p> <ul style="list-style-type: none"> • Microsoft Windows Server 2016 • Microsoft Windows Server 2012, including: <ul style="list-style-type: none"> ◦ Microsoft Windows Server 2012 R2 (Standard Edition) ◦ Microsoft Windows Server 2012 (Standard Edition and Datacenter Edition) • Microsoft Windows Server 2008, including: <ul style="list-style-type: none"> ◦ Microsoft Windows Server 2008 R2 (Standard Edition, Datacenter Edition, and Enterprise Edition) ◦ Microsoft Windows Server 2008 (Standard Edition, Datacenter Edition, and Enterprise Edition) • Microsoft Windows Server 2003, including: <ul style="list-style-type: none"> ◦ Microsoft Windows Server 2003 R2 (Standard Edition, Datacenter Edition, and Enterprise Edition) ◦ Microsoft Windows Server 2003 (Standard Edition, Datacenter Edition, and Enterprise Edition) or later, including Service Pack 1 (SP1) • Microsoft Windows 7, including: <ul style="list-style-type: none"> ◦ Microsoft Windows 7 (Professional Edition) ◦ Microsoft Windows 7 (Enterprise Edition) <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note Support for standard operating systems may be subject to version changes. You can access the ECS console to check the latest supported operating systems.</p> </div> |
| Requirements for the basic system environment | <ul style="list-style-type: none"> • Multi-partition system disks are supported. • Make sure that the system disk has sufficient storage space. • Do not modify critical system files. • Before importing an image, confirm the integrity of the file system. • The NTFS file system with the MBR partition type is supported. |
| Applications | <p>Do not install qemu-ga on an imported image. If it is installed, some of the services that Alibaba Cloud needs may be unavailable.</p> |
| Image file formats | <ul style="list-style-type: none"> • RAW • VHD • qcow2 <p>We recommend that you configure the system disk size for importing images based on the virtual disk size (not the image file size). The system disk size for importing images must range from 40 GiB to 500 GiB.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note We recommend that you import images in the VHD format, which has a smaller transmission footprint.</p> </div> |

2.6.7.2. Convert the image file format

You can only import image files in the RAW, VHD, and qcow2 formats to ECS. If you want to import images in other formats, you must convert the image into a supported format. This topic describes how to convert the image format in Windows and Linux.

Context

You can use the `qemu-img` tool to convert an image from VMDK, VDI, VHDX, qcow1, or QED to RAW, VHD, or qcow2, or implement conversion between RAW, VHD, and qcow2.

 **Note** We recommend that you use the qcow2 format if your application environment supports this format.

Windows

1. Download qemu.

Visit [QEMU Binaries for Windows \(64 bit\)](#) to download the qemu tool. Select a qemu version based on your operating system.

2. Install qemu.

The installation path in this example is `C:\Program Files\qemu`.

3. Configure the environment variables for qemu.

- i. Choose **Start > Computer**, right-click Computer, and choose **Properties** from the shortcut menu.

- ii. In the left-side navigation pane, click **Advanced System Settings**.

- iii. In the **System Properties** dialog box that appears, click the **Advanced** tab and then click **Environment Variables**.

- iv. In the **Environment Variables** dialog box that appears, find the **Path** variable from the **System variables** section.

- If the **Path** variable exists, click **Edit**.

- If the **Path** variable does not exist, click **New**.

- v. Add a system variable value.

- In the **Edit System Variable** dialog box that appears, add `C:\Program Files\qemu` to the **Variable value** field, separate different variable values with semicolons (;), and then click **OK**.

- In the **New System Variable** dialog box that appears, enter `Path` in the **Variable name** field, enter `C:\Program Files\qemu` in the **Variable value** field, and then click **OK**.

4. Open Command Prompt in Windows and run the `qemu-img --help` command. If a successful response is displayed, the tool is installed.

5. In the Command Prompt window, run the `cd [Directory of the source image file]` command to switch to a new file directory,

for example, `cd D:\ConvertImage`.

6. In the Command Prompt window, run the `qemu-img convert -f raw -O qcow2 centos.raw centos.qcow2` command to convert the image file format.

The parameters are described as follows:

- The `-f` parameter is followed by the source image format.

- The `-O` parameter (case-sensitive) is followed by the destination image format, source file name, and destination file name.

After the conversion is complete, the destination file appears in the directory of the source image file.

Linux

1. Install the qemu-img tool.
 - o For Ubuntu, run the `apt install qemu-img` command.
 - o For CentOS, run the `yum install qemu-img` command.
2. Run the `qemu-img convert -f raw -O qcow2 centos.raw centos.qcow2` command to convert the image file format. The parameters are described as follows:
 - o The `-f` parameter is followed by the source image format.
 - o The `-O` parameter (case-sensitive) is followed by the destination image format, source file name, and destination file name.

2.6.7.3. Import a custom image

After you upload a local image to an OSS bucket, you can import the image to the ECS environment as a custom image.

Prerequisites

- An image that meets the limits and requirements for image import has been made. The image must be in the RAW, VHD, or qcow2 format. For more information, see [Limits on importing custom images](#) and [Convert the image file format](#).
- You have been authorized to import images. For more information, see *the RAM authorization section* in ASCM Console User Guide.
- A local image has been uploaded to a bucket by using the OSS console or calling an OSS API operation. For more information, see *the Upload objects section in OSS User Guide* or *the Put Object section* in OSS Developer Guide.

 **Note** Make sure that the bucket is in the same region as the custom image you want to create.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Snapshots and Images > Images**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Click **Import Image**.
5. Configure the parameters of the image.

| Parameter | Required | Description |
|-----------------|----------|--|
| Region | Yes | The region of the custom image to be imported. |
| Organization | Yes | The organization to which the custom image belongs. |
| Resource Set | Yes | The resource set of the custom image. |
| OSS Bucket Name | Yes | The name of the OSS bucket where the image to be imported resides. |
| OSS Object Name | Yes | The endpoint of the OSS object where the image is stored. For information about how to obtain the endpoint of an OSS object, see <i>the Obtain object URL s section</i> in OSS User Guide. |

| Parameter | Required | Description |
|---------------------|----------|--|
| Image Name | Yes | The name of the custom image. The name must be 2 to 128 characters in length. It must start with a letter and can contain letters, periods (.), underscores (_), and hyphens (-). |
| Operating System | Yes | Linux and Windows are available. |
| System Disk | Yes | The size of the system disk of the ECS instance. Unit: GiB. |
| System Architecture | Yes | x86_64 and i386 are available. |
| Platform | Yes | Linux: <ul style="list-style-type: none"> ◦ CentOS ◦ Ubuntu ◦ SUSE ◦ OpenSUSE ◦ Debian ◦ CoreOS ◦ Aliyun ◦ Others Linux ◦ Customized Linux Windows: <ul style="list-style-type: none"> ◦ Windows Server 2003 ◦ Windows Server 2008 ◦ Windows Server 2012 |
| Image Format | Yes | The format of the custom image. RAW , VHD , and qcow2 are available. |
| Description | No | The description of the custom image. |

6. Click **OK**.

Result

You can go to the Images page to view the progress of custom image creation. For more information, see [View images](#). When 100% is displayed in the Progress column of the Images page, the image is created.

2.6.8. Export a custom image

You can export a custom image to an OSS bucket and then download it to your local device.

Prerequisites

- OSS is activated and an OSS bucket is created. For more information, see the *"Create buckets"* section in OSS User Guide.
- You are authorized to export images. For more information, see *the "RAM" chapter* in ASCM Console User Guide.

Context

You can export custom images to the RAW, VHD, or qcow2 format. After a custom image is exported to an OSS bucket, you can download the image to your local device. For more information, see *the "Obtain object URLs"*

section in OSS User Guide.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Snapshots and Images > Images**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the target custom image and click **Export Image** in the **Actions** column.
5. Select a bucket for **OssBucket**. Enter a prefix in the **OSS Prefix** field. Then click **OK**.
The **OSS Prefix** field is optional. The prefix must be 1 to 30 characters in length and can contain digits and letters.

2.6.9. Delete a custom image

You can delete a custom image that is no longer needed. However, public images cannot be deleted.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Snapshots and Images > Images**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Select one of the following methods to delete custom images.
 - To delete one image, find the image and click **Delete Image** in the **Actions** column.
 - To delete multiple images, select images and click **Delete** at the bottom of the image list.
5. Click **OK**.

2.7. Snapshots

2.7.1. Create a snapshot

You can manually create a snapshot for a disk to back up disk data.

Prerequisites

- The instance to which the disk is attached is in the **Running** or **Stopped** state.
- The disk is in the **Running** state.

Context

You can create up to 64 snapshots for each disk.

Snapshots can be used in the following scenarios:

- Restore a disk from one of its snapshots.
- Create a custom image from a system disk snapshot.

For more information, see [Create a custom image from a snapshot](#). Data disk snapshots cannot be used to create custom images.

- Create a new data disk from a data disk snapshot.

To create a data disk from a snapshot, set **Use Snapshot** to **Yes** and specify a snapshot on the **Create Disk** page. For more information, see [Create a disk](#). When you re-initialize a data disk created from a snapshot, the disk is restored to the status of the snapshot.

Note the following considerations when you create a snapshot:

- For each disk, the first snapshot is a full snapshot and subsequent snapshots are incremental snapshots. It takes

an extended period of time to create the first snapshot. It takes a short period of time to create an incremental snapshot. The amount of taken time depends on the volume of data that has been changed since the last snapshot. The more data that has been changed, the more time it takes.

- Avoid creating snapshots during peak hours.

Procedure

1. [Log on to the ECS console.](#)
2. In the left-side navigation pane, choose **Storage > Disks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the target disk and click **Create Snapshot** in the **Actions** column.
5. Enter a snapshot name and description and then click **OK**.

 **Note** The names of manual snapshots cannot start with `auto` because `auto` is a prefix reserved for automatic snapshots.

You can go to the **Snapshots** page to check the progress of snapshot creation. For more information, see [View snapshots](#). When 100% is displayed in the **Progress** column, the snapshot is created.

2.7.2. View snapshots

You can view the list of created snapshots.

Procedure

1. [Log on to the ECS console.](#)
2. In the left-side navigation pane, choose **Snapshots and Images > Snapshots**.
3. In the top navigation bar, select an organization, a resource set, and a region.
The created snapshots that match the specified criteria are displayed.
4. Select a filter option from the drop-down list, enter the relevant information in the search bar, and click **Search**.

You can select multiple filter options to narrow down search results.

| Filter option | Description |
|---------------|---|
| Snapshot Name | Enter a snapshot name to search for the snapshot. |
| Snapshot ID | Enter a snapshot ID to search for the snapshot. |
| Instance ID | Enter an instance ID to search for the snapshots related to the instance. |
| Disk ID | Enter a disk ID to search for the snapshots related to the disk. |
| Snapshot Type | Select a snapshot type to search for the snapshots of that type. Valid values: <ul style="list-style-type: none"> ◦ All ◦ User Snapshots: manual snapshots ◦ Automatic Snapshots: automatic snapshots |
| Creation Time | Enter a time to search for the snapshots that were created at that time. |

2.7.3. Delete a snapshot

You can delete a snapshot that is no longer needed. After the snapshot is deleted, it cannot be recovered. You cannot delete system disk snapshots that have been used to create custom images.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Snapshots and Images > Snapshots**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Use one of the following methods to delete the snapshot:
 - To delete a single snapshot, find the snapshot and click **Delete** in the **Actions** column.
 - To delete one or more snapshots at a time, select the snapshots and click **Delete** in the lower-left corner of the Snapshots page.
5. In the message that appears, click **OK**.

2.8. Automatic snapshot policies

2.8.1. Create an automatic snapshot policy

Automatic snapshot policies can apply to both system disks and data disks and can be used to create periodical snapshots for the disks. Using automatic snapshot policies can improve data availability and operation error tolerance.

Context

Automatic snapshot policies can effectively eliminate the following risks associated with manual snapshot creation:

- When applications such as personal websites or databases deployed on an ECS instance encounter attacks or system vulnerabilities, you may not be able to manually create snapshots. In this case, you can use the latest automatic snapshot to roll back the affected disks to restore your data and reduce losses.
- You can also specify an automatic snapshot policy to create snapshots before regular system maintenance tasks. This eliminates the need to manually create snapshots and ensures that snapshots are always created before maintenance.

You can create up to 64 snapshots for each disk. If the maximum number of snapshots for a disk is reached when a new snapshot is being created, the system deletes the oldest automatic snapshot.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Snapshots and Images > Automatic Snapshot Policies**.
3. Click **Create Policy**.
4. Configure the properties of the automatic snapshot policy.

| Parameter | Required | Description |
|--------------|----------|--|
| Region | Yes | The ID of the region to which the automatic snapshot policy applies. |
| Organization | Yes | The organization to which the automatic snapshot policy applies. |

| Parameter | Required | Description |
|------------------|----------|--|
| Policy Name | Yes | The name of the automatic snapshot policy. The name must be 2 to 128 characters in length and cannot start with a special character or digit. It can contain periods (.), underscores (_), hyphens (-), and colons (:). |
| Creation Time | Yes | <p>The time when a snapshot is automatically created. You can select any hour from 00:00 to 23:00.</p> <div style="border: 1px solid #add8e6; padding: 5px; margin: 5px 0;"> <p> Note The default time zone for the snapshot policy is UTC+8. You can change the time zone based on your business requirements.</p> </div> <p>The creation of an automatic snapshot is canceled if the scheduled time for creating the snapshot is reached but the previous automatic snapshot is still being created. This may occur if the disk stores a large volume of data. For example, you can specify a policy for the system to create automatic snapshots at the following points in time: 00:00, 01:00, and 02:00. When the system starts creating a snapshot at 00:00, it takes 70 minutes for the system to complete the snapshot task. Therefore, the system does not create another snapshot at 01:00. Instead, after the system completes the snapshot task at 01:10, the system creates the next snapshot at 02:00.</p> |
| Frequency | Yes | The day when a snapshot is created. You can select multiple values. The day ranges from Monday to Sunday. |
| Retention Policy | No | <p>The retention policy of the automatic snapshot. The default value of the retention time is 30 days. You can configure the following parameters:</p> <ul style="list-style-type: none"> ◦ Keep for: Specify the number of days during which the snapshots can be retained. Valid values: 1 to 65536. ◦ Always keep the snapshots until the number of snapshots reaches the upper limit: Select this option to retain the snapshots until the maximum number of snapshots is reached. |

5. Click OK.

What's next

After the automatic snapshot policy is created, you need to apply it to a disk to automatically create snapshots for the disk. For more information, see [Configure an automatic snapshot policy for multiple disks](#).

2.8.2. View automatic snapshot policies

You can view the list of automatic snapshot policies.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Snapshots and Images > Automatic Snapshot Policies**.
3. In the top navigation bar, select an organization, a resource set, and a region.
The created automatic snapshot policies that match the specified criteria are displayed.
4. View the list of automatic snapshot policies.

2.8.3. Modify an automatic snapshot policy

You can modify the properties of an automatic snapshot policy, such as the name, creation time, frequency, and retention policy.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Snapshots and Images > Automatic Snapshot Policies**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the automatic snapshot policy and click **Modify Policy** in the **Actions** column.
5. Modify the properties of the policy.

Changes made to the retention duration do not take effect on existing snapshots, but take effect only on newly created snapshots.

6. Click **OK**.

2.8.4. Configure an automatic snapshot policy

After you apply an automatic snapshot policy to a disk, snapshots will be created automatically for the disk based on the policy settings. You can cancel an applied automatic snapshot policy at any time.

Context

We recommend that you configure the automatic snapshot policy to create automatic snapshots during off-peak hours. You can also manually create a snapshot for the disk that already has an automatic snapshot policy applied. When an automatic snapshot is being created, you must wait until the snapshot is complete before you can create a manual snapshot. The automatic snapshot is named in the following format: `auto_yyyyMMdd_1`, such as `auto_20140418_1`.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Storage > Disks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the disk and click **Configure Automatic Snapshot Policy** in the **Actions** column.
5. Select a procedure based on the operation you want to perform on the policy.
 - To apply an automatic snapshot policy, turn on **Automatic Snapshot Policy**, select a policy, and then click **OK**.
 - To cancel an automatic snapshot policy, turn off **Automatic Snapshot Policy** and click **OK**.

2.8.5. Configure an automatic snapshot policy for multiple disks

After you apply an automatic snapshot policy to a disk, snapshots will be created automatically for the disk based on the policy settings. You can cancel an applied automatic snapshot policy at any time.

Context

We recommend that you configure the automatic snapshot policy to create automatic snapshots during off-peak hours. You can also manually create a snapshot for the disk that already has an automatic snapshot policy applied. When an automatic snapshot is being created, you must wait until the snapshot is complete before you can create a manual snapshot. The automatic snapshot is named in the following format: `auto_yyyyMMdd_1`, such as `auto_20140418_1`.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Snapshots and Images > Automatic Snapshot Policies**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the automatic snapshot policy and click **Apply Policy** in the **Actions** column.
5. Select a tab based on the operation you want to perform on the disks.
 - To apply the automatic snapshot policy, click the **Disks Without Policy Applied** tab, select one or more disks, and click **Apply Policy** at the bottom of the disk list.
 - To cancel the automatic snapshot policy, click the **Disks With Policy Applied** tab, select one or more disks, and click **Disable Automatic Snapshot Policy** at the bottom of the disk list.

2.8.6. Delete an automatic snapshot policy

You can delete an automatic snapshot policy that is no longer needed. After you delete the automatic snapshot policy, the policy is automatically canceled for disks that have it applied.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Snapshots and Images > Automatic Snapshot Policies**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the automatic snapshot policy and click **Delete Policy** in the **Actions** column.
5. In the message that appears, click **OK**.

2.9. Security groups

2.9.1. Create a security group

Security groups are an important means for network security isolation. They are used to set network access control for one or more ECS instances.

Prerequisites

A Virtual Private Cloud (VPC) has been created. For more information, see *VPC User Guide*.

Context

Instances that belong to the same account and are in the same region and in the same security group can

communicate with each other over the internal network. If instances that belong to the same account in the same region are in different security groups, you can implement internal network communication by authorizing mutual access between two security groups.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > Security Groups**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Click **New Security Group**.
5. Configure the parameters of the security group.

| Type | Parameter | Required | Description |
|----------------|---------------------|----------|--|
| Region | Organization | Yes | The organization to which the security group belongs. Make sure that the security group and the VPC belong to the same organization. |
| | Resource Set | Yes | The resource set to which the security group belongs. Make sure that the security group and the VPC belong to the same resource set. |
| | Region | Yes | The region to which the security group belongs. Make sure that the security group and the VPC belong to the same region. |
| | Zone | Yes | The ID of the zone where the security group resides. |
| Basic Settings | VPC | Yes | The VPC to which the security group belongs. |
| | Security Group Name | No | The name must be 2 to 128 characters in length and start with a letter. It can contain letters, digits, periods (.), underscores (_), hyphens (-), colons (:), and commas (,). It cannot start with http:// or https://. |
| | | | |

| Type | Parameter | Required | Description |
|------|-------------|----------|--|
| | Description | No | The description of the security group. We recommend that you provide an informational description to simplify future management operations. The name must be 2 to 256 characters in length and start with a letter. It can contain letters, digits, periods (.), underscores (_), hyphens (-), and commas (,). It cannot start with http:// or https://. |

6. Click **Submit**.

2.9.2. View security groups

You can view the list of security groups that you create.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > Security Groups**.
3. In the top navigation bar, select an organization, a resource set, and a region. The created security groups that match the specified criteria are displayed.
4. Select a filtering option from the drop-down list, enter the relevant information in the search bar, and click **Search**.

You can select multiple filtering options to narrow down the search results.

| Filtering option | Description |
|---------------------|--|
| Security Group ID | Enter a security group ID to search for the security group. |
| Security Group Name | Enter a security group name to search for the security group. |
| VPC ID | Enter a VPC ID to search for the security groups that belong to the VPC. |

2.9.3. Modify a security group

You can modify the name and description of a created security group.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > Security Groups**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the security group and click **Modify** in the **Actions** column.

5. Modify the name and description of the security group.
6. Click **OK**.

2.9.4. Add a security group rule

You can use security group rules to control access to and from the ECS instances in a security group over the Internet and the internal network.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > Security Groups**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the target security group and click **Rules** in the **Actions** column.
5. Click **Create Rule**.
6. Configure parameters listed in the following table to create a security group rule.

| Parameter | Required | Description |
|-----------|----------|--|
| ENI Type | Yes | Valid value: Internal Network ENI . In VPCs, you cannot find public NICs in ECS instances and can add only internal security group rules. However, the added security group rules apply to both the Internet and internal network. |
| Direction | Yes | <ul style="list-style-type: none"> ◦ Outbound: access from the ECS instances in the current security group to other ECS instances on the internal network or resources on the Internet. ◦ Inbound: access from other ECS instances on the internal network or resources on the Internet to the ECS instances in the current security group. |
| Action | Yes | <ul style="list-style-type: none"> ◦ Allow: allows access requests on the specified port or ports. ◦ Deny: discards data packets and returns no messages. <p>If two security group rules are different only in the Action parameter, the Deny rule takes effect while the Allow rule is ignored.</p> |
| Protocol | Yes | <ul style="list-style-type: none"> ◦ All: all protocols. This value can be used in total trust scenarios. ◦ TCP: can be used to allow or deny traffic on one or several successive ports. ◦ UDP: can be used to allow or deny traffic on one or several successive ports. ◦ ICMP: can be used when the <code>ping</code> command is used to test the status of the network connection between instances. ◦ ICMPv6: can be used when the <code>ping6</code> command is used to test the status of the network connection between instances. ◦ GRE: can be used for VPN. |

| Parameter | Required | Description |
|----------------------|----------|--|
| Port Range | Yes | <p>The port range depends on the protocol type.</p> <ul style="list-style-type: none"> When you set Protocol to All, -1/-1 is displayed, indicating all ports. You cannot specify a port range in this case. When you set Protocol to TCP, you can use the <start port number>/<end port number> format to specify a port range. Valid port numbers: 1 to 65535. Set the start port number and end port number to the same value to specify a single port. For example, use 22/22 to specify port 22. When you set Protocol to UDP, you can use the <start port number>/<end port number> format to specify a port range. Valid port numbers: 1 to 65535. Set the start port number and end port number to the same value to specify a single port. For example, use 3389/3389 to specify port 3389. When you set Protocol to ICMP, -1/-1 is displayed, indicating all ports. You cannot specify a port range in this case. When you set Protocol to ICMPv6, -1/-1 is displayed, indicating all ports. You cannot specify a port range in this case. When you set Protocol to GRE, -1/-1 is displayed, indicating all ports. You cannot specify a port range in this case. |
| Priority | Yes | <p>The priority of the rule. Valid values: 1 to 100. The default value is 1, indicating the highest priority.</p> |
| Authorization Type | Yes | <ul style="list-style-type: none"> IPv4 Addresses: IPv4 addresses or IPv4 CIDR blocks. IPv6 Addresses: IPv6 addresses or IPv6 CIDR blocks. Security Groups: another security group. This authorization type takes effect only on the internal network. |
| Authorization Object | Yes | <p>Authorization objects depend on the authorization type.</p> <p>When you set Authorization Type to IPv4 Addresses:</p> <ul style="list-style-type: none"> Enter an IPv4 address or IPv4 CIDR block. Example: <i>12.1.1.1</i> or <i>13.1.1.1/25</i>. You can enter up to 10 authorization objects at a time. Separate multiple objects with commas (,). If you specify <i>0.0.0.0/0</i>, all IPv4 addresses will be allowed or denied based on the Action parameter. Exercise caution when you specify <i>0.0.0.0/0</i>. <p>When you set Authorization Type to IPv6 Addresses:</p> <ul style="list-style-type: none"> Enter an IPv6 address or IPv6 CIDR block. Example: <i>2001:0db8::1428:****</i> or <i>2001:0db8::1428:****/128</i>. You can enter up to 10 authorization objects at a time. Separate multiple objects with commas (,). If you specify <i>::/0</i>, all IPv6 addresses will be allowed or denied based on the Action parameter. Exercise caution when you specify <i>::/0</i>. <p>When you set Authorization Type to Security Groups, select a security group ID. If the current security group is of the VPC type, the selected security group must be in the same VPC as the current security group.</p> |

| Parameter | Required | Description |
|-------------|----------|---|
| Description | No | The description of the security group rule. We recommend that you provide an informational description to simplify future management operations. The description must be 2 to 256 characters in length and cannot start with http:// or https://. |

7. Click **OK**.

2.9.5. Clone a security group rule

You can clone a security group rule to quickly create a similar rule.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > Security Groups**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the target security group and click **Rules** in the **Actions** column.
5. On the Rules page that appears, click the **Inbound** or **Outbound** tab.
6. Find the target security group rule and click **Clone** in the **Actions** column.
7. In the **Clone Security Group Rule** dialog box, modify the attributes of the security group rule.
For more information about the attributes of security group rules, see [Add a security group rule](#).
8. Click **OK**.

2.9.6. Modify a security group rule

You can modify improper rules in a security group to ensure the security of ECS instances in the security group.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > Security Groups**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the target security group and click **Rules** in the **Actions** column.
5. On the Rules page that appears, click the **Inbound** or **Outbound** tab.
6. Find the target security group rule and click **Modify** in the **Actions** column.
7. In the **Modify Security Group Rule** dialog box, modify the attributes of the security group rule.
For more information about the attributes of security group rules, see [Add a security group rule](#).
8. Click **OK**.

2.9.7. Export security group rules

You can export security group rules of a security group to a local device for backup.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > Security Groups**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the target security group and click **Rules** in the **Actions** column.

5. On the Rules page that appears, click the **Inbound** or **Outbound** tab.
6. Click **Export** in the upper-right corner to download and save the rules to a local device.

2.9.8. Import security group rules

You can import a local backup file of security group rules into a security group to quickly create or restore security group rules.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > Security Groups**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the target security group and click **Rules** in the **Actions** column.
5. On the Rules page that appears, click the **Inbound** or **Outbound** tab.
6. Click **Import** in the upper-right corner.
7. In the **Import Rule** dialog box, click **Choose File**.
8. Select the target local backup file of security group rules and click **Open**. Then, click **OK**.

The local backup file must be in the CSV format. You can download a template file from the **Import Rule** dialog box.

2.9.9. Add an instance

You can add an existing instance to a security group in the same region. After the instance is added, the security group rules of the security group automatically apply to the instance.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > Security Groups**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the security group and click **Manage Instances** in the **Actions** column.
5. Click **Add Instance**.
6. Select an instance and click **OK**.

2.9.10. Remove instances from a security group

You can remove instances from a security group, but each of the instances must always belong to at least one security group.

Prerequisites

The instances to be removed are added to two or more security groups.

Context

After an ECS instance is removed from a security group, the instance will be isolated from the other ECS instances in the security group. We recommend that you perform a full test in advance to ensure that services can run properly after you remove the instance.

Procedure

1. [Log on to the ECS console](#).

2. In the left-side navigation pane, choose **Networks and Security > Security Groups**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the target security group and click **Manage Instances** in the **Actions** column.
5. On the Instances page that appears, select one or more instances and click **Remove** in the lower-left corner.
6. Click **OK**.

2.9.11. Delete a security group

You can delete a security group that is no longer needed.

Prerequisites

No instances exist in the security group.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > Security Groups**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Use one of the following methods to delete the security group:
 - To delete a single security group, find the security group and click **Delete** in the **Actions** column.
 - To delete one or more security groups at a time, select the security groups and click **Delete** in the lower-left corner of the Security Groups page.
5. In the message that appears, click **OK**.

2.10. Elastic Network Interfaces

2.10.1. Create an ENI

You can bind elastic network interfaces (ENIs) to instances to create high-availability clusters and implement fine-grained network management. You can also unbind an ENI from an instance and then bind the ENI to another instance to implement a low-cost failover solution.

Prerequisites

- A virtual private cloud (VPC) and a VSwitch are created. For more information, see [Create a VPC](#) and [Create a VSwitch](#) in *Apsara Stack VPC User Guide*.
- A security group is available in the VPC. If no security group is available in the VPC, create a security group. For more information, see [Create a security group](#).

Context

ENIs are classified into primary and secondary ENIs.

A primary ENI is created by default when an instance is created in a VPC. This primary ENI has the same lifecycle as the instance and cannot be unbound from the instance.

ENIs created separately are secondary ENIs. You can bind secondary ENIs to or unbind them from instances. This topic describes how to create a secondary ENI.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > ENIs**.
3. In the top navigation bar, select an organization, a resource set, and a region.

4. Click **Create ENI**.
5. Configure parameters listed in the following table to create an ENI.

| Section | Parameter | Required | Description |
|----------------|----------------|----------|---|
| Region | Organization | Yes | The organization in which to create the ENI. |
| | Resource Set | Yes | The resource set in which to create the ENI. |
| | Region | Yes | The region in which to create the ENI. |
| | Zone | Yes | The zone in which to create the ENI. |
| Basic Settings | VPC | Yes | <p>The VPC in which to create the ENI. The secondary ENI can be bound only to an instance in the same VPC.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note After the ENI is created, you cannot change its VPC.</p> </div> |
| | VSwitch | Yes | <p>The VSwitch to be associated with the ENI. The secondary ENI can be bound only to an instance in the same VPC. Select a VSwitch that is in the same zone as the instance to which the ENI will be bound. The VSwitch of the ENI can be different from that of the instance.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note After an ENI is created, you cannot change its VSwitch.</p> </div> |
| | Security Group | Yes | The security group in which to create the ENI within the specified VPC. The rules of the security group automatically apply to the ENI. |

| Section | Parameter | Required | Description |
|---------|--------------------|----------|--|
| | ENI Name | Yes | The name of the ENI. The name must be 2 to 128 characters in length. It must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . It can contain letters, digits, periods (.), underscores (_), hyphens (-), colons (:), and commas (,). |
| | Description | No | The description of the ENI. We recommend that you provide an informational description to simplify future management operations. The description must be 2 to 256 characters in length. It must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . It can contain letters, digits, periods (.), underscores (_), hyphens (-), colons (:), and commas (,). |
| | Primary Private IP | No | The primary private IPv4 address of the ENI. The IPv4 address must be within the CIDR block of the specified VSwitch. If you do not specify a primary private IP address, the system automatically assigns a private IP address to the ENI. |

6. Click **Submit**.

Result

The created ENI is displayed on the ENIs page and is in the **Available** state.

2.10.2. View ENIs

You can view the list of created ENIs.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > ENIs**.
3. In the top navigation bar, select an organization, a resource set, and a region. The created ENIs that match the specified criteria are displayed.
4. Select a filter option from the drop-down list, enter the relevant information in the search bar, and click **Search**.

You can select multiple filter options to narrow down search results.

| Filter option | Description |
|-------------------|---|
| ENI Name | Enter an ENI name to search for the ENI. |
| ENI ID | Enter an ENI ID to search for the ENI. |
| VSwitch ID | Enter a VSwitch ID to search for the ENIs that are associated with the VSwitch. |
| Security Group ID | Enter a security group ID to search for the ENIs that belong to the security group. |
| Instance ID | Enter an instance ID to search for the ENIs that are bound to the instance. |

2.10.3. Modify a secondary ENI

You can modify the attributes of a secondary elastic network interface (ENI), including the name, security group, and description.

Prerequisites

The secondary ENI is in the **Available** state.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > ENIs**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the secondary ENI and click **Modify** in the **Actions** column.
5. In the Modify ENI dialog box that appears, modify the name, security group, and description of the ENI.
6. Click **OK**.

2.10.4. Bind a secondary ENI to an instance

You can bind a secondary elastic network interface (ENI) to an instance. After the ENI is bound to the instance, the instance can process the traffic on the ENI.

Prerequisites

- The secondary ENI is in the **Available** state.
- The instance to which you want to bind the secondary ENI is in the **Running** or **Stopped** state.
- The instance and the secondary ENI belong to the same VPC.
- The VSwitch with which the secondary ENI is associated is in the same zone as the VSwitch to which the instance is connected. An ENI can be bound only to an instance in the same zone. The VSwitches of the ENI and of the instance can be different but must be in the same zone.

Context

The following limits apply when you bind an ENI to an instance:

- You can manually bind only secondary ENIs. Primary ENIs share the same lifecycle as instances and cannot be manually bound.
- An ENI can only be bound to a single ECS instance. However, an ECS instance can be bound with multiple ENIs. The maximum number of ENIs that can be bound to an instance depends on the instance type. For more

information about the number of ENIs that can be bound to an instance of each instance type, see Instance families in *ECS Product Introduction*.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > ENIs**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the target secondary ENI and click **Bind** in the **Actions** column.
5. In the Bind dialog box that appears, select an instance and click **OK**.

Result

In the **Status/Creation Time** column, the status of the secondary ENI changes to **Bound**.

2.10.5. Unbind a secondary ENI from an instance

You can unbind a secondary elastic network interface (ENI) from an instance. After the secondary ENI is unbound from the instance, the instance no longer processes the traffic on the ENI.

Prerequisites

- The secondary ENI is in the **Bound** state.
- The instance is in the **Running** or **Stopped** state.

Context

Only secondary ENIs can be unbound. Primary ENIs share the same lifecycle as instances and cannot be unbound.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > ENIs**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the secondary ENI and click **Unbind** in the **Actions** column.
5. Click **OK**.

Result

In the **Status/Creation Time** column, the status of the secondary ENI changes to **Available**.

2.10.6. Delete a secondary ENI

You can delete a secondary elastic network interface (ENI) that is no longer needed.

Prerequisites

The secondary ENI is in the **Available** state.

Context

You can delete only secondary ENIs. Primary ENIs share the same lifecycle as instances and cannot be deleted.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, choose **Networks and Security > ENIs**.
3. In the top navigation bar, select an organization, a resource set, and a region.

4. Find the secondary ENI and click **Delete** in the **Actions** column.
5. Click **OK**.

2.11. Deployment sets

2.11.1. Create a deployment set

You can use a deployment set to distribute or aggregate instances involved in your business. You can select hosts, racks, or network switches to improve service availability or network performance based on your needs.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Deployment Sets**.
3. Click **Create Deployment Set**.
4. Configure the parameters of the deployment set.

| Category | Parameter | Required | Description |
|----------------|-------------------|----------|---|
| Region | Organization | Yes | The organization to which the deployment set belongs. |
| | Resource Set | Yes | The resource set to which the deployment set belongs. |
| | Region | Yes | The region where the deployment set is located. |
| | Zone | Yes | The zone where the deployment set is located. |
| Basic Settings | Deployment Domain | Yes | This parameter setting determines the Deployment Target options. Valid values: <ul style="list-style-type: none"> ◦ Default: When Default is selected, the deployment target options are Host, Rack, and Network Switch. ◦ Switch: When Switch is selected, the deployment target options are Host and Rack. |
| | Deployment Target | Yes | The basic unit that can be scheduled when you deploy instances. <ul style="list-style-type: none"> ◦ Host: Instances are distributed or aggregated at the host level. ◦ Rack: Instances are distributed or aggregated at the rack level. ◦ VSwitch: Instances are distributed or aggregated at the VSwitch level. |

| Category | Parameter | Required | Description |
|----------|---------------------|----------|--|
| | Deployment Policy | Yes | The dispersion policies are used to improve service availability to avoid business impact when a host, rack, or switch fails. The aggregation policies are used to improve network performance to minimize the access latency between instances. Options are: <ul style="list-style-type: none"> Loose Dispersion Strict Dispersion Loose Aggregation Strict Aggregation |
| | Deployment Set Name | No | The name of the deployment set. The name must be 2 to 128 characters in length. It must start with a letter but cannot start with http:// or https://. It can contain digits, periods (.), underscores (_), hyphens (-), colons (:), and commas (,). |
| | Description | No | The description of the deployment set. We recommend that you provide an informational description to simplify future management operations. The description must be 2 to 256 characters in length. It must start with a letter but cannot start with http:// or https://. It can contain digits, periods (.), underscores (_), hyphens (-), colons (:), and commas (,). |

- Click **Submit**.

2.11.2. View deployment sets

You can view the list of created deployment sets.

Procedure

- Log on to the [ECS console](#).
- In the left-side navigation pane, click **Deployment Sets**.
- In the top navigation bar, select an organization, a resource set, and a region. The created deployment sets that match the specified criteria are displayed.
- Select a filter option from the drop-down list, enter the relevant information in the search bar, and then click **Search**.

You can select multiple filter options to narrow down search results.

| Filter option | Description |
|---------------------|--|
| Deployment Set Name | Enter a deployment set name to search for the deployment set. |
| Deployment Set ID | Enter a deployment set ID to search for the deployment set. |
| Resource Set | Enter a resource set name to search for the deployment sets that belong to the resource set. |

2.11.3. Modify a deployment set

You can modify the name and description of a deployment set.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Deployment Sets**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the deployment set and click **Modify** in the **Actions** column.
5. In the Change Deployment Set dialog box, change the name of the deployment set.
6. Click **OK**.

2.11.4. Delete a deployment set

You can delete a deployment set that is no longer needed.

Prerequisites

No instances exist in the deployment set.

Procedure

1. [Log on to the ECS console](#).
2. In the left-side navigation pane, click **Deployment Sets**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the deployment set and click **Delete** in the **Actions** column.
5. Click **OK**.

2.12. Install FTP software

2.12.1. Overview

File Transfer Protocol (FTP) transfers files between a client and a server by establishing two TCP connections. One is the command link for transferring commands between a client and a server. The other is the data link used to upload or download data. Before uploading files to an instance, you must build an FTP site for the instance.

2.12.2. Install and configure vsftpd in CentOS

This topic describes how to install and configure vsftpd in CentOS to transfer files.

Procedure

1. Install vsftpd.

```
yum install vsftpd -y
```

2. Add an FTP account and a directory.
 - i. Check the location of the *nologin* file, which is usually under the */usr/sbin* or */sbin* directory.

- ii. Create an FTP account.

Run the following commands to create the `/alidata/www/wwwroot` directory and specify this directory as the home directory of the account `pwftp`. You can also customize the account name and directory.

```
mkdir -p /alidata/www/wwwroot
useradd -d /alidata/www/wwwroot -s /sbin/nologin pwftp
```

- iii. Modify the account password.

```
passwd pwftp
```

- iv. Modify the permissions on the specified directory.

```
chown -R pwftp.pwftp /alidata/www/wwwroot
```

3. Configure vsftpd.

- i. Open the vsftpd configuration file.

```
vi /etc/vsftpd/vsftpd.conf
```

- ii. Change the value of `anonymous_enable` from `YES` to `NO`.
- iii. Delete the comment delimiter (`#`) from the following configuration lines:

```
local_enable=YES
write_enable=YES
chroot_local_user=YES
```

- iv. Press the Esc key to exit the edit mode, and enter `:wq` to save the modifications and exit.

4. Modify the shell configuration.

- i. Open the shell configuration file.

```
vi /etc/shells
```

- ii. If the file does not contain `/usr/sbin/nologin` or `/sbin/nologin`, add it to the file.

5. Start vsftpd and perform a logon test.

- i. Start vsftpd.

```
service vsftpd start
```

- ii. Use the account `pwftp` to perform an FTP logon test.

This example uses the directory `/alidata/www/wwwroot`.

2.12.3. Install vsftpd in Ubuntu or Debian

This topic describes how to install and configure vsftpd in an instance running Ubuntu or Debian to transfer files.

Procedure

1. Update the software source.

```
apt-get update
```

2. Install vsftpd.

```
apt-get install vsftpd -y
```

3. Add an FTP account and a directory.

- i. Check the location of the `nologin` file, which is typically under the `/usr/sbin` or `/sbin` directory.

- ii. Create an FTP account.

Run the following commands to create the `/alidata/www/wwwroot` directory and specify this directory as the home directory of the account `pwftp`. You can also customize the account name and directory.

```
mkdir -p /alidata/www/wwwroot
useradd -d /alidata/www/wwwroot -s /sbin/nologin pwftp
```

- iii. Modify the account password.

```
passwd pwftp
```

- iv. Modify the permissions on the specified directory.

```
chown -R pwftp.pwftp /alidata/www/wwwroot
```

4. Configure vsftpd.

- i. Open the vsftpd configuration file.

```
vi /etc/vsftpd.conf
```

- ii. Change the value of `anonymous_enable` from `YES` to `NO`.
- iii. Delete the comment delimiter (`#`) from the following configuration lines:

```
local_enable=YES
write_enable=YES
chroot_local_user=YES
chroot_list_enable=YES
chroot_list_file=/etc/vsftpd.chroot_list
```

- iv. Press the `Esc` key to exit the edit mode, and enter `:wq` to save the modifications and exit.
- v. Open the `/etc/vsftpd.chroot_list` file and add the FTP account name to the file. Save the modifications and exit.

You can follow steps a to d to open and save the file.

5. Modify shell configurations.

- i. Open the shell configuration file.

```
vi /etc/shells
```

- ii. If the file does not contain `/usr/sbin/nologin` or `/sbin/nologin`, add it to the file.

6. Start vsftpd and perform a logon test.

- i. Start vsftpd.

```
service vsftpd restart
```

- ii. Use the account `pwftp` to perform an FTP logon test.

This example uses the directory `/alidata/www/wwwroot`.

2.12.4. Build an FTP site in Windows Server 2008

This topic describes how to build an FTP site on an instance running Windows Server 2008.

Prerequisites

You have added the Web Server (IIS) role and installed FTP on an instance.

Procedure

1. [Connect to an instance.](#)

2. Choose **Start > Administrative Tools > Internet Information Services (IIS) Manager**.
3. Right-click the server name and select **Add FTP Site** from the shortcut menu.
4. Enter an FTP site name and a physical path, and then click **Next**.
5. Set **IP Address** to **All Unassigned** and **SSL** to **No SSL**, and then click **Next**.
6. Set **Authentication** to **Basic**, **Authorization** to **All Users**, and **Permissions** to **Read and Write**, and click **Finish**.

Result

Then you can use the administrator account and password to upload and download files through FTP. Make sure that the following conditions are met:

- The port for the FTP site is not in use by other applications, and Windows firewall is not blocking the port.
- The security group of the instance contains a security group rule that allows inbound access to the FTP port.

2.12.5. Build an FTP site in Windows Server 2012

This topic describes how to build an FTP site on an instance running Windows Server 2012.

Prerequisites

You have added the Web Server (IIS) role and installed FTP on an instance.

Procedure

1. [Connect to an instance](#).
2. Click the **Server Manager** icon.
3. In the left-side navigation pane, click **IIS**.
4. In the **Server** area, right-click the server name and select **Internet Information Services (IIS) Manager** from the shortcut menu.
5. Right-click the server name and select **Add FTP Site** from the shortcut menu.
6. Enter an FTP site name and a physical path, and then click **Next**.
7. Set **IP Address** to **All Unassigned** and **SSL** to **No SSL**, and then click **Next**.
8. Set **Authentication** to **Basic**, **Authorization** to **All Users**, and **Permissions** to **Read and Write**, and click **Finish**.

Result

Then you can use the administrator account and password to upload and download files through FTP. Make sure that the following conditions are met:

- The port for the FTP site is not in use by other applications, and Windows firewall is not blocking the port.
- The security group of the instance contains a security group rule that allows inbound access to the FTP port.

3. Container Service

3.1. Container Service support for Kubernetes

1.18

Container Service strictly conforms to the terms of the Certified Kubernetes Conformance Program. This topic describes the changes that Container Service has made to support Kubernetes 1.18.

Version upgrades

Container Service has upgraded and optimized all of its components to support Kubernetes 1.18.8.

| Key component | Version | Description |
|---------------|-----------------------------|--|
| Kubernetes | 1.18.8 | Some frequently used API versions are deprecated in Kubernetes 1.18. Before you upgrade a Kubernetes cluster, we recommend that you upgrade the deprecated API versions that are listed in this topic. |
| Docker | 19.03.5 (containerd 1.2.10) | No. |
| etcd | 3.4.3 | No. |
| CoreDNS | 1.6.7 | No. |

Version details

- **Resource changes and deprecation**

The following APIs are deprecated in Kubernetes 1.18:

- The APIs `apps/v1beta1` and `apps/v1beta2` of all the resources are replaced by `apps/v1`.
- The API extensions/`v1beta1` of `DaemonSets`, `Deployments`, and `ReplicaSets` is replaced by `apps/v1`.
- The API extensions/`v1beta1` of `NetworkPolicies` resources is replaced by `networking.k8s.io/v1`.
- The API extensions/`v1beta1` of pod security policies is replaced by `policy/v1beta1`.

The label that specifies the regions of a node is changed to `topology.kubernetes.io/region`. The label that specifies the zone of a node is changed to `topology.kubernetes.io/zone`. We recommend that you update the related configurations for your workloads.

- **Feature upgrades**

- **Server-side Apply Beta 2** is introduced. You can view the relationships between the configuration items of a resource in the `metadata.managedFields` field of the resource.
- The **Node Local DNS Cache** feature is released to improve the DNS availability and performance of your cluster.
- The **Volume Snapshot** feature is in public preview and supports operations such as data volume backup, recovery, and scheduled backup.

Container Service upgrades for Kubernetes 1.18.8

In Kubernetes 1.18.8, Container Service enables the following feature in the kubelet configuration file: Users who use raw data volumes can upgrade clusters without the need to drain the nodes.

3.2. What is Container Service?

Container Service provides high-performance, scalable, and enterprise-class management service for Kubernetes containerized applications throughout the application lifecycle.

Container Service simplifies the deployment and scaling operations on Kubernetes clusters. Integrated with services such as virtualization, storage, network, and security, Container Service aims to provide the optimal cloud environment for Kubernetes containerized applications. Alibaba Cloud is a Kubernetes Certified Service Provider (KCSP). As one of the first services to participate in the Certified Kubernetes Conformance Program, Container Service provides you with professional support and services.

3.3. ACK@Edge overview

ACK@Edge is released for commercial use. ACK@Edge is a cloud-managed solution that is provided by Container Service to coordinate cloud and edge computing. This topic describes the background and features of edge Kubernetes clusters.

Overview

With the rapid growth of smart devices connected to the Internet and the advent of 5G and IoT, computing and storage services provided by traditional cloud computing platforms can no longer satisfy the needs of edge devices for time-efficient computing, larger storage capacity, and enhanced computing capacity. Edge Kubernetes clusters are intended for bringing cloud computing to edges (clients). Edge Kubernetes clusters can be created, managed, and maintained in the Container Service console. This is the trend of cloud computing.

An edge Kubernetes cluster is a standard, secure, and highly-available Kubernetes cluster deployed in the cloud. This type of cluster is integrated with features of Alibaba Cloud, such as virtualization, storage, networking, and security. This simplifies the management and maintenance of clusters and allows you to focus on your business development. ACK@Edge provides the following features:

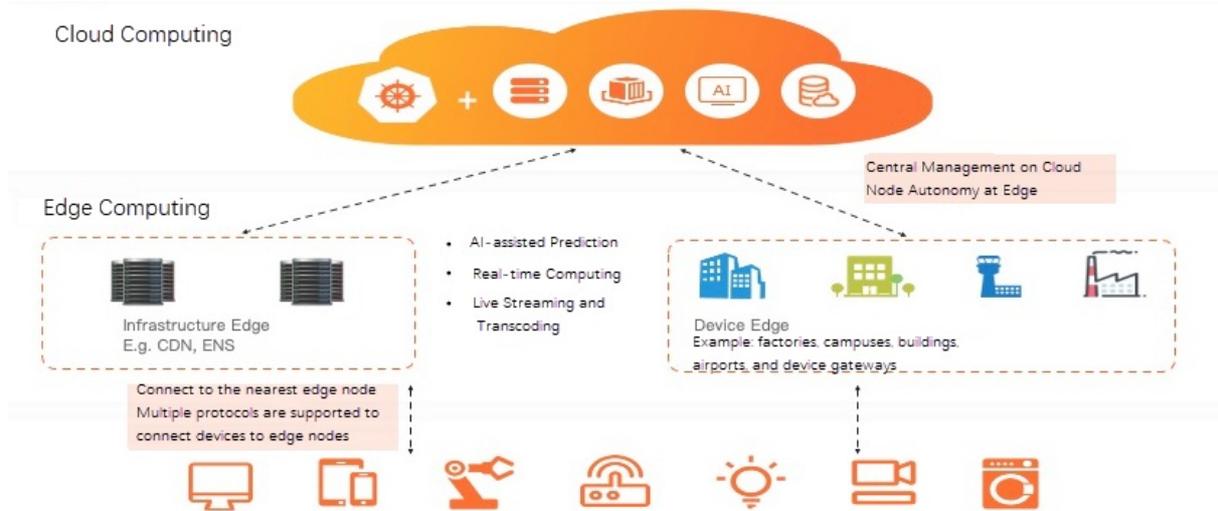
- Allows you to build a cloud-native infrastructure for edge computing with a few clicks.
- Allows you to quickly connect edge computing resources to the cloud. These resources include IoT gateway devices, terminals, Content Delivery Network (CDN) resources, and data centers.
- Applies to diverse scenarios, such as edge intelligence, intelligent buildings, intelligent factories, audio and video live streaming, online education, and CDN.

Edge Kubernetes clusters support features such as node autonomy, cell-based management, and native APIs for the management and maintenance of resources at the edge side. To use these features, you do not need to rewrite the logic of your services. This provides a native and centralized method for application lifecycle management and resource scheduling in edge computing scenarios.

Features

Edge Kubernetes clusters provide the following features to support lifecycle management for containerized applications and resources in edge computing scenarios:

- Allows you to create highly available edge Kubernetes clusters with a few clicks and provides lifecycle management on edge Kubernetes clusters, such as scaling cloud nodes, adding edge nodes to clusters, upgrading, logging, and monitoring. You can perform the preceding operations in the Container Service console.
- Supports access to various heterogeneous resources, such as data centers and IoT devices. Hybrid scheduling of heterogeneous resources is also supported.
- Supports node autonomy and network autonomy to ensure the reliability of edge nodes and services in edge computing scenarios where the network connection is weak.
- Supports reverse tunneling for management and maintenance of edge nodes.



3.4. Planning and preparation

Before you start using Container Service, you need to create cloud resources such as VPC networks, VSwitches, disks, and OSS buckets based on your application requirements.

Before you create a Kubernetes cluster, make the following preparations:

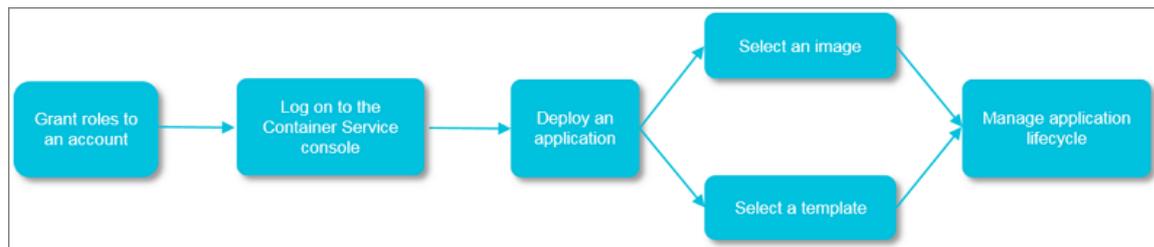
- **Create a VPC network (optional)**
To create a cluster in an existing VPC network, you must create the VPC network and VSwitches in advance.
- **Create a volume (optional)**
To create a stateful application with network storage, you must create disks or OSS buckets in advance.

3.5. Quick start

3.5.1. Procedure

You can perform the following steps to use the Container Service service.

The following diagram shows the procedure to use the Container Service service.



Step 1: Authorize the default role

Authorize the default role of Container Service to perform operations on the resources that belong to the specified organization.

Step 2: Log on to the Container Service console

Log on to the Container Service console. For more information, see [Log on to the Container Service console](#).

Step 3: Create an Container Service cluster

Set the network environment and the number of nodes, and configure node details.

Step 4: Deploy an application by using an image or orchestration template

You can use an existing image or orchestration template, or create a new image or orchestration template. To create an application that consists of services based on different images, use an orchestration template.

Step 5: Manage the application lifecycle

3.5.2. Log on to the Container Service console

You can perform the following steps to log on to the Container Service console.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

Note When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top navigation bar, choose **Products > Elastic Computing > Container Service for Kubernetes**.
5. Select the required organization and region.
6. Click **ACK** to go to the Container Service console.

3.5.3. Create a Kubernetes cluster

This topic describes how to create a Kubernetes cluster in the Container Service console.

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**. On the Clusters page that appears, click **Create Kubernetes Cluster** in the upper-right corner.
3. On the **Create Kubernetes Cluster** page, set the parameters.

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|--------------------|--|
| Cluster Name | <p>Enter a name for the cluster. The name must be 1 to 63 characters in length, and can contain digits, letters, and hyphens (-).</p> <p> Note The cluster name must be unique among clusters that belong to the same Alibaba Cloud account.</p> |
| Region | Select the region where you want to deploy the cluster. |
| VPC | <p>You can select a virtual private cloud (VPC) from the drop-down list.</p> <ul style="list-style-type: none">◦ If the specified VPC is already associated with a NAT gateway, the cluster uses this NAT gateway.◦ Otherwise, the system automatically creates a NAT gateway. If you do not want the system to create a NAT gateway, clear Configure SNAT for VPC. <p> Note If you disable the system to automatically create a NAT gateway and want the VPC to access the Internet, you must manually associate the VPC with a NAT gateway or create Source Network Address Translation (SNAT) rules for the VPC.</p> |
| VSwitch | <p>Select one or more vSwitches for the cluster.</p> <p>You can select up to three vSwitches that are deployed in different zones.</p> |
| Kubernetes Version | Select a Kubernetes version. |
| Container Runtime | You can select Docker or Sandboxed-Container. |
| Billing Method | Only pay-as-you-go nodes are supported. |

| Parameter | Description |
|-----------------------|--|
| Master Configurations | <p>Set the Instance Type and System Disk parameters:</p> <ul style="list-style-type: none"> ◦ Master Node Quantity: You can add up to three master nodes. ◦ Instance Type: You can select one or more instance types. For more information, see the <i>Instance types</i> chapter of <i>ECS User Guide</i>. ◦ System Disk: SSD Disk and Ultra Disk are supported. <div style="background-color: #e1f5fe; padding: 5px; margin-top: 10px;"> <p> Note You can select Enable Backup to back up disk data.</p> </div> |
| Worker Instance | <p>You can select Create Instance or Add Existing Instance.</p> |
| Worker Configurations | <p>If Worker Instance is set to Create Instance, set the following parameters:</p> <ul style="list-style-type: none"> ◦ Instance Type: You can select one or more instance types. For more information, see the <i>Instance types</i> chapter of <i>ECS User Guide</i>. ◦ Selected Types: The selected instance types are displayed. ◦ Quantity: Set the number of worker nodes. ◦ System Disk: SSD Disk and Ultra Disk are supported. <div style="background-color: #e1f5fe; padding: 5px; margin-top: 10px;"> <p> Note You can select Enable Backup to back up disk data.</p> </div> <ul style="list-style-type: none"> ◦ Mount Data Disk: SSD Disk and Ultra Disk are supported. <div style="background-color: #e1f5fe; padding: 5px; margin-top: 10px;"> <p> Note</p> <ul style="list-style-type: none"> ▪ You can select Encrypt Disk to encrypt disks. ▪ You can select Enable Backup to back up disk data. </div> |
| Operating System | <p>The CentOS and Aliyun Linux operating systems are supported.</p> |

| Parameter | Description |
|---------------------------------|--|
| Password | <p>Set a password that is used to log on to the nodes.</p> <p>Note The password must be 8 to 30 characters in length, and must contain at least three of the following types of characters: uppercase letters, lowercase letters, digits, and special characters.</p> |
| Confirm Password | Enter the password again. |
| Network Plug-in | Flannel and Terway are supported. By default, Flannel is selected. |
| Pod CIDR Block and Service CIDR | <p>These parameters are optional. For more information, see <i>Network planning</i> in <i>VPC User Guide</i>.</p> <p>Note These parameters are available only when you select an existing VPC.</p> |
| Configure SNAT | This parameter is optional. If you clear Configure SNAT for VPC, you must create a NAT gateway or configure SNAT rules for the VPC. |
| Access to the Internet | <p>Specify whether to expose the API server with an elastic IP address (EIP). The Kubernetes API server provides multiple HTTP-based RESTful APIs that can be used to create, delete, modify, query, and watch resource objects such as pods and Services.</p> <ul style="list-style-type: none"> If you select this check box, an EIP is created and associated with an internal-facing Server Load Balancer (SLB) instance. Port 6443 used by the API server is exposed on the master nodes. You can connect to and manage the cluster by using kubeconfig over the Internet. If you clear this check box, no EIP is created. You can connect to and manage the cluster only by using kubeconfig within the VPC. |
| Ingress | Specify whether to Install Ingress Controllers . By default, Install Ingress Controllers is selected. |
| Log Service | If you enable Log Service, you can select an existing project or create a project. If you select Enable Log Service , the Log Service plug-in is automatically installed in the cluster. If you select Create Ingress Dashboard , Ingress access logs are collected and displayed on dashboards. |
| Volume Plug-in | By default, CSI is selected. |
| Deletion Protection | This check box prevents you from accidentally deleting the cluster in the console or by calling API operations. |

| Parameter | Description |
|-----------------|---|
| RDS Whitelist | <p>Add the IP addresses of the nodes to the whitelist of the ApsaraDB RDS instance that is allowed to access the Kubernetes cluster.</p> <p>Note To enable an ApsaraDB RDS instance to access the Kubernetes cluster, you must deploy the ApsaraDB RDS instance in the same VPC as the Kubernetes cluster.</p> |
| Node Protection | This check box is selected by default to prevent you from accidentally deleting the nodes in the console or by calling API operations. |
| Label | Add labels to the cluster. |

4. Complete the advanced settings of the cluster.

| Parameter | Description |
|-----------------------|--|
| IP Addresses per Node | The number of IP addresses that are assigned to a node. |
| Custom Image | You can select a custom image. After you select a custom image, all nodes in the cluster are deployed by using this image. |
| Kube-proxy Mode | <p>iptables and IPVS are supported.</p> <ul style="list-style-type: none"> iptables is a tested and stable kube-proxy mode. It uses iptables rules to conduct service discovery and load balancing. The performance of this mode is limited by the size of the Kubernetes cluster. This mode is suitable for Kubernetes clusters that manage a small number of Services. IPVS is a high-performance kube-proxy mode. It uses Linux Virtual Server (LVS) to conduct service discovery and load balancing. This mode is suitable for Kubernetes clusters that manage a large number of Services. We recommend that you use this mode in scenarios where high-performance load balancing is required. |
| Custom Node Name | Specify whether to use a custom node name. |
| Node Port Range | Specify the value of Node Port Range . |
| Taints | Add taints to all worker nodes in the Kubernetes cluster. |
| CPU Policy | <p>Specify the CPU policy. Valid values:</p> <ul style="list-style-type: none"> None: indicates that the default CPU affinity is used. This is the default policy. Static: allows pods with specific resource characteristics on the node to be granted with enhanced CPU affinity and exclusivity. |

| Parameter | Description |
|----------------|---|
| Cluster Domain | The default domain name of the cluster is cluster.local. You can specify a custom domain name. |
| Cluster CA | Specify whether to enable the cluster CA certificate. |
| User Data | <p>Customize the startup behaviors of ECS instances and import data to the ECS instances. The user data can be used to perform the following operations:</p> <ul style="list-style-type: none"> ◦ Run scripts during instance startup. ◦ Import user data as normal data to an ECS instance for future reference. |

5. Click **Create Cluster** in the upper-right corner of the page.
6. On the **Confirm** page, after all check items are verified, select the terms of service and disclaimer and click **OK** to start the deployment.

Result

After the cluster is created, you can find the cluster on the **Clusters** page in the Container Service console.

3.5.4. Create an application from an orchestration template

Container Service provides orchestration templates that you can use to create applications. You can also modify the templates based on YAML syntax to customize applications.

Prerequisites

A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).

Context

This topic describes how to use an orchestration template to create an NGINX application that consists of a Deployment and a Service. The Deployment provisions pods for the application and the Service manages access to the pods at the backend.

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Workloads > Deployments**.
5. On the **Deployments** page, select the namespace and click **Create from YAML** in the upper-right corner.
6. Configure the template and click **Create**.
 - **Sample Template**: Container Service provides YAML templates of various resource types to help you quickly deploy resource objects. You can also create a custom template based on YAML syntax to describe the resource that you want to define.
 - **Add Deployment**: This feature allows you to quickly define a YAML template.
 - **Use Existing Template**: You can import an existing template to the configuration page.

Based on an orchestration template provided by Container Service, the following sample template is used to create a Deployment of an NGINX application.

Note Container Service supports YAML syntax. You can use the `---` symbol to separate multiple resource objects. This allows you to create multiple resource objects in a single template.

```

apiVersion: apps/v1beta2 # for versions before 1.8.0 use apps/v1beta1
kind: Deployment
metadata:
  name: nginx-deployment
  labels:
    app: nginx
spec:
  replicas: 2
  selector:
    matchLabels:
      app: nginx
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
        - name: nginx
          image: nginx:1.7.9 # replace it with your exactly <image_name:tags>
          ports:
            - containerPort: 80
---
apiVersion: v1 # for versions before 1.8.0 use apps/v1beta1
kind: Service
metadata:
  name: my-service1 #TODO: to specify your service name
  labels:
    app: nginx
spec:
  selector:
    app: nginx #TODO: change label selector to match your backend pod
  ports:
    - protocol: TCP
      name: http
      port: 30080 #TODO: choose an unique port on each node to avoid port conflict
      targetPort: 80
  type: LoadBalancer ##In this example, the type is changed from Nodeport to LoadBalancer.

```

- Click **Create**. A message that indicates the deployment status appears.

After the application is created, choose **Services and Ingresses > Services** in the left-side navigation pane. On the Services page, you can find that a Service named `my-service1` is created for the application. The external endpoint of the Service is also displayed on the page. Click the endpoint in the **External Endpoint** column.

- You can visit the NGINX welcome page in the browser.



3.6. Kubernetes clusters

3.6.1. Authorizations

3.6.1.1. Assign RBAC roles to a RAM user

This topic describes how to assign role-based access control (RBAC) roles to Resource Access Management (RAM) users. By default, RBAC is enabled for Kubernetes 1.6 and later. RBAC is important for you to manage clusters. You can use RBAC to specify the types of operations that are allowed for specific users based on their roles in an organization.

Procedure

1. [Log on to the Container Service console.](#)
2. In the left-side navigation pane, click **Authorizations**.
3. On the **RAM Users** tab, select the RAM user to which you want to grant permissions and click **Modify Permissions**.

Note If you log on to the Container Service console as a RAM user, make sure that the RAM user is assigned the predefined RBAC administrator role or the cluster-admin role.

4. On the **Configure Role-Based Access Control (RBAC)** wizard page, click **Add Permissions** to add cluster-scoped or namespace-scoped permissions and select a predefined or custom RBAC role in the Permission column. You can also click the minus sign (-) to delete permissions. After you add the permissions, click **Next Step**.

Note For each RAM user, you can assign only one predefined RBAC role but one or more custom RBAC roles to manage the same cluster or namespace.

The following table describes the permissions that the predefined and custom RBAC roles have on clusters and namespaces.

Roles and permissions

| Role | RBAC permissions on cluster resources |
|-----------------|--|
| Administrator | Read and write permissions on resources in all namespaces. |
| O&M Engineer | Read and write permissions on resources in all namespaces and read-only permissions on nodes, persistent volumes (PVs), namespaces, and service quotas within a cluster. |
| Developer | Read and write permissions on resources in a specified namespace or all namespaces. |
| Restricted User | Read-only permissions on resources in a specified namespace or all namespaces. |
| Custom | The cluster role that you select for a custom role determines what permissions the custom role has. Before you select a cluster role, make sure that you are aware of the permissions that the cluster role has in case the RAM user is granted excessive permissions. |

After the authorization is complete, you can use the account of the specified RAM user to log on to the Container Service console. For more information, see [Log on to the Container Service console](#).

Predefined and custom RBAC roles

Container Service provides the following predefined RBAC roles: administrator, O&M engineer, developer, and restricted user. You can use these roles to regulate Container Service access control in most scenarios. In addition, you can use custom roles to customize permissions on clusters.

Container Service provides a set of custom RBAC roles.

Note The cluster-admin role is similar to a super administrator. By default, the cluster-admin role has the permissions to manage all resources within a cluster.

The screenshot shows the 'Authorizations' page in the console. It has a progress bar with three steps: 'Select RAM User', 'Configure Role-Based Access Control (RBAC)', and 'Submit Authorization'. The 'Configure RBAC' step is active. Below the progress bar, there's a section for 'Modify Permissions' with a sub-header 'Manage the permissions of the sub-account, you can add new permissions, delete / modify existing permissions'. A 'RAM User: loginName' field is present. Below that, there's a table with columns 'Cluster/namespace' and 'Permission'. The 'Cluster' dropdown is set to 'All Clusters'. The 'Permission' section has radio buttons for 'Administrator' (selected), 'O&M Engineer', 'Developer', and 'Restricted User'. An 'Add Permissions' button is below the table. A 'Permission Description' section follows, listing roles and their descriptions: Administrator (Read/write access to resources in all namespaces...), O&M Engineer (Read/write access to resources that are both visible in the console and in all namespaces...), Developer (Read/write access to resources that are both visible in the console and in the specified namespaces...), Restricted User (Read-only access to resources that are both visible in the console and in the specified namespaces...), and Custom (Different cluster roles have different permissions...).

You can log on to a master node of a cluster and run the following command to view the custom RBAC roles that are assigned to the current account :

```
# kubectl get clusterrole
```

```
# kubectl get clusterrole
NAME                AGE
admin               13d
alibaba-log-controller      13d
alicloud-disk-controller-runner  13d
cluster-admin        13d
cs:admin            13d
edit                13d
flannel             13d
kube-state-metrics   22h
node-exporter       22h
prometheus-k8s      22h
prometheus-operator  22h
system:aggregate-to-admin  13d
....
system:volume-scheduler  13d
view                13d
```

Run the following command to view the details of a role, for example, the cluster-admin role:

```
# kubectl get clusterrole cluster-admin -o yaml
```

 **Notice** After a RAM user is assigned the cluster-admin role, the RAM user has the same permissions as the Alibaba Cloud account to which the RAM user belongs. The RAM user has full control over all resources within the cluster. Proceed with caution.

```
# kubectl get clusterrole cluster-admin -o yaml
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  annotations:
    rbac.authorization.kubernetes.io/autoupdate: "true"
  creationTimestamp: 2018-10-12T08:31:15Z
  labels:
    kubernetes.io/bootstrapping: rbac-defaults
  name: cluster-admin
  resourceVersion: "57"
  selfLink: /apis/rbac.authorization.k8s.io/v1/clusterroles/cluster-admin
  uid: 2f29f9c5-cdf9-11e8-84bf-00163e0b2f97
rules:
- apiGroups:
  - "*"
  resources:
  - "*"
  verbs:
  - "*"
- nonResourceURLs:
  - "*"
  verbs:
  - "*"
```

3.6.2. Clusters

3.6.2.1. Create a Kubernetes cluster

This topic describes how to create a Kubernetes cluster in the Container Service console.

Procedure

1. [Log on to the Container Service console.](#)
2. In the left-side navigation pane, click **Clusters**. On the Clusters page, click **Create Kubernetes Cluster** in the upper-right corner.
3. On the **Create Kubernetes Cluster** page, set the parameters.

| Parameter | Description |
|--------------|---|
| Cluster Name | Enter a name for the cluster. The name must be 1 to 63 characters in length, and can contain digits, letters, and hyphens (-).  Note The names of the clusters that belong to the same Alibaba Cloud account must be unique. |
| Region | Select the region where you want to deploy the cluster. |

| Parameter | Description |
|-----------------------|--|
| VPC | <p>You can select a virtual private cloud (VPC) from the drop-down list.</p> <ul style="list-style-type: none"> ◦ If the specified VPC is already associated with a NAT gateway, the cluster uses the associated NAT gateway. ◦ Otherwise, the system automatically creates a NAT gateway. If you do not want the system to create a NAT gateway, clear Configure SNAT for VPC. <p>Note If you disable the system to automatically create a NAT gateway and want the VPC to access the Internet, you must manually associate the VPC with a NAT gateway or create Source Network Address Translation (SNAT) rules for the VPC.</p> |
| VSwitch | <p>Select one or more vSwitches for the cluster.</p> <p>You can select at most three vSwitches that are deployed in different zones.</p> |
| Kubernetes Version | Select a Kubernetes version. |
| Container Runtime | You can select Docker or Sandboxed-Container. |
| Billing Method | Only pay-as-you-go nodes are supported. |
| Master Configurations | <p>Set the Instance Type and System Disk parameters:</p> <ul style="list-style-type: none"> ◦ Master Node Quantity: You can add at most three master nodes. ◦ Instance Type: You can select one or more instance types. For more information, see the <i>Instance types</i> chapter of <i>ECS User Guide</i>. ◦ System Disk: SSD Disk and Ultra Disk are supported. <p>Note You can select Enable Backup to back up disk data.</p> |
| Worker Instance | You can select Create Instance or Add Existing Instance . |

| Parameter | Description |
|---------------------------------|--|
| Worker Configurations | <p>If Worker Instance is set to Create Instance, set the following parameters:</p> <ul style="list-style-type: none"> Instance Type: You can select one or more instance types. For more information, see the <i>Instance types</i> chapter of <i>ECS User Guide</i>. Selected Types: The selected instance types are displayed. Quantity: Set the number of worker nodes. System Disk: SSD Disk and Ultra Disk are supported. <p>Note You can select Enable Backup to back up disk data.</p> <ul style="list-style-type: none"> Mount Data Disk: SSD Disk and Ultra Disk are supported. <p>Note</p> <ul style="list-style-type: none"> You can select Encrypt Disk to encrypt disks. You can select Enable Backup to back up disk data. |
| Operating System | The CentOS and Aliyun Linux operating systems are supported. |
| Password | <p>Set a password that is used to log on to the nodes.</p> <p>Note The password must be 8 to 30 characters in length, and must contain at least three of the following types of character: uppercase letters, lowercase letters, digits, and special characters.</p> |
| Confirm Password | Enter the password again. |
| Network Plug-in | Flannel and Terway are supported. By default, Flannel is selected. |
| Pod CIDR Block and Service CIDR | <p>These parameters are optional. For more information, see <i>Network planning</i> in <i>VPC User Guide</i>.</p> <p>Note These parameters are available only when you select an existing VPC.</p> |
| Configure SNAT | This parameter is optional. If you clear Configure SNAT for VPC, you must create a NAT gateway or configure SNAT rules for the VPC. |

| Parameter | Description |
|------------------------|---|
| Access to the Internet | <p>Specify whether to expose the API server with an elastic IP address (EIP). The Kubernetes API server provides multiple HTTP-based RESTful APIs that can be used to create, delete, modify, query, and watch resource objects such as pods and Services.</p> <ul style="list-style-type: none"> ◦ If you select this check box, an EIP is created and attached to an internal-facing Server Load Balancer (SLB) instance. Port 6443 used by the API server is exposed on the master nodes. You can connect to and manage the cluster by using kubeconfig over the Internet. ◦ If you clear this check box, no EIP is created. You can connect to and manage the cluster only by using kubeconfig from within the VPC. |
| Ingress | Specify whether to Install Ingress Controllers . By default, Install Ingress Controllers is selected. |
| Log Service | If you enable Log Service, you can select an existing project or create a project. If you select Enable Log Service , the Log Service plug-in is automatically installed in the cluster. If you select Create Ingress Dashboard , the Ingress access log is collected and displayed on dashboards. |
| Volume Plug-in | By default, CSI is selected. |
| Deletion Protection | This check box prevents you from accidentally deleting the cluster in the console or by calling API operations. |
| RDS Whitelist | <p>Add the IP addresses of the nodes to the whitelist of the ApsaraDB RDS instance that is allowed to access the Kubernetes cluster.</p> <div style="background-color: #e6f2ff; padding: 5px; border: 1px solid #d9e1f2;"> <p>Note To enable an ApsaraDB RDS instance to access the Kubernetes cluster, you must deploy the ApsaraDB RDS instance in the same VPC as the Kubernetes cluster.</p> </div> |
| Node Protection | This check box is selected by default to prevent you from accidentally deleting the nodes in the console or by calling API operations. |
| Label | Add labels to the cluster. |

4. Complete the advanced settings of the cluster.

| Parameter | Description |
|-----------------------|--|
| IP Addresses per Node | The number of IP addresses that are assigned to a node. |
| Custom Image | You can select a custom image. After you select a custom image, all nodes in the cluster are deployed by using this image. |

| Parameter | Description |
|------------------|--|
| Kube-proxy Mode | <p>iptables and IPVS are supported.</p> <ul style="list-style-type: none"> ◦ iptables is a tested and stable kube-proxy mode. It uses iptables rules to conduct service discovery and load balancing. The performance of this mode is limited by the size of the Kubernetes cluster. This mode is suitable for Kubernetes clusters that manage a small number of Services. ◦ IPVS is a high-performance kube-proxy mode. It uses Linux Virtual Server (LVS) to conduct service discovery and load balancing. This mode is suitable for Kubernetes clusters that manage a large number of Services. We recommend that you use this mode in scenarios where high-performance load balancing is required. |
| Custom Node Name | Specify whether to use a custom node name. |
| Node Port Range | Specify the value of Node Port Range . |
| Taints | Add taints to all worker nodes in the Kubernetes cluster. |
| CPU Policy | <p>Specify the CPU policy. Valid values:</p> <ul style="list-style-type: none"> ◦ None: indicates that the default CPU affinity is used. This is the default policy. ◦ Static: allows pods with specific resource characteristics on the node to be granted with enhanced CPU affinity and exclusivity. |
| Cluster Domain | The default domain name of the cluster is cluster.local. You can specify a custom domain name. |
| Cluster CA | Specify whether to enable the cluster CA certificate. |
| User Data | <p>You can customize the startup behaviors of ECS instances and import data to the ECS instances. The user data can be used to perform the following operations:</p> <ul style="list-style-type: none"> ◦ Run scripts during instance startup. ◦ Import user data as normal data to an ECS instance for future reference. |

5. Click **Create Cluster** in the upper-right corner of the page.
6. On the **Confirm** page, after all check items are verified, select the terms of service and disclaimer and click **OK** to start the deployment.

Result

After the cluster is created, you can view the created cluster on the **Clusters** page in the Container Service console.

3.6.2.2. Create an edge Kubernetes cluster

Edge Kubernetes clusters are intended for bringing cloud computing to edges (clients). Edge Kubernetes clusters can be created, managed, and maintained in the Container Service console. Container Service is a platform built on top of the edge computing infrastructure. It is also integrated with cloud computing and edge computing. This topic describes how to create an edge Kubernetes cluster.

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**. On the Clusters page that appears, click **Create Kubernetes Cluster** in the upper-right corner.
3. On the Create Cluster page, click the **Managed Edge Kubernetes** tab and set the cluster parameters.

| Parameter | Description |
|--------------------|--|
| Cluster Name | <p>Enter a name for the cluster. The name must be 1 to 63 characters in length, and can contain digits, letters, and hyphens (-).</p> <p> Note The cluster name must be unique among clusters that belong to the same Alibaba Cloud account.</p> |
| Region | Select the region where you want to deploy the cluster. |
| VPC | <p>You can select a virtual private cloud (VPC) from the drop-down list.</p> <ul style="list-style-type: none"> ◦ If the specified VPC is already associated with a NAT gateway, the cluster uses this NAT gateway. ◦ Otherwise, the system automatically creates a NAT gateway. If you do not want the system to create a NAT gateway, clear Configure SNAT for VPC. <p> Note If you disable the system to automatically create a NAT gateway and want the VPC to access the Internet, you must manually associate the VPC with a NAT gateway or create Source Network Address Translation (SNAT) rules for the VPC.</p> |
| VSwitch | <p>Select one or more vSwitches for the cluster.</p> <p>You can select up to three vSwitches that are deployed in different zones.</p> |
| Kubernetes Version | Select a Kubernetes version. |

| Parameter | Description |
|-----------------------|---|
| Master Configurations | <p>Set the Instance Type and System Disk parameters:</p> <ul style="list-style-type: none"> Master Node Quantity: You can add up to three master nodes. Instance Type: You can select one or more instance types. For more information, see the <i>Instance types</i> chapter of <i>ECS User Guide</i>. System Disk: SSD Disk and Ultra Disk are supported. <p>Note You can select Enable Backup to back up disk data.</p> |
| Worker Instance | <p>You can select Create Instance or Add Existing Instance.</p> |
| Worker Configurations | <p>If Worker Instance is set to Create Instance, set the following parameters:</p> <ul style="list-style-type: none"> Instance Type: You can select one or more instance types. For more information, see the <i>Instance types</i> chapter of <i>ECS User Guide</i>. Selected Types: The selected instance types are displayed. Quantity: Set the number of worker nodes. By default, worker nodes are not required in edge Kubernetes clusters. Therefore, the value of this parameter is set to 0. System Disk: SSD Disk and Ultra Disk are supported. <p>Note You can select Enable Backup to back up disk data.</p> <ul style="list-style-type: none"> Mount Data Disk: SSD Disk and Ultra Disk are supported. <p>Note</p> <ul style="list-style-type: none"> You can select Encrypt Disk to encrypt disks. You can select Enable Backup to back up disk data. |
| Operating System | <p>The CentOS and Aliyun Linux operating systems are supported.</p> |
| Network Plug-in | <p>Edge Kubernetes clusters support only Flannel. You do not need to set this parameter.</p> |

| Parameter | Description |
|---------------------------------|--|
| Password | <p>Set a password that is used to log on to the nodes.</p> <p>Note The password must be 8 to 30 characters in length, and must contain at least three of the following types of character: uppercase letters, lowercase letters, digits, and special characters.</p> |
| Confirm Password | Enter the password again. |
| Pod CIDR Block and Service CIDR | <p>These parameters are optional. For more information, see <i>Network planning in VPC User Guide</i>.</p> <p>Note These parameters are available only when you select an existing VPC.</p> |
| Configure SNAT | <p>This parameter is optional. If you clear Configure SNAT for VPC, you must create a NAT gateway or configure SNAT rules for the VPC.</p> <p>Note For edge Kubernetes clusters, we recommend that you select Configure SNAT for VPC.</p> |
| Access to the Internet | <p>Specify whether to expose the API server with an elastic IP address (EIP). The Kubernetes API server provides multiple HTTP-based RESTful APIs that can be used to create, delete, modify, query, and watch resource objects such as pods and Services.</p> <p>Note We recommend that you expose the API server with an EIP.</p> <ul style="list-style-type: none"> ◦ If you select this check box, an EIP is created and associated with an internal-facing Server Load Balancer (SLB) instance. Port 6443 used by the API server is exposed on the master nodes. You can connect to and manage the cluster by using kubectl over the Internet. ◦ If you clear this check box, no EIP is created. You can connect to and manage the cluster only by using kubectl within the VPC. |

| Parameter | Description |
|---------------------|---|
| SSH Logon | <p>To enable Secure Shell (SSH) logon, you must first select Expose API Server with EIP.</p> <ul style="list-style-type: none"> ◦ If you select Use SSH to Access the Cluster from the Internet, you can access the cluster through SSH. ◦ If you clear Use SSH to Access the Cluster from the Internet, you cannot access the cluster through SSH or kubectl. If you want to access an Elastic Compute Service (ECS) instance in the cluster through SSH, you must manually bind an elastic IP address (EIP) to the ECS instance and configure security group rules to open SSH port 22. |
| Log Service | <p>If you enable Log Service, you can select an existing project or create a project. If you select Enable Log Service, the Log Service plug-in is automatically installed in the cluster. If you select Create Ingress Dashboard, Ingress access logs are collected and displayed on dashboards.</p> |
| Deletion Protection | <p>If you select this check box, the cluster cannot be deleted in the console or by calling API operations.</p> |
| Node Protection | <p>This check box is selected by default to prevent nodes from being accidentally deleted in the console or by calling API operations.</p> |

4. Complete advanced settings of the cluster.

| Parameter | Description |
|-----------------------|--|
| IP Addresses per Node | <p>The number of IP addresses that is assigned to a node. We recommend that you use the default value.</p> |
| Custom Image | <p>You can select a custom image. After you select a custom image, all nodes in the cluster are deployed by using this image.</p> |
| Kube-proxy Mode | <p>iptables and IPVS are supported.</p> <ul style="list-style-type: none"> ◦ iptables is a tested and stable kube-proxy mode. It uses iptables rules to conduct service discovery and load balancing. The performance of this mode is limited by the size of the Kubernetes cluster. This mode is suitable for Kubernetes clusters that manage a small number of Services. ◦ IPVS is a high-performance kube-proxy mode. It uses Linux Virtual Server (LVS) to conduct service discovery and load balancing. This mode is suitable for Kubernetes clusters that manage a large number of Services. We recommend that you use this mode in scenarios where high-performance load balancing is required. |
| Node Port Range | <p>Specify the value of Node Port Range.</p> |

| Parameter | Description |
|----------------|--|
| Taints | Add taints to all of the worker nodes in the cluster. We recommend that you do not add additional taints in case the system components cannot be deployed in the edge Kubernetes cluster. |
| CPU Policy | Specify the CPU policy. Valid values: <ul style="list-style-type: none"> None: indicates that the default CPU affinity is used. This is the default policy. Static: allows pods with specific resource characteristics on the node to be granted with enhanced CPU affinity and exclusivity. |
| Cluster Domain | The default domain name of the cluster is cluster.local. You can specify a custom domain name. |
| Cluster CA | Specify whether to enable the cluster CA certificate. |
| User Data | Customize the startup behaviors of ECS instances and import data to the ECS instances. The user data can be used to perform the following operations: <ul style="list-style-type: none"> Run scripts during instance startup. Import user data as normal data to an ECS instance for future reference. |

5. Click **Create Cluster** in the upper-right corner of the page.
6. On the **Confirm** page, after all check items are verified, select the terms of service and disclaimer and click **OK** to start the deployment.

Result

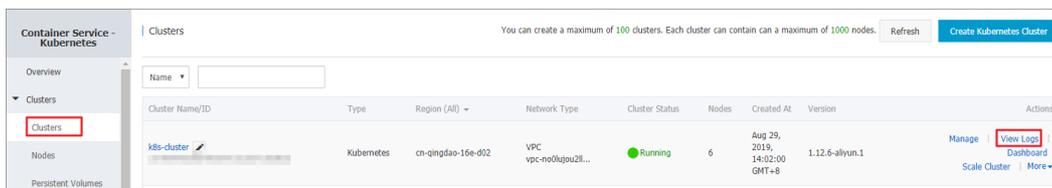
After the cluster is created, you can find the cluster on the **Clusters** page in the Container Service console.

3.6.2.3. View log files of a cluster

You can view the operation log of a cluster in the Container Service console.

Procedure

1. [Log on to the Container Service console.](#)
2. In the left-side navigation pane, click **Clusters**.
3. Find the cluster that you want to manage and click **View Logs** in the Actions column.



You can view operations performed on the cluster.

| Time | Description |
|------------------------------|--|
| Aug 29, 2019, 14:34:41 GMT+8 | start to update cluster status CREATE_COMPLETE |
| Aug 29, 2019, 14:26:37 GMT+8 | Stack CREATE completed successfully: |
| Aug 29, 2019, 14:02:02 GMT+8 | Successfully to CreateStack with response &ros.CreateStackResponse[Id:"02c493b2-29c2-496b-942f-e66873595298", Name:"k8s-for-cs-..."] |
| Aug 29, 2019, 14:02:02 GMT+8 | Start to wait stack ready |
| Aug 29, 2019, 14:02:00 GMT+8 | Start create cluster certificate |
| Aug 29, 2019, 14:02:00 GMT+8 | Start to validateCIDR |
| Aug 29, 2019, 14:02:00 GMT+8 | Start to create cluster task |
| Aug 29, 2019, 14:02:00 GMT+8 | Initial cluster info |
| Aug 29, 2019, 14:02:00 GMT+8 | Start to CreateK8sCluster |
| Aug 29, 2019, 14:02:00 GMT+8 | Start to CreateStack |

3.6.2.4. Connect to a cluster through kubectl

You can use the Kubernetes command line tool, [kubectl](#), to connect to a Kubernetes cluster from a local computer.

Procedure

1. Download the latest [kubectl](#) client from the [Kubernetes change log page](#).
2. Install and set up the [kubectl](#) client.
For more information, see [Install and set up kubectl](#).
3. Configure the cluster credentials.

You can use the `scp` command to securely copy the master node configuration file from the `/etc/kubernetes/kube.conf` directory of the master VM and paste it to the `$HOME/.kube/config` directory of the local computer, where the `kubectl` credentials are expected to be stored.

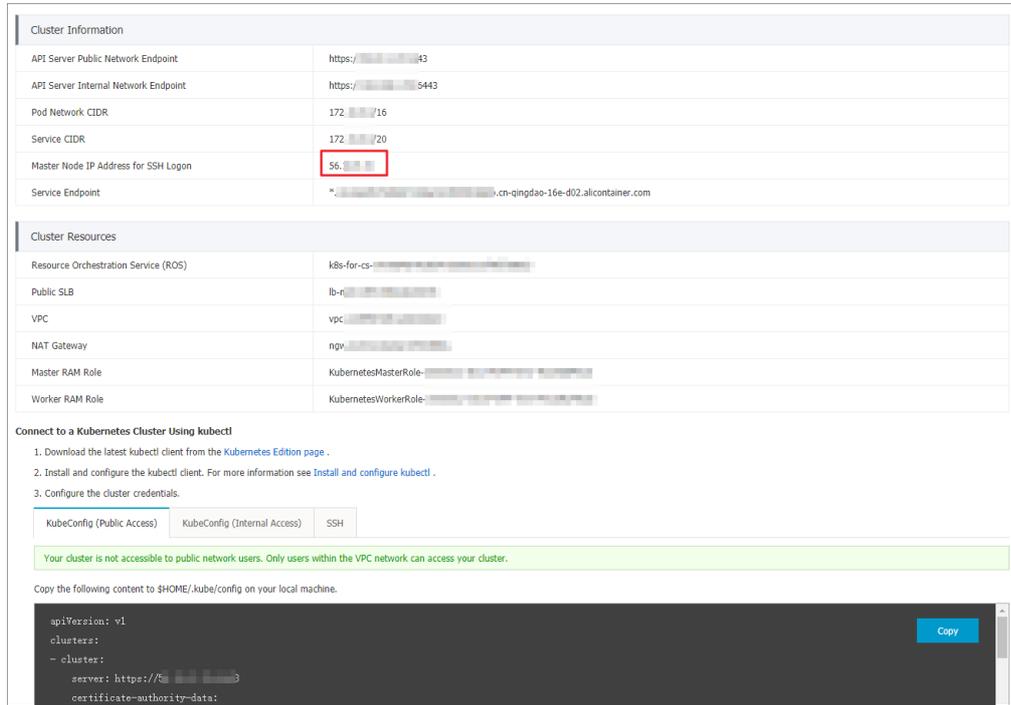
```
mkdir $HOME/.kube
scp root@<master-public-ip>:/etc/kubernetes/kube.conf $HOME/.kube/config
```

You can find `master-public-ip` on the cluster details page.

- i. [Log on to the Container Service console](#).
- ii. In the left-side navigation pane, click **Clusters**. The Clusters page appears.

iii. Find the target cluster and click **Manage** in the Actions column.

In the **Cluster Information** section, you can find the master node IP address.



3.6.2.5. Connect to a master node by using SSH

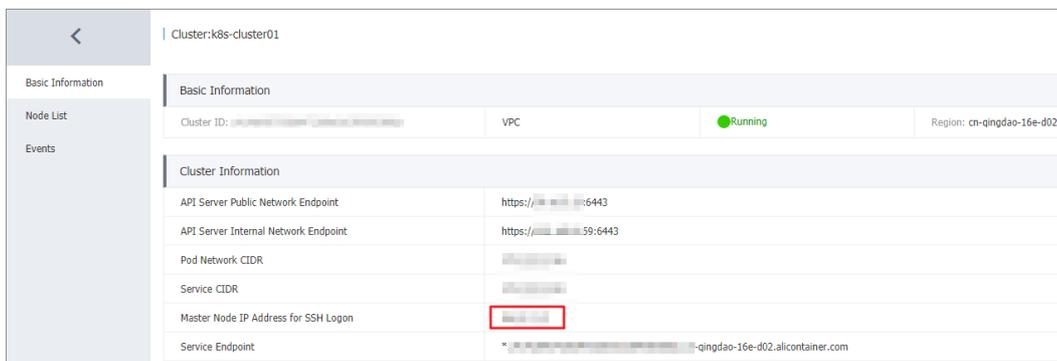
You can access a master node in a cluster by using a Secure Shell (SSH) client.

Prerequisites

- A Kubernetes cluster is created and **Use SSH to Access the Cluster from the Internet** is selected for the cluster. For more information, see [Create a Kubernetes cluster](#).
- The SSH client can connect to the network where the cluster is deployed.

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, choose **Clusters > Clusters** to go to the Clusters page. Find the cluster that you want to manage, and click **Manage** in the Actions column for the cluster.
3. The Basic Information page appears. In the Cluster Information section, you can find the IP address that is displayed in the **Master Node IP Address for SSH Logon** field.



4. Use SSH to connect to the cluster from an SSH client that has access to the cluster network.
 - o If you have a leased line that connects to the cluster network over the Internet, you can use tools such as PuTTY to create an SSH connection.
 - o If you have an Elastic Compute Service (ECS) instance that is connected to the Virtual Private Cloud (VPC) network of the cluster, run the following command to create an SSH connection:

```
ssh root@ssh_ip #ssh_ip specifies the IP address of the master node for SSH connection.
```

3.6.2.6. Expand a cluster

This topic describes how to scale out the worker nodes of a Kubernetes cluster in the Container Service console.

Context

You cannot scale out the master nodes of a Kubernetes cluster.

Procedure

1. [Log on to the Container Service console.](#)
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to expand and choose **More > Expand** in the **Actions** column.
4. On the **Node Pools** page, select the node pool that you want to scale out and click **Scale Out** in the **Actions** column.
5. Go to the Expand page and set the required parameters.

In this example, the number of worker nodes in the cluster is increased from three to five. The following table describes the required parameters.

| Parameter | Description |
|-------------------|--|
| Nodes to Add | Specify the number of nodes to be added to the cluster. |
| Region | By default, the region where the cluster is deployed is displayed. |
| Container Runtime | By default, the container runtime of the cluster is displayed. |
| VPC | By default, the virtual private cloud (VPC) of the cluster is displayed. |
| VSwitch | Select one or more vSwitches for the cluster. You can select at most three vSwitches that are deployed in different zones . |
| Instance Type | You can select one or more instance types. For more information, see the <i>Instance types</i> topic of <i>ECS User Guide</i> . |
| Selected Types | The selected instance types. |
| System Disk | Standard SSDs, enhanced SSDs (ESSDs), and ultra disks are supported. |
| Mount Data Disk | Standard SSDs, ESSDs, and ultra disks are supported. |
| Operating System | The operating system of the cluster. |
| Password | <ul style="list-style-type: none"> o Password: Enter the password that is used to log on to the nodes. o Confirm Password: Enter the password again. |

| Parameter | Description |
|---------------|--|
| ECS Label | You can add labels to the ECS instances. |
| Node Label | Add labels to nodes. |
| Taints | Add taints to the worker nodes in the cluster. |
| Custom Image | You can select a custom image. After you select a custom image, all nodes in the cluster are deployed by using this image. |
| RDS Whitelist | Set the Relational Database Service (RDS) whitelist. Add the IP addresses of the nodes in the cluster to the RDS whitelist. |
| User Data | Customize the startup behaviors of ECS instances and import data to the ECS instances. The user data can be used in the following ways: <ul style="list-style-type: none"> Run scripts during instance startup. Import user data as normal data to an ECS instance for future reference. |

6. Click **Submit**.

What's next

After the cluster is expanded, choose **Nodes > Nodes** in the left-side navigation pane. On the Nodes page, you can find that the number of worker nodes is increased from 3 to 5.

3.6.2.7. Renew a certificate

This topic describes how to renew a Kubernetes cluster certificate in the console.

Prerequisites

A Kubernetes cluster is created and the cluster certificate is about to expire.

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. Select the cluster for which you want to renew the certificate and click **Update Certificate**. The **Update Certificate** message appears.

 **Note** The **Update Certificate** button will be displayed two months before your cluster certificate expires.

4. Click **Update** and the **Confirm** page appears.
5. Click **OK**.

Result

- On the **Update Certificate** page, the following message appears: **The certificate has been updated**.
- On the **Clusters** page, the **Update Certificate** button disappears.

3.6.2.8. Delete a Kubernetes cluster

This topic describes how to delete a Kubernetes cluster in the Container Service console.

Procedure

1. [Log on to the Container Service console.](#)
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to delete and choose **More > Delete** in the **Actions** column.
4. In the **Delete Cluster** dialog box that appears, select the resources that you want to retain, select **I understand the above information and want to delete the specified cluster**, and then click **OK**.

What's next

Resource Orchestration Service (ROS) does not have permissions to delete resources that are manually added to resource created by ROS. For example, if you manually add a vSwitch to a virtual private cloud (VPC) created by ROS, ROS cannot delete the VPC and therefore the cluster cannot be deleted.

Container Service allows you to forcibly delete clusters. If your first attempt to delete a cluster fails, you can forcibly delete the cluster and ROS resource stack. However, you still need to manually release the resources that are manually added.

An error message appears when an attempt to delete a cluster fails.

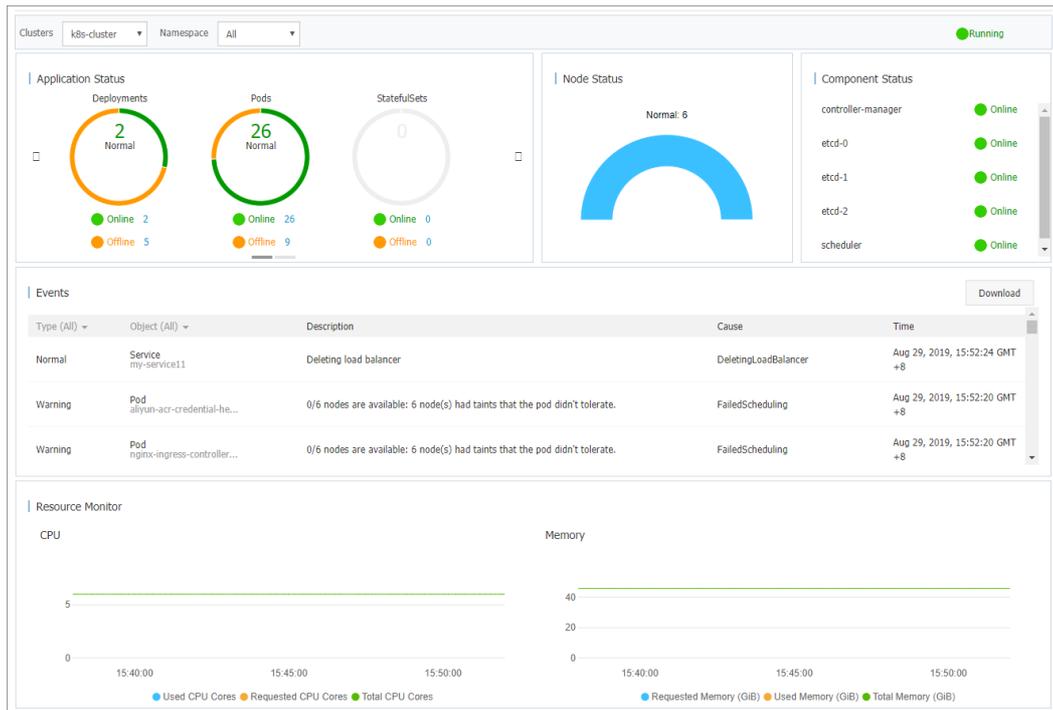
Select the cluster that you failed to delete and choose **More > Delete** in the **Actions** column. In the dialog box that appears, you can view the resources that are manually added. Select the **Force Delete** check box and click **OK** to delete the cluster and ROS resource stack.

3.6.2.9. View cluster overview

The Container Service console provides a cluster overview page. This page displays the information such as application status, component status, and resource monitoring status. This allows you to check the health status of your cluster at your convenience.

Procedure

1. [Log on to the Container Service console.](#)
2. In the left-side navigation pane, click **Overview**. The Overview page appears.
3. Select the target cluster and namespace. You can view the application status, component status, and resource monitoring charts.
 - **Application Status**: displays the statuses of the deployments, pods, and replica sets that are running in the cluster. Green sections indicate a normal state and yellow sections indicate an exception state.
 - **Node Status**: displays the statuses of the nodes in the cluster.
 - **Component Status**: Components are deployed in the kube-system namespace. Core components are used, such as the scheduler, controller-manager, and etcd.
 - **Events**: displays events such as warnings and errors. If no events are displayed, the cluster is running in the normal state.
 - **Monitoring**: displays CPU and memory monitoring charts. CPU usage is measured in cores or millicores and accurate to three decimal places. A millicore is one thousandth of a core. Memory usage is measured in GiB and accurate to three decimal places. For more information, see [Meaning of CPU](#) and [Meaning of memory](#).



3.6.3. Nodes

3.6.3.1. Add existing nodes to a Kubernetes cluster

Container Service allows you to add an existing Elastic Compute Service (ECS) instance to a Kubernetes cluster. You can add only worker nodes to clusters.

Prerequisites

- A Kubernetes cluster is created. For more information, see [Log on to the Container Service console](#).
- An ECS instance is created. Make sure that the region, zone, organization, project, security group, virtual private cloud (VPC), and operating system settings of the ECS instance are the same as those of the cluster.

Context

- By default, a cluster can contain at most 50 nodes. To increase the quota, submit a ticket.
- The ECS instance that you add must be in the same region and VPC as the cluster.
- The ECS instance must belong to the same Apsara Stack tenant account as the cluster.
- The ECS instance must be running the CentOS operating system.

Procedure

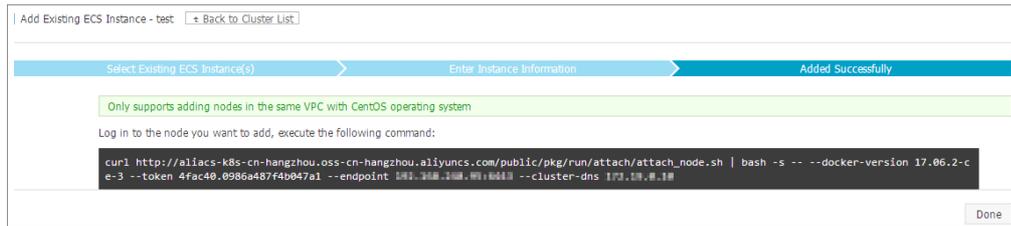
1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Nodes** > **Nodes**.
5. In the upper-right corner of the page, click **Add Existing Node**.
6. On the page that appears, you can manually add existing ECS instances to the cluster.

To manually add an ECS instance, you must obtain the installation command and log on to the ECS instance to run the command. You can add only one ECS instance at a time.

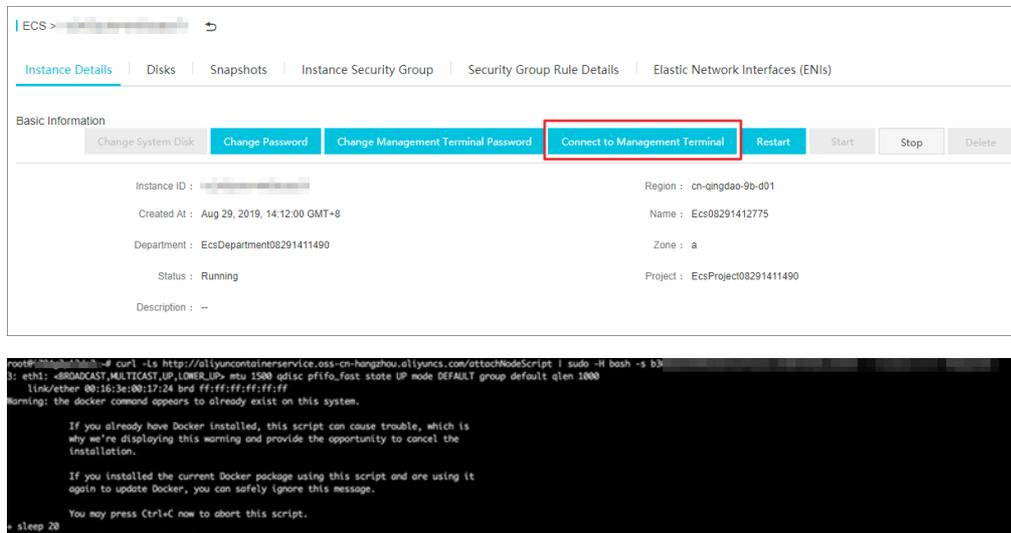
- i. Select **Manual**, find and select the ECS instance that you want to add, and then click **Next Step**. You can add only one ECS instance at a time.
- ii. Confirm the instance information and click **Next Step**.



- iii. On the Complete wizard page, copy the command.



- iv. Click **Done**.
- v. Go to the Apsara Uni-manager Management Console. In the top navigation bar, choose **Products > Elastic Compute Service**. On the **Instances** page, select the organization and region of the cluster, and then find the ECS instance that you want to add to the cluster.
- vi. Click the instance name to go to the Instance Details tab. Click **Connect to VNC**. In the dialog box that appears, enter the VNC password and then click **OK**. After you log on to the instance, paste the copied command and click **OK** to run the script.



- vii. After the script is executed, the ECS instance is added to the cluster. You can go to the Clusters page and click the cluster ID to view nodes in the cluster. Check whether the ECS instance has been added to the cluster.

3.6.3.2. Add nodes to an edge Kubernetes cluster

You can add worker nodes to an edge Kubernetes cluster in the Container Service console. However, you must make sure that the added nodes can communicate with the Kubernetes API server of the cluster. This topic describes how to add nodes to an edge Kubernetes cluster.

Prerequisites

Create an edge Kubernetes cluster

Context

By default, a cluster can contain at most 50 nodes.

Procedure

1. [Log on to the Container Service console](#)
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, click the name of the cluster that you want to manage.
4. In the left-side navigation pane of the cluster details page, choose **Nodes > Nodes**. On the page that appears, click **Add Existing Node** in the upper-right corner.
5. On the **Select Existing ECS Instance** wizard page, click **Next Step**.
You can only manually add nodes to edge Kubernetes clusters.
6. On the **Specify Instance Information** wizard page, set the parameters and click **Next Step**.

| Parameter | Description | Default |
|----------------------|---|--|
| flannelface | The name of the network interface controller (NIC) that is used by the Flannel plug-in. | The name of the NIC that is specified in the default route entry of the node. |
| enableiptables | Specifies whether to enable iptables. | false |
| quiet | Specifies whether to answer all questions with yes when you add nodes. | false |
| manageRuntime | Specifies whether to use edgeadm to install and manage the runtime. | false |
| nodeNameOverride | The name of the node. | <ul style="list-style-type: none"> ○ "". This is the default value. This value specifies that the hostname is used as the node name. ○ "**". This value specifies that a random string that contains six characters is used as the node name. ○ "*.X.XX". This value specifies that a random string that is followed by a suffix is used as the node name. The random string contains six characters. |
| allowedClusterAddons | The list of add-ons to be installed. By default, this parameter is empty. This indicates that no add-on is installed. For a standard edge node, set this parameter to ["kube-proxy","flannel","coredns"]. | [] |

| Parameter | Description | Default |
|--------------------|--|--|
| gpuVersion | Specifies whether the node to be added is a GPU-accelerated node. By default, this parameter is empty. Supported GPU models are Nvidia_Tesla_T4, Nvidia_Tesla_P4, and Nvidia_Tesla_P100. | "". This is the default value. This value specifies that the node to be added is not a GPU-accelerated node. |
| inDedicatedNetwork | Specifies whether an Express Connect circuit is used to connect to the managed edge Kubernetes cluster. | false |
| labels | Specifies the labels to be added to the node. | {} |
| annotations | Specifies the annotations to be added to the node. | {} |
| nodeface | <p>This parameter specifies the following information:</p> <ul style="list-style-type: none"> Specifies the node IP address that kubelet retrieves from the specified network interface. If you do not specify this parameter, kubelet attempts to retrieve the node IP address in the following order: <ul style="list-style-type: none"> Searches <i>/etc/hosts</i> for the node whose name is the same as the specified hostname. Finds the IP address of the network interface that is specified in the default route entry of the node. Specifies the name of the NIC that is used by the Flannel plugin. In this case, this parameter is equivalent to the flannelface parameter. This parameter will soon replace the flannelface parameter. | "" |

7. On the **Complete** wizard page, copy the script to the node that you want to add to the edge Kubernetes cluster and click **Done**.



8. Log on to the edge node and execute the script. This way, the node is added to the edge Kubernetes cluster.

3.6.3.3. View nodes

You can view nodes of a Kubernetes cluster by using the Container Service console, kubectl, or Kubernetes Dashboard.

View nodes by using kubectl

 **Note** To view the nodes in a cluster by using kubectl, you must [Connect to a Kubernetes cluster through kubectl](#).

Connect to a cluster by using kubectl and run the following command to view the nodes in the cluster:

```
kubectl get nodes
```

Sample output:

```
$ kubectl get nodes
NAME                STATUS AGE  VERSION
iz2ze2n6ep53tch701yh9zz Ready 19m  v1.6.1-2+ed9e3d33a07093
iz2zeaf762wibijx39e5az Ready 7m   v1.6.1-2+ed9e3d33a07093
iz2zeaf762wibijx39e5bz Ready 7m   v1.6.1-2+ed9e3d33a07093
iz2zef4dnn9nos8elyr32kz Ready 14m  v1.6.1-2+ed9e3d33a07093
iz2zeitvvo8enoreufstkmz Ready 11m  v1.6.1-2+ed9e3d33a07093
```

View nodes by using the Container Service console

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Nodes > Nodes** to view the nodes in the cluster.

3.6.3.4. Manage node labels

You can manage node labels in the Container Service console. You can add a label to multiple nodes at a time, filter nodes by label, and delete labels.

You can use labels to schedule nodes. For more information, see [Set node scheduling](#).

Prerequisites

A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).

Add a label to multiple nodes at a time

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Nodes > Nodes**.
5. On the **Nodes** page, click **Manage Labels and Taints** in the upper-right corner.
6. Select multiple nodes and click **Add Label**.
7. In the dialog box that appears, enter the name and value of the label and click **OK**.



The image shows a dialog box titled "Add" with a close button (X) in the top right corner. It contains two input fields: "Name" with the text "group" and "Value" with the text "worker". At the bottom right, there are two buttons: "OK" (highlighted in blue) and "Close".

On the Labels tab, you can find that the selected nodes have the same label.

Delete a label

1. Log on to the Container Service console.
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Nodes > Nodes**.
5. On the **Nodes** page, click **Manage Labels and Taints** in the upper-right corner.
6. Select a node, find the label that you want to delete, and then click the  icon. In the message that appears, click **Confirm**.

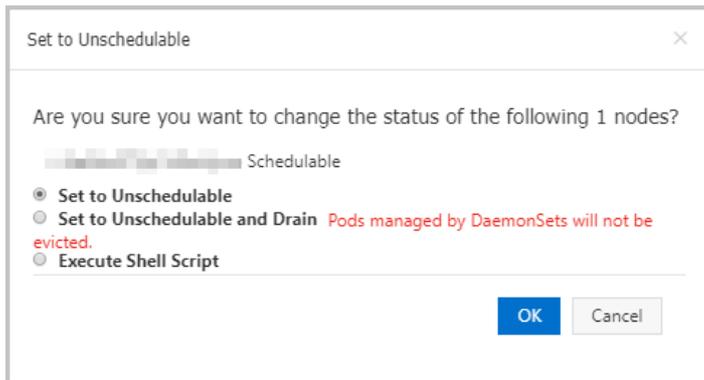
After the label is deleted, it is removed from the Labels column.

3.6.3.5. Set node schedulability

You can mark a node as schedulable or unschedulable in the Container Service console. This allows you to optimize the distribution of the loads on each node. This topic describes how to set node schedulability.

Procedure

1. Log on to the Container Service console.
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Nodes > Nodes**.
5. On the **Nodes** page, select the node that you want to manage and click **Drain/Set to Unscheduleable** in the **Actions** column.
6. In the dialog box that appears, you can change the status of the node.
 - If you select **Set to Unscheduleable**, pods will not be scheduled to this node when you deploy new applications.
 - If you select **Set to Unscheduleable and Drain**, pods will not be scheduled to this node when you deploy new applications. Pods on this node will be evicted, except for the pods that are managed by DaemonSets.In this example, **Set to Unscheduleable** is selected.



7. Click **OK**.

The status of the node is changed to Unscheduleable.

What's next

Pods will not be scheduled to the node when you deploy new applications.

3.6.3.6. Remove a node

To restart or release an Elastic Compute Service (ECS) node in a cluster, you must first remove the node from the cluster. This topic describes how to remove a node.

Prerequisites

- A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).
- You can connect to the Kubernetes cluster by using `kubectl`. For more information, see [Connect to a Kubernetes cluster through `kubectl`](#).

Context

- When you remove a node, pods that run on the node are migrated to other nodes. This may cause service interruption. We recommend that you remove nodes during off-peak hours.
- Unknown errors may occur when you remove nodes. Before you remove nodes, we recommend that you back up data on these nodes.
- Nodes remain in the unscheduleable state when they are being removed.
- You can remove only worker nodes. You cannot remove master nodes.

Procedure

1. Run the following command to migrate the pods on the node that you want to remove to other nodes.

Note Make sure that the other nodes have sufficient resources for these pods.

```
kubectl drain node-name
```

Note `node-name` must be in the format of `your-region-name.node-id`.

- `your-region-name` specifies the region where the cluster that you want to manage is deployed.
- `node-id` specifies the ID of the ECS instance where the node to be removed is deployed. Example: `cn-hangzhou.i-xxx`.

2. [Log on to the Container Service console](#).
3. In the left-side navigation pane, click **Clusters**.

4. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
5. In the left-side navigation pane of the details page, choose **Nodes > Node**.
6. Find the node that you want to remove and choose **More > Remove** in the **Actions** column.

 **Note** To remove multiple nodes at a time, select the nodes that you want to remove on the **Nodes** page and click **Batch Remove**.

7. (Optional) Set parameters in the Remove Node dialog box.
 - o Select **Drain the Node** to migrate the pods on the node to other nodes. If you select this option, make sure that the other nodes have sufficient resources for these pods.
 - o Select **Release ECS Instance** to release the ECS instance where the node is deployed.

 **Note**

- o Select this option to release only pay-as-you-go ECS instances.
- o Subscription ECS instances are automatically released after the subscription expires.
- o If you do not select **Release ECS Instance**, you are still billed for the ECS instance where the node is deployed.

8. In the **Remove Node** message, click **OK**.

3.6.3.7. View node resource usage

You can view the resource usage of the nodes in a cluster in the Container Service console.

Prerequisites

A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Nodes > Nodes**.
5. On the **Nodes** page, find the node that you want to manage and choose **More > Details** in the **Actions** column.

You can view the request rate and usage rate of CPU and memory resources on each node.

- o CPU request rate = $\text{sum}(\text{The amount of CPU resources requested by all pods on the node}) / \text{Total CPU resources of the node}$
- o CPU usage rate = $\text{sum}(\text{The amount of CPU resources used by all pods on the node}) / \text{Total CPU resources of the node}$
- o Memory request rate = $(\text{The amount of memory resources requested by all pods on the node}) / \text{Total memory resources of the node}$
- o Memory usage rate = $\text{sum}(\text{The amount of memory resources used by all pods on the node}) / \text{Total memory resources of the node}$

Note

- You can adjust the workload of a node based on the resource usage. For more information, see [Set node scheduling](#).
- When the request or usage rate of a node reaches 100%, pods are not scheduled to the node.

3.6.3.8. Upgrade the NVIDIA driver on a GPU node

This topic describes how to upgrade the NVIDIA driver on a GPU node when workloads are deployed on the node and when no workload is deployed on the node.

Upgrade the NVIDIA driver on a GPU node where workloads are deployed

1. [Connect to a Kubernetes cluster through kubectl](#)
2. Run the following command to set the target node to unschedulable.

```
kubectl cordon node-name
```

Note

- Currently, you can only upgrade the NVIDIA driver on worker nodes.
- *node-name* must be in the format of *your-region-name.node-id*.
 - *your-region-name* represents the region where your cluster is deployed.
 - *node-id* represents the ID of the ECS instance where the target node is deployed.

You can run the following command to query *node-name*.

```
kubectl get node
```

```
[root@gpu-test ~]# kubectl cordon cn-hangzhou.i-  
node/cn-hangzhou.i- already cordoned
```

3. Run the following command to migrate pods from the target node to other nodes:

```
kubectl drain node-name --grace-period=120 --ignore-daemonsets=true
```

```
[root@gpu-test ~]# kubectl drain cn-hangzhou.i-  
node/cn-hangzhou.i- --grace-period=120 --ignore-daemonsets=true  
node/cn-hangzhou.i- cordoned  
WARNING: Ignoring DaemonSet-managed pods: flexvolume-  
pod/domain-nginx- evicted  
pod/old-nginx- evicted  
pod/new-nginx- evicted  
pod/old-nginx- evicted
```

4. Run the following command to log on to the target node:

```
ssh root@xxx.xxx.x.xx
```

5. Run the following command to check the current NVIDIA driver version:

```
nvidia-smi
```



```
cd /tmp
```

```
curl -O https://cn.download.nvidia.cn/tesla/384.111/NVIDIA-Linux-x86_64-384.111.run
```

```
chmod u+x NVIDIA-Linux-x86_64-384.111.run
```

```
./NVIDIA-Linux-x86_64-384.111.run --uninstall -a -s -q
```

4. Run the following command to restart the target node.

```
reboot
```

5. Download the driver that you want to use from the official NVIDIA website. In this example, version 410.79 is used.

6. Run the following command to install the downloaded driver under the directory where it was saved:

```
sh ./NVIDIA-Linux-x86_64-410.79.run -a -s -q
```

7. Run the following commands to configure the driver:

```
nvidia-smi -pm 1 || true
```

```
nvidia-smi -acp 0 || true
```

Result

Run the following command on a master node to check the NVIDIA driver version on the target node. The driver version is now 410.79.

Note Replace *node-name* with the target node name.

```
kubectl exec -n kube-system -t nvidia-device-plugin-node-name nvidia-smi
```

```

[root@gpu-test ~]# kubectl exec -n kube-system -t nvidia-device-plugin-cn-... nvidia-smi
Mon Jan 21 03:14:48 2019
+-----+
| NVIDIA-SMI 410.79      | Driver Version: 410.79      | CUDA Version: N/A      |
+-----+-----+
| GPU  Name          Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp   Perf   Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+
|  0   Tesla P4             0n      | 00000000:00:08.0 Off  |    0%      Default  |
|N/A   21C    P8             6W / 75W |  0MiB / 7611MiB |           |
+-----+-----+
+-----+
| Processes:                        | GPU Memory Usage |
| GPU       PID  Type  Process name                        | Usage             |
+-----+-----+
| No running processes found      |                   |
+-----+

```

3.6.3.9. GPU scheduling for Kubernetes clusters with GPU-accelerated nodes

This topic describes GPU scheduling for Kubernetes clusters with GPU-accelerated nodes.

Prerequisites

Container Service, Resource Orchestration Service (ROS), and Resource Access Management (RAM) are activated.

 **Note** Container Service uses ROS to deploy applications in Kubernetes clusters. To create a Kubernetes cluster, you must first activate ROS.

Context

Starting from version 1.8, Kubernetes adds support for the following hardware acceleration devices by using **device plug-ins**: NVIDIA GPUs, InfiniBand devices, and field-programmable gate arrays (FPGAs). GPU solutions developed by the community will be phased out in version 1.10, and removed from the master code in version 1.11. Container Service enables you to use a Kubernetes cluster with GPU-accelerated nodes to run compute-intensive tasks such as machine learning and image processing. You can deploy applications and achieve auto scaling without the need to install NVIDIA drivers or Compute Unified Device Architecture (CUDA) in advance.

The system performs the following operations when a Kubernetes cluster is created:

- Creates Elastic Compute Service (ECS) instances, configures a public key to enable Secure Shell (SSH) logon from master nodes to other nodes, and then configures the Kubernetes cluster through CloudInit.
- Creates a security group that allows access to the virtual private cloud (VPC) over Internet Control Message Protocol (ICMP).
- If you do not specify an existing VPC, the system creates a VPC and a vSwitch and creates SNAT entries for the vSwitch.
- Adds route entries to the VPC.
- Creates a NAT gateway and an elastic IP address (EIP).
- Creates a RAM user and grants it permissions to query, create, and delete ECS instances, and permissions to add and delete disks. The RAM user is also granted full permissions on Server Load Balancer (SLB) instances, CloudMonitor, VPC, Log Service, and Apsara File Storage NAS (NAS). The system also creates an AccessKey pair for the RAM user. The system automatically creates SLB instances, disks, and VPC route entries based on your configuration.
- Creates an internal-facing SLB instance and opens port 6443.
- Creates an Internet-facing SLB instance and open ports 6443, 8443, and 22. If you enable SSH logon when you create the cluster, port 22 is open. Otherwise, port 22 is not open.

Limits

- Kubernetes clusters support only VPCs.
- By default, you can create only a limited amount of cloud resources with each account. You cannot create clusters if the quota on clusters is reached. Make sure that you have sufficient quota before you create a cluster. To request a quota increase, submit a ticket.
 - By default, you can create at most five clusters across all regions with each account. Each cluster can contain at most 40 nodes. To increase the quota of clusters or nodes, submit a ticket.

 **Note** In a Kubernetes cluster, you can create at most 48 route entries for each VPC. This means that a cluster can contain at most 48 nodes. To increase the quota of nodes, submit a ticket to increase the quota of route entries first.

- By default, you can create at most 100 security groups with each account.
- By default, you can create at most 60 pay-as-you-go SLB instances with each account.
- By default, you can create at most 20 EIPs with each account.
- Limits on ECS instances:
 - Only CentOS is supported.

Create a GN5 Kubernetes cluster

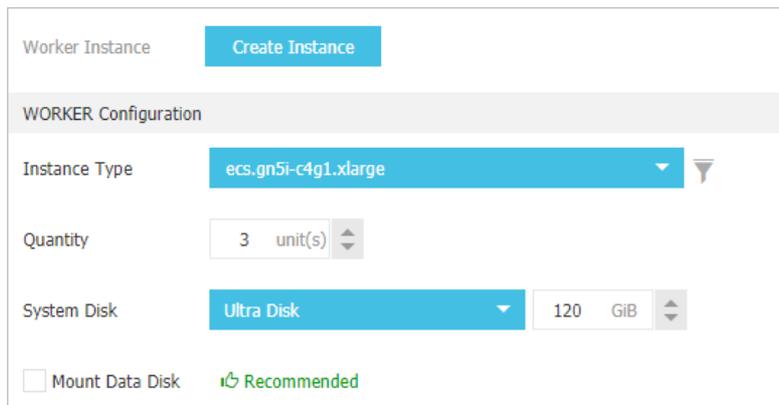
GN5 Kubernetes clusters support only Kubernetes 1.12.6-aliyun.1. Kubernetes 1.11.5 is not supported.

1. Log on to the Container Service console.
2. In the left-side navigation pane, click **Clusters**.
3. In the upper-right corner of the **Clusters** page, click **Create Kubernetes Cluster**.

On the cluster configuration page, set the parameters.

Note To create a cluster with GPU-accelerated nodes, select GPU-accelerated ECS instance types as worker nodes. For more information about other parameters, see [Cluster parameters](#).

4. Configure worker nodes. In this example, worker nodes are used to run GPU computing tasks and the gn5i-c4g1 instance type is selected.
 - i. If you choose to create worker instances, you must set Instance Type and Quantity. In this example, three GPU-accelerated worker nodes are created.



Note We recommend that you use standard SSDs.

- ii. If you choose to add existing instances, you must create GPU-accelerated ECS instances in the region where you want to create the cluster in advance.
5. Specify the other parameters and click **Create Cluster** to start the deployment.
After the cluster is created, choose **Clusters > Nodes** to go to the Nodes page.
Select a worker nodes that was configured for the cluster and choose **More > Details** in the Actions column to view the GPU devices that are attached to the node.

Create a GPU experimental environment to run TensorFlow

Jupyter is a standard tool that is used by data scientists to create an experimental environment to run TensorFlow. The following example shows how to deploy a Jupyter application.

1. Log on to the Container Service console. In the left-side navigation pane, choose **Applications > Deployments** to go to the **Deployments** page.
2. In the upper-right corner of the page, click **Create from Template**.
3. Select the required cluster and namespace. Select a sample template, or set Sample Template to Custom and customize the template in the Template field. Then, click **Create** to create the application.

Deploy templates

Only Kubernetes versions 1.8.4 and above are supported. For clusters of version 1.8.1, you can perform "upgrade cluster" operation in the cluster list

Clusters: xuntest2

Namespace: default

Resource Type: Custom

Template

```
1 ---
2 # Define the tensorflow deployment
3 apiVersion: apps/v1
4 kind: Deployment
5 metadata:
6   name: tf-notebook
7   labels:
8     app: tf-notebook
9 spec:
10  replicas: 1
11  selector: # define how the deployment finds the pods it mangages
12    matchLabels:
13      app: tf-notebook
14  template: # define the pods specifications
15    metadata:
16      labels:
17        app: tf-notebook
18    spec:
19      containers:
20        - name: tf-notebook
21          image: tensorflow/tensorflow:1.4.1-gpu-py3
22          resources:
23            limits:
24              nvidia.com/gpu: 1
25          ports:
26            - containerPort: 8888
27              hostPort: 8888
28          env:
29            - name: PASSWORD
```

Add Deployment

Deploy with exist template

Save Template DEPLOY

In this example, the template uses a Deployment and a Service to create a Jupyter application.

```

---
# Define the tensorflow deployment
apiVersion: apps/v1
kind: Deployment
metadata:
  name: tf-notebook
  labels:
    app: tf-notebook
spec:
  replicas: 1
  selector: # define how the deployment finds the pods it manages
    matchLabels:
      app: tf-notebook
  template: # define the pods specifications
    metadata:
      labels:
        app: tf-notebook
    spec:
      containers:
      - name: tf-notebook
        image: tensorflow/tensorflow:1.4.1-gpu-py3
        resources:
          limits:
            nvidia.com/gpu: 1          #The number of NVIDIA GPUs that are requested by the application.
        ports:
          - containerPort: 8888
            hostPort: 8888
        env:
          - name: PASSWORD          #The password that is used to access the Jupyter application. You can modify the password as required.
            value: mypassw0rd
# Define the tensorflow service
---
apiVersion: v1
kind: Service
metadata:
  name: tf-notebook
spec:
  ports:
  - port: 80
    targetPort: 8888
  name: jupyter
  selector:
    app: tf-notebook
  type: LoadBalancer          #An SLB instance is used to route internal traffic and perform load balancing.

```

4. In the left-side navigation pane, choose **Ingresses and Load Balancing > Services**. Select the required cluster and namespace, find the tf-notebook Service, and then check its external endpoint.

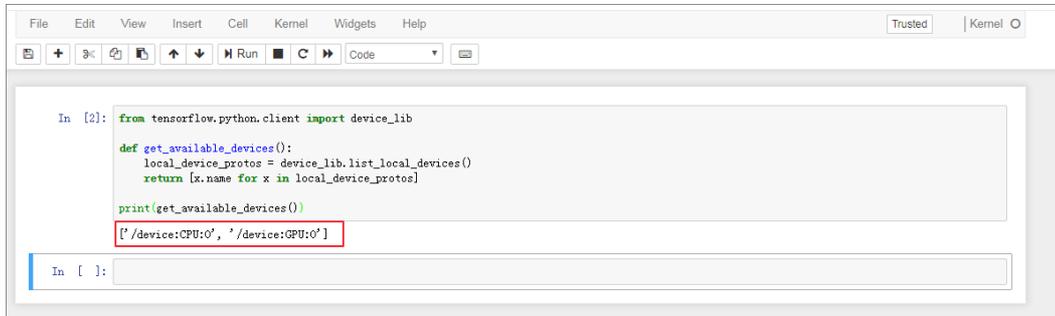
| Name | Label | Type | Time Created | ClustersIP | InternalEndpoint | ExternalEndpoint | Action |
|-------------|--|--------------|---------------------|------------|---|------------------|---|
| kubernetes | component:apiserver provider:kubernetes | ClusterIP | 05/17/2019,18:12:33 | [Redacted] | kubernetes:443 TCP | - | Details Update View YAML Delete |
| tf-notebook | - | LoadBalancer | 05/23/2019,10:46:02 | [Redacted] | tf-notebook:80 TCP tf-notebook:30708 TCP | [Redacted] | Details Update View YAML Delete |

5. To connect to the Jupyter application, enter `http://EXTERNAL-IP` into the address bar of your browser and enter the password specified in the template.
6. You can run the following program to verify that the Jupyter application is allowed to use GPU resources. The program lists all devices that can be used by TensorFlow:

```

from tensorflow.python.client import device_lib
def get_available_devices():
    local_device_protos = device_lib.list_local_devices()
    return [x.name for x in local_device_protos]
print(get_available_devices())

```



```

In [2]: from tensorflow.python.client import device_lib
def get_available_devices():
    local_device_protos = device_lib.list_local_devices()
    return [x.name for x in local_device_protos]
print(get_available_devices())

```

['/device:CPU:0', '/device:GPU:0']

3.6.3.10. Use labels to schedule pods to GPU-accelerated nodes

To use Kubernetes clusters for GPU computing, you must schedule pods to GPU-accelerated nodes. Container Service allows you to schedule pods to specific GPU-accelerated nodes by adding labels to the GPU-accelerated nodes.

Context

When Kubernetes deploys nodes with NVIDIA GPUs, the attributes of these GPUs are discovered and exposed as node labels. These labels have the following benefits:

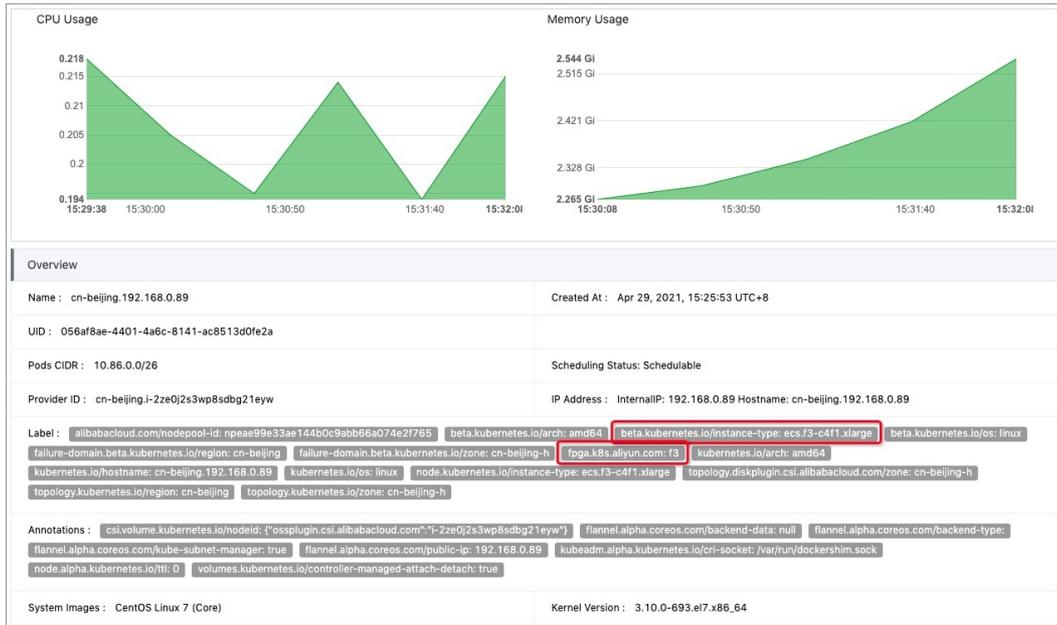
- You can use the labels to filter GPU-accelerated nodes.
- The labels can be used as conditions to schedule pods.

Procedure

1. [Log on to the Container Service console.](#)
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Nodes** > **Nodes**.

Note In this example, the Kubernetes cluster contains three master nodes, among which two are equipped with GPUs. Record the node IP addresses.

5. On the **Nodes** page, select the node that is equipped with GPUs and choose **More** > **Details** in the **Actions** column to view the labels that are added to the node.



You can also log on to a master node and run the following command to view the labels of GPU-accelerated nodes:

```
# kubectl get nodes
NAME                                STATUS ROLES AGE  VERSION
cn-beijing.i-2ze2dy2h9w97v65u**** Ready  master 2d   v1.12.6-aliyun.1
cn-beijing.i-2ze8o1a45qdv5q8a**** Ready  <none> 2d   v1.12.6-aliyun.1 # Compare these nodes with the
nodes displayed in the console to identify GPU-accelerated nodes.
cn-beijing.i-2ze8o1a45qdv5q8a**** Ready  <none> 2d   v1.12.6-aliyun.1
cn-beijing.i-2ze9xylyn11vop7g**** Ready  master 2d   v1.12.6-aliyun.1
cn-beijing.i-2zed5sw8snjniq6m**** Ready  master 2d   v1.12.6-aliyun.1
cn-beijing.i-2zej9s0zjykp9pw**** Ready  <none> 2d   v1.12.6-aliyun.1
```

Select a GPU-accelerated node and run the following command to query the labels of the node:

```
# kubectl describe node cn-beijing.i-2ze8o1a45qdv5q8a****
Name:      cn-beijing.i-2ze8o1a45qdv5q8a7luz
Roles:     <none>
Labels:    aliyun.accelerator/nvidia_count=1 #Note
           aliyun.accelerator/nvidia_mem=12209MiB
           aliyun.accelerator/nvidia_name=Tesla-M40
           beta.kubernetes.io/arch=amd64
           beta.kubernetes.io/instance-type=ecs.gn4-c4g1.xlarge
           beta.kubernetes.io/os=linux
           failure-domain.beta.kubernetes.io/region=cn-beijing
           failure-domain.beta.kubernetes.io/zone=cn-beijing-a
           kubernetes.io/hostname=cn-beijing.i-2ze8o1a45qdv5q8a****
.....
```

In this example, the following labels are added to the GPU-accelerated node.

| key | value |
|---------------------------------|--|
| aliyun.accelerator/nvidia_count | The number of GPU cores. |
| aliyun.accelerator/nvidia_mem | The size of the GPU memory. Unit: MiB. |

| key | value |
|--------------------------------|-----------------------------|
| aliyun.accelerator/nvidia_name | The name of the NVIDIA GPU. |

GPU-accelerated nodes of the same type have the same GPU name. You can use this label to locate GPU-accelerated nodes.

```
# kubectl get no -l aliyun.accelerator/nvidia_name=Tesla-M40
NAME                STATUS  ROLES  AGE   VERSION
cn-beijing.i-2ze8o1a45qdv5q8a**** Ready  <none> 2d   v1.12.6-aliyun.1
cn-beijing.i-2ze8o1a45qdv5q8a**** Ready  <none> 2d   v1.12.6-aliyun.1
```

- In the left-side navigation pane of the details page, choose **Workloads > Deployments**.
- On the **Deployments** page, select the namespace and click **Create from Template** in the upper-right corner.
- Create a Deployment for a TensorFlow job. The Deployment is used to schedule pods to a GPU-accelerated node.
- You can also exclude an application from GPU-accelerated nodes. The following example shows how to schedule a pod based on node affinity for an NGINX application. For more information, see the section that describes node affinity in [Create an application from an image](#).

The following YAML template is used as an example:

```
apiVersion: v1
kind: Pod
metadata:
  name: not-in-gpu-node
spec:
  affinity:
    nodeAffinity:
      requiredDuringSchedulingIgnoredDuringExecution:
        nodeSelectorTerms:
          - matchExpressions:
              - key: aliyun.accelerator/nvidia_name
                operator: DoesNotExist
  containers:
    - name: not-in-gpu-node
      image: nginx
```

Result

In the left-side navigation pane of the details page, choose **Workloads > Pods**. On the Pods page, select the namespace to view pods created in the namespace. On the Pods page, you can find that the pods in the preceding examples are scheduled to the desired nodes. This means that labels can be used to schedule pods to GPU-accelerated nodes.

3.6.3.11. Manually upgrade the kernel of a GPU node in a cluster

This topic describes how to manually upgrade the kernel of a GPU node in a cluster.

Context

The current kernel version is earlier than `3.10.0-957.21.3`.

Procedure

1. [Connect to a Kubernetes cluster through kubectl](#).
2. Run the following command to set the target GPU node to unschedulable. This example uses node cn-beijing.i-2ze19qyi8votgjz12345 as the target node.

```
kubectl cordon cn-beijing.i-2ze19qyi8votgjz12345
node/cn-beijing.i-2ze19qyi8votgjz12345 already cordoned
```

3. Run the following command to drain the target GPU node:

```
# kubectl drain cn-beijing.i-2ze19qyi8votgjz12345 --grace-period=120 --ignore-daemonsets=true
node/cn-beijing.i-2ze19qyi8votgjz12345 cordoned
WARNING: Ignoring DaemonSet-managed pods: flexvolume-9scb4, kube-flannel-ds-r2qmh, kube-proxy-worker-l62sf
, logtail-ds-f9vbg
pod/nginx-ingress-controller-78d847fb96-5fkkw evicted
```

4. Uninstall the existing nvidia-driver.

 **Note** This step uninstalls the version 384.111 driver. If your driver version is not 384.111, you need to download a driver from the official NVIDIA website and replace 384.111 with your actual version number.

- i. Log on to the target GPU node and run the `nvidia-smi` command to query the driver version.

```
# nvidia-smi -a | grep 'Driver Version'
Driver Version          : 384.111
```

- ii. Run the following commands to download the driver installation package:

```
cd /tmp/
curl -O https://cn.download.nvidia.cn/tesla/384.111/NVIDIA-Linux-x86_64-384.111.run
```

 **Note** The installation package is required to uninstall the driver.

- iii. Run the following commands to uninstall the existing nvidia-driver:

```
chmod u+x NVIDIA-Linux-x86_64-384.111.run
./NVIDIA-Linux-x86_64-384.111.run --uninstall -a -s -q
```

5. Run the following commands to upgrade kernel:

```
yum clean all && yum makecache
yum update kernel -y
```

6. Run the following command to restart the GPU node:

```
reboot
```

7. Log on to the GPU node and run the following command to install the kernel-level package.

```
yum install -y kernel-devel-$(uname -r)
```

8. Run the following commands to download the required driver and install it on the target node. In this example, version 410.79 is used.

```
cd /tmp/
curl -O https://cn.download.nvidia.cn/tesla/410.79/NVIDIA-Linux-x86_64-410.79.run
chmod u+x NVIDIA-Linux-x86_64-410.79.run
sh ./NVIDIA-Linux-x86_64-410.79.run -a -s -q
# warm up GPU
nvidia-smi -pm 1 || true
nvidia-smi -acp 0 || true
nvidia-smi --auto-boost-default=0 || true
nvidia-smi --auto-boost-permission=0 || true
nvidia-modprobe -u -c=0 -m || true
```

9. Check the `/etc/rc.d/rc.local` file and check whether the following configurations are included. If not, add the following content.

```
nvidia-smi -pm 1 || true
nvidia-smi -acp 0 || true
nvidia-smi --auto-boost-default=0 || true
nvidia-smi --auto-boost-permission=0 || true
nvidia-modprobe -u -c=0 -m || true
```

10. Run the following commands to restart kubelet and Docker.

```
service kubelet stop
service docker restart
service kubelet start
```

11. Run the following command to set the GPU node to schedulable:

```
# kubectl uncordon cn-beijing.i-2ze19qyi8votgjz12345
node/cn-beijing.i-2ze19qyi8votgjz12345 already uncordoned
```

12. Run the following command on the `nvidia-device-plugin` container to check the driver version:

```
kubectl exec -n kube-system -t nvidia-device-plugin-cn-beijing.i-2ze19qyi8votgjz12345 nvidia-smi
Thu Jan 17 00:33:27 2019
+-----+
| NVIDIA-SMI 410.79   Driver Version: 410.79   CUDA Version: N/A   |
+-----+
| GPU Name   Persistence-M| Bus-Id  Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|  Memory-Usage | GPU-Util  Compute M. |
+-----+-----+
| 0 Tesla P100-PCIE... On  | 00000000:00:09:0 Off |          0          |
| N/A   27C   P0   28W / 250W | 0MiB / 16280MiB | 0%      Default  |
+-----+-----+
+-----+
| Processes:                                     GPU Memory |
|  GPU   PID  Type  Process name      Usage   |
+-----+-----+
| No running processes found                       |
+-----+
```

3.6.3.12. Node pools

3.6.3.12.1. Create a node pool

You can use a node pool to manage multiple nodes in a Kubernetes cluster as a group. For example, you can centrally manage the labels and taints of the nodes in a node pool. This topic describes how to create a node pool in the Container Service console.

Prerequisites

- A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).
- The Kubernetes cluster must run Kubernetes 1.9 or later.

Notice

- By default, a cluster can contain at most 100 nodes.
- Before you add an existing Elastic Compute Service (ECS) instance that is deployed in a virtual private cloud (VPC), make sure that an elastic IP address (EIP) is associated with the ECS instance, or a NAT gateway is created for the VPC. In addition, make sure that the ECS instance can access the Internet. Otherwise, you cannot add the ECS instance.

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Nodes > Node Pools**.
5. On the **Node Pools** page, click **Create Node Pool**.
6. In the **Create Node Pool** dialog box, set the parameters.

For more information about the parameters, see [Create a Kubernetes cluster](#). The following list describes some of the parameters:

- Name: the name of the node pool.
 - Quantity: the initial number of nodes in the node pool. If you do not want to add nodes to the node pool, set this parameter to 0.
 - ECS Label: the labels of the ECS instances.
 - Node Label: the labels of nodes in the node pool.
 - Custom Resource Group: the resource group to which the nodes of the node pool belong.
 - Custom Security Group: the custom security group of the ECS instance.
7. Click **Confirm Order**.
On the **Node Pools** page, check the **Status** column of the managed node pool. If the node pool is in the **Initializing** state, it indicates that the node pool is being created. After the node pool is created, the **state** of the node pool changes to **Active**.

What's next

After the node pool is created, find the node pool on the **Node Pools** page and click **Details** in the **Actions** column to view details of the node pool.

3.6.3.12.2. Scale out a node pool

You can use a node pool to manage multiple nodes in a Kubernetes cluster as a group. For example, you can centrally manage the labels and taints of the nodes in a node pool. This topic describes how to scale out a node pool in the Container Service console.

Prerequisites

- A node pool is created. For more information, see [Create a node pool](#).
- The Kubernetes cluster must run Kubernetes 1.9 or later.

 Notice

- By default, a cluster can contain at most 100 nodes.
- Before you add an existing Elastic Compute Service (ECS) instance that is deployed in a virtual private cloud (VPC), make sure that an elastic IP address (EIP) is associated with the ECS instance, or a NAT gateway is created for the VPC. In addition, make sure that the ECS instance can access the Internet. Otherwise, you cannot add the ECS instance.

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Nodes > Node Pools**.
5. On the **Node Pools** page, select the node pool that you want to manage and click **Scale Out** in the **Actions** column.
6. In the dialog box that appears, set the number of nodes that you want to add to the node pool.
7. (Optional) In the dialog box that appears, click **Modify Node Pool Settings** to modify the node pool.

For more information, see [Expand a Container Service cluster](#). The following list describes some of the parameters:

- **ECS Label:** You can add labels to the ECS instances in the node pool.
- **Node Label:** You can add labels to the nodes in the node pool.
- **Taints:** You can add one or more taints to the specified nodes.

 **Note** If you select **Synchronize Node Labels and Taints**, the specified labels and taints are synchronized to the existing and newly added nodes.

- **CloudMonitor Agent:** You can install the CloudMonitor agent on the nodes and view monitoring information about the nodes in the CloudMonitor console.
8. Click **Submit**.
On the **Node Pools** page, the **status** of the node pool is **Scaling**. This indicates that the scale-out event is in progress. After the scale-out event is completed, the **status** of the node pool changes to **Active**.

What's next

Click **Details** in the **Actions** column for the node pool. On the **Nodes** tab, you can check the nodes that are added to the node pool.

3.6.3.12.3. Schedule an application pod to a specific node pool

Label is an important concept of Kubernetes. Services, Deployments, and pods are associated with each other by labels. You can configure pod scheduling policies based on node labels. This allows you to schedule pods to nodes that have specific labels. This topic describes how to schedule an application pod to a specific node pool.

Procedure

1. Add a label to the nodes in a node pool.

Container Service allows you to manage a group of cluster nodes by using a node pool. For example, you can centrally manage the labels and taints of the nodes in a node pool. For more information about how to create a node pool, see [Create a node pool](#).

- i. [Log on to the Container Service console](#)
- ii. In the left-side navigation pane, click **Clusters**.

- iii. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
- iv. In the left-side navigation pane of the details page, choose **Nodes > Node Pools**.
- v. In the upper-right corner of the **Node Pools** page, click **Create Node Pool**.
- vi. In the Create Node Pool dialog box, click **Show Advanced Options** and click  on the right of **Node Label** to add labels to nodes.

In this example, the pod: nginx label is added.

You can also click **Scale Out** on the right side of a node pool to update or add labels for the nodes. If automatic scaling is enabled for a node pool, click **Modify** on the right side of the node pool to update or add labels for the nodes.

2. Configure a scheduling policy for an application pod.

After the preceding step is completed, the pod: nginx label is added to the nodes in the node pool. You can set the `nodeSelector` or `nodeAffinity` field in pod configurations to ensure that an application pod is scheduled to nodes with matching labels in a node pool. Perform the following steps:

- o Set `nodeSelector`.

`nodeSelector` is a field in the `spec` section of pod configurations. Add the pod: nginx label to `nodeSelector`. Sample template:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: nginx-deployment-basic
  labels:
    app: nginx
spec:
  replicas: 2
  selector:
    matchLabels:
      app: nginx
  template:
    metadata:
      labels:
        app: nginx
    spec:
      nodeSelector:
        pod: nginx # After you add the label in this field, the application pod can run only on nodes with this label in
the node pool.
      containers:
        - name: nginx
          image: nginx:1.7.9
          ports:
            - containerPort: 80
```

- o Set `nodeAffinity`.

You can also use `nodeAffinity` to schedule an application pod based on your requirements. `nodeAffinity` supports the following scheduling policies:

- `requiredDuringSchedulingIgnoredDuringExecution`

If this policy is used, a pod can be scheduled only to a node that meets the match rules. If no node meets the match rules, the system retries until a node that meets the rules is found. `IgnoreDuringExecution` indicates that if the label of the node where the pod is deployed changes and no longer meets the match rules, the pod continues to run on the node.

- requiredDuringSchedulingRequiredDuringExecution

If this policy is used, the pod can be scheduled only to a node that meets the match rules. If no node meets the rules, the system retries until a node that meets the rules is found. RequiredDuringExecution indicates that if the label of the node where the pod is deployed changes and no longer meets the match rules, the system redeploys the pod to another node that meets the rules.

- preferredDuringSchedulingIgnoredDuringExecution

If this policy is used, the pod is preferably scheduled to a node that meets the match rules. If no node meets the rules, the system ignores the rules.

- preferredDuringSchedulingRequiredDuringExecution

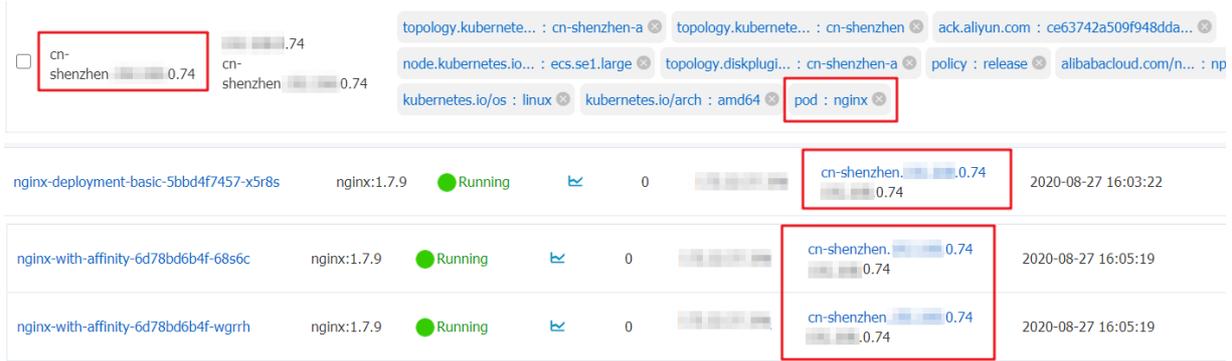
If this policy is used, the pod is preferably scheduled to a node that meets the match rules. If no node meets the rules, the system ignores the rules. RequiredDuringExecution indicates that if the label of a node where the pod is deployed changes and still meets the match rules, the system reschedules the pod to a node that meets the match rules.

In the following example, the requiredDuringSchedulingIgnoredDuringExecution policy is used to ensure that the application pod always runs on a node in a specific node pool.

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: nginx-with-affinity
  labels:
    app: nginx-with-affinity
spec:
  replicas: 2
  selector:
    matchLabels:
      app: nginx-with-affinity
  template:
    metadata:
      labels:
        app: nginx-with-affinity
    spec:
      affinity:
        nodeAffinity:
          requiredDuringSchedulingIgnoredDuringExecution:
            nodeSelectorTerms:
              - matchExpressions:
                  - key: pod
                    operator: In # This policy ensures that the application pod can run only on a node that has the pod: nginx label.
                    values:
                      - nginx
      containers:
        - name: nginx-with-affinity
          image: nginx:1.7.9
          ports:
            - containerPort: 80
```

Result

In the preceding example, all application pods are scheduled to the xxx.xxx.0.74 node. This node has the pod: nginx label.



3.6.4. Storage

3.6.4.1. Overview

In the Container Service console, you can create volumes of other Apsara Stack services, enabling you to create stateful applications and use Apsara Stack disks and OSS to implement persistent storage.

Both static and dynamic volumes are supported. The following table shows how static and dynamic volumes are supported.

| Apsara Stack storage | Static volume | Dynamic volume |
|----------------------|--|----------------|
| Apsara Stack disk | <p>You can use a static disk volume through either of the following methods:</p> <ul style="list-style-type: none"> Use a volume directly Use a volume through a PV and PVC | Supported |
| Apsara Stack NAS | <p>You can use a static NAS volume through either of the following methods:</p> <ul style="list-style-type: none"> Use a volume through the FlexVolume plug-in <ul style="list-style-type: none"> Use a volume directly Use a volume through a PV or PVC Use a volume through the Kubernetes NFS driver | Supported |
| Apsara Stack OSS | <p>You can use a static OSS volume through either of the following methods:</p> <ul style="list-style-type: none"> Use a volume directly Use a volume through a PV or PVC | Not supported |

3.6.4.2. Mount disk volumes

You can mount disks as volumes.

Container Service allows you to mount disks as persistent volumes (PVs) in Kubernetes clusters.

Disks can be mounted to Kubernetes clusters as the following volume types:

- **Statically provisioned disk volumes**

You can use statically provisioned disk volumes in the following ways:

- **Mount disks as volumes**
- **Mount disk volumes by creating a PV and a persistent volume claim (PVC)**

- **Dynamically provisioned disk volumes**

Usage notes

- You can mount a disk only to one pod.
- Before you mount a disk to a pod, you must create the disk and obtain its disk ID.
The disk must meet the following capacity requirements:
 - If the disk is a basic disk, it must be at least 5 GiB in size.
 - If the disk is an ultra disk, it must be at least 20 GiB in size.
 - If the disk is a standard SSD, it must be at least 20 GiB in size.
- `volumeId`: the ID of the disk that you want to mount. The value must be the same as those of `volumeName` and `PV Name`.
- A disk can be mounted only to a node that is deployed in the same zone as the disk.
- Only pay-as-you-go disks can be mounted. If you change the billing method of an Elastic Compute Service (ECS) instance in the cluster from pay-as-you-go to subscription, you cannot change the billing method of its disks to subscription. Otherwise, the disks cannot be mounted to the cluster.

Statically provisioned disk volumes

You can mount disks as volumes or by creating PVs and PVCs.

Prerequisites

A disk is created in the ECS console.

- **Mount a disk as a volume**

Use the following `disk-deploy.yaml` file to create a pod:

```
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: nginx-disk-deploy
spec:
  replicas: 1
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
      - name: nginx-flexvolume-disk
        image: nginx
        volumeMounts:
        - name: "d-bp1j17ifxfasvts3tf40"
          mountPath: "/data"
      volumes:
      - name: "d-bp1j17ifxfasvts3tf40"
        flexVolume:
          driver: "alicloud/disk"
          fsType: "ext4"
          options:
            volumeId: "d-bp1j17ifxfasvts3tf40"
```

- **Mount disk volumes by creating a PV and a PVC**

- i. **Create a PV of the disk type**

You can create a PV of the disk type in the Container Service console or by using a YAML file.

- **Create a PV by using a YAML file**

Use the following `disk-pv.yaml` file to create a PV:

 **Note** The PV name must be the same as the disk ID.

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: d-bp1j17ifxfasvts3tf40
  labels:
    failure-domain.beta.kubernetes.io/zone: cn-hangzhou-b
    failure-domain.beta.kubernetes.io/region: cn-hangzhou
spec:
  capacity:
    storage: 20Gi
  storageClassName: disk
  accessModes:
    - ReadWriteOnce
  flexVolume:
    driver: "alicloud/disk"
    fsType: "ext4"
    options:
      volumelId: "d-bp1j17ifxfasvts3tf40"
```

- **Create a PV in the console**

- Log on to the [Container Service console](#).
- In the left-side navigation pane, click **Clusters**.
- On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
- In the left-side navigation pane of the details page, choose **Volumes > Persistent Volumes**.
- On the **Persistent Volumes** page, click **Create** in the upper-right corner.
- In the Create PV dialog box, set the parameters. PV parameters

| Parameter | Description |
|------------------|--|
| PV Type | In this example, Cloud Disk is selected. |
| Volume Plug-in | Displays the supported storage drivers. |
| Access Mode | By default, ReadWriteOnce is selected. |
| Disk ID | Select a disk that is in the same region and zone as your cluster. |
| File System Type | Select the file system of the disk. Supported file systems are ext4 , ext3 , xf s, and vf at. Default value: ext4 . |
| Label | Add labels to the PV. |

- Click **Create**.

ii. Create a PVC

Use the following `disk-pvc.yaml` file to create a PVC:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: pvc-disk
spec:
  accessModes:
    - ReadWriteOnce
  storageClassName: disk
  resources:
    requests:
      storage: 20Gi
```

iii. Create a pod

Use the following `disk-pod.yaml` file to create a pod:

```
apiVersion: v1
kind: Pod
metadata:
  name: "flexvolume-alicloud-example"
spec:
  containers:
    - name: "nginx"
      image: "nginx"
      volumeMounts:
        - name: pvc-disk
          mountPath: "/data"
  volumes:
    - name: pvc-disk
      persistentVolumeClaim:
        claimName: pvc-disk
```

Dynamically provisioned disk volumes

To mount a disk as a dynamically provisioned volume, you must create a StorageClass of the disk type and specify the StorageClass in the `storageClassName` field of the PVC.

1. Create a StorageClass

```
kind: StorageClass
apiVersion: storage.k8s.io/v1beta1
metadata:
  name: alicloud-disk-common-hangzhou-b
provisioner: alicloud/disk
parameters:
  type: cloud_ssd
  regionid: cn-hangzhou
  zoneid: cn-hangzhou-b
```

Required parameters:

- `provisioner`: Set this parameter to `alicloud/disk`. This indicates that the Provisioner plug-in is used to create the StorageClass.
- `type`: Specify the disk type. Valid values: `cloud`, `cloud_efficiency`, `cloud_ssd`, and `available`. If you set this parameter to `available`, the system attempts to create a disk in the following order: ultra disk, standard SSD, and basic disk. The system stops trying until a disk is created.
- `regionid`: Specify the region of the disk.

- zoneid: Specify the zone of the disk.

2. Create a PVC

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: disk-common
spec:
  accessModes:
    - ReadWriteOnce
  storageClassName: alicloud-disk-common-hangzhou-b
  resources:
    requests:
      storage: 20Gi
---
kind: Pod
apiVersion: v1
metadata:
  name: disk-pod-common
spec:
  containers:
    - name: disk-pod
      image: nginx
      volumeMounts:
        - name: disk-pvc
          mountPath: "/mnt"
  restartPolicy: "Never"
  volumes:
    - name: disk-pvc
      persistentVolumeClaim:
        claimName: disk-common
```

Default options

By default, Kubernetes clusters support the following types of StorageClass:

- alicloud-disk-common: basic disk.
- alicloud-disk-efficiency: ultra disk.
- alicloud-disk-ssd: standard SSD.
- alicloud-disk-available: This option ensures high availability. The system attempts to create an ultra disk first. If no ultra disk is available in the specified zone, the system attempts to create a standard SSD. If no standard SSD is available, the system attempts to create a basic disk.

3. Create a multi-instance StatefulSet by using a disk

We recommend that you use the volumeClaimTemplates parameter. This parameter allows the system to dynamically create PVCs and PVs. PVCs are associated with corresponding PVs.

```
apiVersion: v1
kind: Service
metadata:
  name: nginx
labels:
  app: nginx
spec:
  ports:
  - port: 80
    name: web
  clusterIP: None
  selector:
    app: nginx
---
apiVersion: apps/v1beta2
kind: StatefulSet
metadata:
  name: web
spec:
  selector:
    matchLabels:
      app: nginx
  serviceName: "nginx"
  replicas: 2
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
      - name: nginx
        image: nginx
        ports:
        - containerPort: 80
          name: web
        volumeMounts:
        - name: disk-common
          mountPath: /data
  volumeClaimTemplates:
  - metadata:
      name: disk-common
    spec:
      accessModes: [ "ReadWriteOnce" ]
      storageClassName: "alicloud-disk-common"
    resources:
      requests:
        storage: 10Gi
```

3.6.4.3. Mount NAS volumes

Container Service allows you to mount Apsara File Storage NAS (NAS) file systems as persistent volumes (PVs) in Container Service clusters.

NAS file systems can be mounted to Kubernetes clusters as the following volume types:

- [Statically provisioned NAS volumes](#)

You can statically provision NAS volumes in the following ways:

- Use the FlexVolume plug-in
 - Directly mount NAS file systems as volumes
 - Mount NAS file systems by creating a PV and a persistent volume claim (PVC)
- Use the Network File System (NFS) driver
- **Dynamically provisioned NAS volumes**

Prerequisites

A NAS file system is created in the NAS console and a mount target is added. The mount target is used to mount the NAS file system to the Kubernetes cluster. The NAS file system and your cluster are deployed in the same virtual private cloud (VPC).

Statically provisioned NAS volumes

You can use the FlexVolume plug-in provided by Alibaba Cloud or the NFS driver provided by Kubernetes to mount NAS file systems.

Use the FlexVolume plug-in

You can use the FlexVolume plug-in to directly mount a NAS file system as a volume. You can also mount a NAS file system by creating a PV and a PVC.

Note

- NAS is a shared storage service. You can mount a NAS file system to multiple pods.
- server: the mount target of the NAS volume.
- path: the directory to which the NAS volume is mount. You can specify a subdirectory. If the specified subdirectory does not exist, the system automatically creates the subdirectory.
- vers: the version of the NFS protocol. Version 4.0 is supported.
- mode: the access permissions on the directory of the NAS file system. If the root directory of the NAS file system is specified, you cannot modify the access permissions. If the NAS file system stores a large amount of data, the mounting operation may be time-consuming or even fail. Therefore, we recommend that you do not set the mode parameter.

Mount a NAS file system as a volume

Use the following `nas-deploy.yaml` file to create a pod:

```
apiVersion: v1
kind: Pod
metadata:
  name: "flexvolume-nas-example"
spec:
  containers:
    - name: "nginx"
      image: "nginx"
      volumeMounts:
        - name: "nas1"
          mountPath: "/data"
  volumes:
    - name: "nas1"
      flexVolume:
        driver: "alicloud/nas"
        options:
          server: "0cd8b4a576-grs79.cn-hangzhou.nas.aliyuncs.com"
          path: "/k8s"
          vers: "4.0"
```

Mount a NAS file system by creating a PV and a PVC

Step 1: Create a PV

You can create a PV in the Container Service console or by using a YAML file.

- Create a PV by using a YAML file.

Use the following `nas-pv.yaml` file to create a PV:

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: pv-nas
spec:
  capacity:
    storage: 5Gi
  storageClassName: nas
  accessModes:
    - ReadWriteMany
  flexVolume:
    driver: "alicloud/nas"
    options:
      server: "0cd8b4a576-uih75.cn-hangzhou.nas.aliyuncs.com"
      path: "/k8s"
      vers: "4.0"
```

- Create a PV in the console
 - [Log on to the Container Service console.](#)
 - In the left-side navigation pane, click **Clusters**.
 - On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
 - In the left-side navigation pane of the details page, choose **Volumes > Persistent Volumes**.
 - On the **Persistent Volumes** page, click **Create** in the upper-right corner.
 - In the Create PV dialog box, set the parameters.

| Parameter | Description |
|---------------------------------|--|
| PV Type | Select NAS . |
| Volume Name | Enter a name for the PV. The name must be unique in the cluster. |
| Volume Plug-in | By default, CSI is selected. |
| Capacity | Specify the capacity of the PV. The capacity of the PV cannot exceed the capacity of the disk. |
| Access Mode | By default, ReadWriteMany is selected. |
| Mount Target Domain Name | The domain name of the mount target that is used to mount the NAS file system to the cluster. |

| Parameter | Description |
|--------------|--|
| Subdirectory | <p>Enter a subdirectory in the NAS file system. The subdirectory must start with a forward slash (/). If you set this parameter, the PV is mounted to the specified subdirectory.</p> <ul style="list-style-type: none"> ▪ If the specified subdirectory does not exist, the system automatically creates the subdirectory in the NAS file system. ▪ If you do not set this parameter, the PV is mounted to the root directory of the NAS file system. |
| Version | The version of the NFS protocol. Version 3 and 4.0 are supported. By default, version 3 is used. We recommend that you use version 3. |
| Label | Add labels to the PV. |

vii. Click **Create**.

Step 2: Create a PVC

Use the following *nas-pvc.yaml* file to create a PVC:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: pvc-nas
spec:
  accessModes:
    - ReadWriteMany
  storageClassName: nas
resources:
  requests:
    storage: 5Gi
```

Step 3: Create a pod

Use the following *nas-pod.yaml* file to create a pod:

```
apiVersion: v1
kind: Pod
metadata:
  name: "flexvolume-nas-example"
spec:
  containers:
    - name: "nginx"
      image: "nginx"
      volumeMounts:
        - name: pvc-nas
          mountPath: "/data"
  volumes:
    - name: pvc-nas
      persistentVolumeClaim:
        claimName: pvc-nas
```

Use the NFS driver

Step 1: Create a NAS file system

Log on to the NAS console. For more information about how to log on to the NAS console, see the *Create a NAS file system* chapter of the *NAS User Guide*.

Note The NAS file system that you want to create and the Kubernetes cluster must be deployed in the same region.

In this example, the following mount target is used: `055f84ad83-ixxxx.cn-hangzhou.nas.aliyuncs.com`.

Step 2: Create a PV

You can create a PV in the console or by using a YAML file.

- **Create a PV by using a YAML template**

Use the following *nas-pv.yaml* file to create a PV.

Run the following command to create a PV that uses the NAS file system:

```
root@master # cat << EOF | kubectl apply -f -
apiVersion: v1
kind: PersistentVolume
metadata:
  name: nas
spec:
  capacity:
    storage: 8Gi
  accessModes:
    - ReadWriteMany
  persistentVolumeReclaimPolicy: Retain
  nfs:
    path: /
    server: 055f84ad83-ixxxx.cn-hangzhou.nas.aliyuncs.com
EOF
```

- **Create a PV by using the console**

For more information, see [Mount a NAS file system by creating a PV and a PVC](#).

Step 3: Create a PVC

Create a PVC and associate the PVC with the PV that is created in Step 2.

```
root@master # cat << EOF | kubectl apply -f -
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: nasclaim
spec:
  accessModes:
    - ReadWriteMany
  resources:
    requests:
      storage: 8Gi
EOF
```

Step 4: Create a pod

Create an application to use the PV.

```
root@master # cat << EOF |kubectl apply -f-
apiVersion: v1
kind: Pod
metadata:
  name: mypod
spec:
  containers:
  - name: myfrontend
    image: registry.aliyuncs.com/spacexnice/netdia:latest
    volumeMounts:
    - mountPath: "/var/www/html"
      name: mypd
  volumes:
  - name: mypd
    persistentVolumeClaim:
      claimName: nasclaim
EOF
```

The NAS file system is mounted to the application that runs in the pod.

Dynamically provisioned NAS volumes

If you want to dynamically provision NAS volumes, you must install a driver plug-in and configure a NAS mount target.

Install the plug-in

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: alicloud-nas
provisioner: alicloud/nas
---
apiVersion: v1
kind: ServiceAccount
metadata:
  name: alicloud-nas-controller
  namespace: kube-system
---
kind: ClusterRoleBinding
apiVersion: rbac.authorization.k8s.io/v1beta1
metadata:
  name: run-alicloud-nas-controller
subjects:
  - kind: ServiceAccount
    name: alicloud-nas-controller
    namespace: kube-system
roleRef:
  kind: ClusterRole
  name: alicloud-disk-controller-runner
  apiGroup: rbac.authorization.k8s.io
---
kind: Deployment
apiVersion: extensions/v1beta1
metadata:
  name: alicloud-nas-controller
  namespace: kube-system
spec:
  replicas: 1
  strategy:
```

```
type: Recreate
template:
  metadata:
    labels:
      app: alicloud-nas-controller
  spec:
    tolerations:
      - effect: NoSchedule
        operator: Exists
        key: node-role.kubernetes.io/master
      - effect: NoSchedule
        operator: Exists
        key: node.cloudprovider.kubernetes.io/uninitialized
    nodeSelector:
      node-role.kubernetes.io/master: ""
    serviceAccount: alicloud-nas-controller
    containers:
      - name: alicloud-nas-controller
        image: registry.cn-hangzhou.aliyuncs.com/acs/alibabacloud-nas-controller:v1.8.4
        volumeMounts:
          - mountPath: /persistentvolumes
            name: nfs-client-root
        env:
          - name: PROVISIONER_NAME
            value: alicloud/nas
          - name: NFS_SERVER
            value: 0cd8b4a576-mmi32.cn-hangzhou.nas.aliyuncs.com
          - name: NFS_PATH
            value: /
    volumes:
      - name: nfs-client-root
        nfs:
          server: 0cd8b4a576-mmi32.cn-hangzhou.nas.aliyuncs.com
          path: /
```

Dynamically provision a NAS volume

```
apiVersion: apps/v1beta1
kind: StatefulSet
metadata:
  name: web
spec:
  serviceName: "nginx"
  replicas: 2
  volumeClaimTemplates:
  - metadata:
    name: html
    spec:
      accessModes:
      - ReadWriteOnce
      storageClassName: alicloud-nas
      resources:
        requests:
          storage: 2Gi
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
      - name: nginx
        image: nginx:alpine
        volumeMounts:
        - mountPath: "/usr/share/nginx/html/"
          name: html
```

3.6.4.4. Mount OSS volumes

You can mount Object Storage Service (OSS) buckets as volumes in Kubernetes clusters.

You can mount OSS buckets in the following ways:

- Mount an OSS bucket as a volume.
- Mount an OSS bucket by creating a PV and a PVC.

Prerequisites

To mount an OSS bucket as a statically provisioned volume, you must create an OSS bucket in the OSS console.

Background

- OSS buckets can be mounted only as statically provisioned volumes.
- An OSS bucket can be shared by multiple pods.
- bucket: Only buckets can be mounted to a Kubernetes cluster. The subdirectories or files in a bucket cannot be mounted to a Kubernetes cluster.
- url: specifies the endpoint of the OSS bucket. The endpoint is the domain name that is used to mount the OSS bucket to the Kubernetes cluster.
- akId: specifies your AccessKey ID.
- akSecret: specifies your AccessKey secret.
- otherOpts: the custom parameters that are used to mount the OSS bucket. The parameters must be in the following format: `-o *** -o ***`.

 **Note** To mount an OSS bucket as a volume, you must create a Secret with your AccessKey pair when you deploy the flexvolume Service.

Mount an OSS bucket as a statically provisioned volume

- **Mount an OSS bucket as a volume**

Use the following *oss-deploy.yaml* file to create a pod:

```
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: nginx-oss-deploy
spec:
  replicas: 1
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
        - name: nginx-flexvolume-oss
          image: nginx
          volumeMounts:
            - name: "oss1"
              mountPath: "/data"
      volumes:
        - name: "oss1"
          flexVolume:
            driver: "alicloud/oss"
            options:
              bucket: "docker"
              url: "oss-cn-hangzhou.aliyuncs.com"
              akId: ***
              akSecret: ***
              otherOpts: "-o max_stat_cache_size=0 -o allow_other"
```

- **Create a static PV and a PVC**

- i. **Create a PV**

You can create a persistent volume (PV) in the Container Service console or by using a YAML file.

- **Create a PV by using a YAML file**

Use the following `oss-pv.yaml` file to create a PV:

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: pv-oss
spec:
  capacity:
    storage: 5Gi
  accessModes:
    - ReadWriteMany
  storageClassName: oss
  flexVolume:
    driver: "alicloud/oss"
    options:
      bucket: "docker"
      url: "oss-cn-hangzhou.aliyuncs.com"
      akId: ***
      akSecret: ***
      otherOpts: "-o max_stat_cache_size=0 -o allow_other"
```

- **Create a PV in the Container Service console**

- a. [Log on to the Container Service console](#).
- b. In the left-side navigation pane, click **Clusters**.
- c. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
- d. In the left-side navigation pane of the details page, choose **Volumes > Persistent Volumes**.
- e. On the **Persistent Volumes** page, click **Create** in the upper-right corner.

f. In the Create PV dialog box, set the parameters.

| Parameter | Description |
|-----------------------------------|---|
| PV Type | In this example, OSS is selected. |
| Name | Enter a name for the PV. The name must be unique in the cluster. In this example, pv-oss is used. |
| Volume Plug-in | By default, CSI is selected. |
| Capacity | Specify the capacity of the PV. |
| Access Mode | By default, ReadWriteMany is selected. |
| AccessKey ID and AccessKey Secret | The AccessKey pair that is required to access OSS buckets. To obtain your AccessKey pair, go to the Apsara Uni-manager Management Console, choose Enterprise > Organizations , click  on the right side of the organization. Then, click AccessKey and copy the AccessKey pair. |
| Optional Parameters | Enter custom parameters in the format of <code>-o *** -o ***</code> . |
| Bucket ID | Enter the name of the OSS bucket that you want to mount. Click Select Bucket . In the dialog box that appears, select the OSS bucket that you want to mount and click Select . |
| Endpoint | Internal Endpoint is recommended. |
| Label | Add labels to the PV. |

g. Click **Create**.

ii. Create a PVC

Use the following `oss-pvc.yaml` file to create a persistent volume claim (PVC):

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: pvc-oss
spec:
  storageClassName: oss
  accessModes:
    - ReadWriteMany
  resources:
    requests:
      storage: 5Gi
```

iii. Create a pod

Use the following `oss-pod.yaml` file to create a pod:

```
apiVersion: v1
kind: Pod
metadata:
  name: "flexvolume-oss-example"
spec:
  containers:
    - name: "nginx"
      image: "nginx"
      volumeMounts:
        - name: pvc-oss
          mountPath: "/data"
  volumes:
    - name: pvc-oss
      persistentVolumeClaim:
        claimName: pvc-oss
```

Can I mount an OSS bucket as a dynamically provisioned volume?

No. Dynamically provisioned OSS volumes are not supported.

3.6.4.5. Create a PVC

This topic describes how to create a persistent volume claim (PVC).

Prerequisites

- A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).
- A persistent volume (PV) is created. In this example, a PV created from a disk is used. For more information, see [Use Apsara Stack disks](#).

By default, PVCs are associated with PVs that have the alicloud-pvname label. This label is added to all PVs that are created in the Container Service console. If a PV does not have this label, manually add the label to the PV before you can associate the PV with a PVC.

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Volumes > Persistent Volume Claims**.
5. In the upper-right corner of the **Persistent Volume Claims** page, click **Create**.
6. In the **Create PVC** dialog box, set the parameters and click **Create**.
 - **PVC Type**: The PVC and PV must be of the same type. You can select Cloud Disk, NAS, and OSS.
 - **Name**: Enter the name of the PVC.
 - **Allocation Mode**: **Use StorageClass**, **Existing Volumes**, and **Create Volume** are supported. In this example, **Use StorageClass** or **Existing Volumes** is selected.
 - **Existing Storage Class**: Click **Select**. In the Select Storage Class dialog box, find the Storage Class that you want to use and click **Select** in the Actions column. This parameter is required only when you set Allocation Mode to **Use StorageClass**.
 - **Existing Volumes**: Click **Select PV**. In the Select PV dialog box, find the PV that you want to use and click **Select** in the Actions column. This parameter is required only when you set Allocation Mode to **Existing Volumes**.
 - **Capacity**: Specifies the capacity used by the PVC. The value cannot be larger than the total capacity of the associated PV.

- **Access Mode:** The default value is ReadWriteOnce. This parameter is required only when you set Allocation Mode to Use StorageClass.

Note If your cluster has a PV that is not used, but you cannot find it in the Select PV dialog box, the reason may be that the PV does not have the alicloud-pvname label.

If you cannot find available PVs, click **Persistent Volumes** in the left-side navigation pane. On the Persistent Volumes page, find the PV that you want to use, click **Manage Labels** in the Actions column, and then add a label to the PV. On the Manage Labels dialog box, set the key of the label to alicloud-pvname and the value to the name of the PV. If the PV is created from a disk, the disk ID is used as the name of the PV.

| Name | Value |
|--|-----------------------|
| alicloud-pvname | d-bp1450080c7emk51w0e |
| failure-domain.beta.kubernetes.io/zone | cn-hangzhou-g |
| failure-domain.beta.kubernetes.io/region | cn-hangzhou |

7. Go to the Persistent Volume Claims page. You can find the newly created PVC in the list.

3.6.4.6. Use PVCs

You can use persistent volume claims (PVCs) to mount volumes for applications.

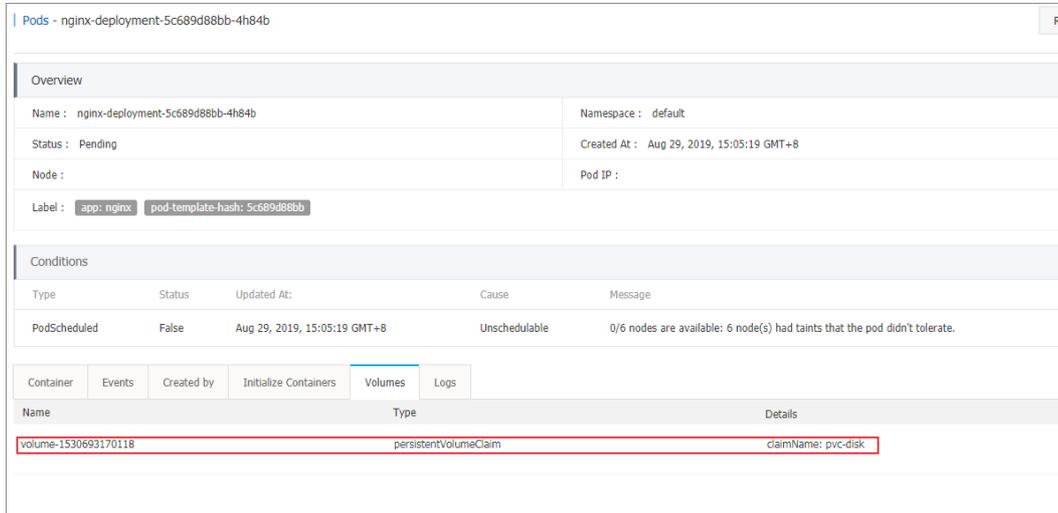
Prerequisites

- A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).
- A PVC is created. In this example, a PVC named pvc-disk is created to mount a disk volume. For more information, see [Create PVCs](#).

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Workloads > Deployments**.
5. On the **Deployments** page, click **Create from Image** in the upper-right corner.
6. On the **Basic Information** wizard page, specify the application name, number of replicas, type, and labels, select whether to synchronize the time zone from nodes to containers, and then click **Next**.
7. On the **Container** wizard page, select an image and configure volumes of the cloud storage type. You can select Cloud Disk, Apsara File Storage NAS (NAS), and Object Service Storage (OSS). In this example, select the PVC named pvc-disk and click **Next**. For more information, see [Container configurations](#).
8. On the **Advanced** wizard page, create a Service for the test-nginx application and click **Create**.

9. On the **Complete** wizard page, click View Details to view detailed information about the application. You are redirected to the details page of the test-nginx application.
10. On the **Pods** tab, you can find the pods to which the application belongs. Select a pod and click **View Details** to view detailed information about the pod.
11. On the details page of the pod, click the **Volumes** tab. You can find that the pod is associated with the pvc-disk PVC.



3.6.5. Network management

3.6.5.1. Set access control for pods

This topic describes how to use network policies to control access between pods.

Prerequisites

You have created a Kubernetes cluster and selected the **Terway network plug-in**. For more information, see [Create a Kubernetes cluster](#).

Context

You can declare network policies to control access between pods and thus prevent applications from interfering each other.

Procedure

For more information about standard Kubernetes network policies, see [Network policies](#).

1. Create a pod that runs as a server and attach `label run=nginx` to the pod. For more information, see [Create an application from an orchestration template](#).

The sample YAML file is as follows:

```

apiVersion: v1
kind: Pod
metadata:
  name: server
  labels:
    run: nginx
spec:
  containers:
    - name: nginx
      image: registry.acs.intranet.env22.com/nginx:1.8

```

2. Create a network policy. For more information, see [Create an application from an orchestration template](#).

The sample YAML file is as follows:

```

kind: NetworkPolicy
apiVersion: networking.k8s.io/v1
metadata:
  name: access-nginx
spec:
  podSelector:
    matchLabels:
      run: nginx # Apply the network policy to pods with the run=nginx label
  ingress:
    - from:
      - podSelector:
          matchLabels:
            access: "true" # Only pods with the access=true label are accessible

```

3. Use the *client.yaml* and *client-label* files to create two pods that run as clients.

One pod has the required label and the other does not.

- i. Create the *client.yaml* and *client-label* files with the following contents respectively.

```

# This pod has no label
apiVersion: v1
kind: Pod
metadata:
  name: client
spec:
  containers:
    - name: busybox
      image: registry.acs.intranet.env22.com/acs/busybox
      command: ["sh", "-c", "sleep 200000"]

```

```

# This pod has the label
apiVersion: v1
kind: Pod
metadata:
  name: client-label
  labels:
    access: "true"
spec:
  containers:
    - name: busybox
      image: registry.acs.intranet.env22.com/acs/busybox
      command: ["sh", "-c", "sleep 200000"]

```

- ii. Run the following commands to create these pods:

```
kubectl apply -f client.yaml
kubectl apply -f client-label.yaml
```

You can see that only the pod with the required label can access the server.

3.6.5.2. Set bandwidth limits for pods

This topic describes how to limit the bandwidth of inbound and outbound traffic that flows through a pod.

Prerequisites

You have created a Kubernetes cluster and selected the **Terway network plug-in**. For more information, see [Create a Kubernetes cluster](#).

Context

Throttling pods helps prevent performance degradation of the host or other workloads when certain pods occupy excessive resources.

Method

You can use the `k8s.aliyun.com/ingress-bandwidth` and `k8s.aliyun.com/egress-bandwidth` annotations for pod throttling.

- `k8s.aliyun.com/ingress-bandwidth` : limits the pod inbound bandwidth.
- `k8s.aliyun.com/egress-bandwidth` : limits the pod outbound bandwidth.
- The bandwidth limit is measured in m and k, which represent Mbit/s and Kbit/s respectively.

Procedure

1. Create a pod that runs as a server in the console. For more information, see [Create an application from an orchestration template](#).

The sample YAML file is as follows:

```
apiVersion: v1
kind: Pod
metadata:
  name: server
  annotations:
    k8s.aliyun.com/ingress-bandwidth: 10m # Set the inbound bandwidth limit to 10 Mbit/s
    k8s.aliyun.com/egress-bandwidth: 10m
spec:
  containers:
    - name: nginx
      image: registry.acs.intranet.env22.com/nginx:1.8
```

2. Run the `kubectl exec` command to connect to the pod. To verify that pod throttling is effective, run the following commands to create a file on the pod. Assume that the IP address of the pod created in [step 1](#) is 172.16.XX.XX.

```
cd /usr/share/nginx/html
dd if=/dev/zero of=bigfile bs=1M count=1000
```

3. Use the `client-deploy.yaml` file to create a pod that runs as a client.

i. Create the `client-deploy.yaml` file with the following content :

```

apiVersion: v1
kind: Pod
metadata:
  name: client
  annotations:
    k8s.aliyun.com/ingress-bandwidth: 10m # Set the inbound bandwidth limit to 10 Mbit/s
    k8s.aliyun.com/egress-bandwidth: 10m
spec:
  containers:
    - name: busybox
      image: registry.acs.intranet.env22.com/acs/netdia
      command: ["sh", "-c", "sleep 200000"]

```

ii. Run the following command to create the pod:

```
kubectl apply -f client-deploy.yaml
```

4. Run the following command to check whether bandwidth is limited:

```
kubectl exec -it client sh
```

3.6.5.3. Work with Terway

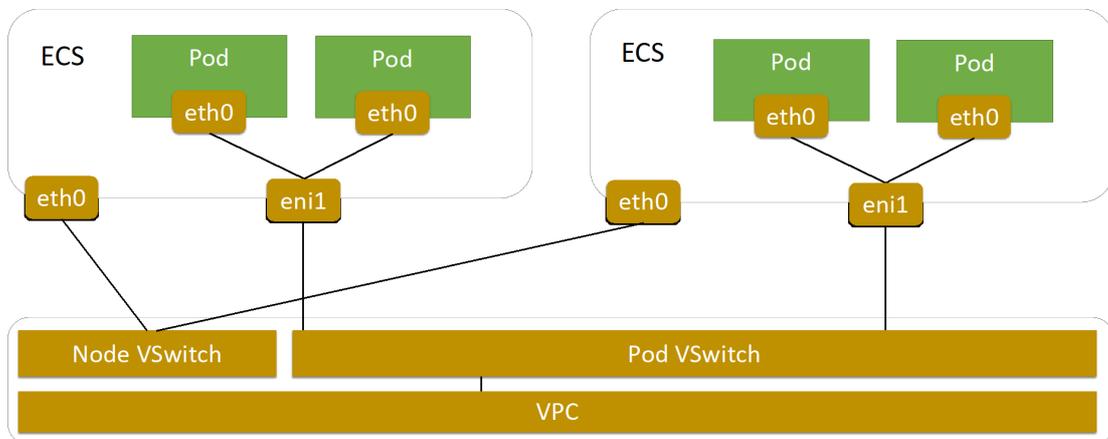
Terway is an open source Container Network Interface (CNI) plug-in developed by Alibaba Cloud. Terway works with Virtual Private Cloud (VPC) and allows you to use standard Kubernetes network policies to regulate how containers communicate with each other. You can use Terway to enable internal communication within a Kubernetes cluster. This topic describes how to use Terway.

Context

Terway is a network plug-in developed by Alibaba Cloud for Container Service. Terway allows you to set up networks for pods by associating Alibaba Cloud elastic network interfaces (ENIs) with the pods. Terway also allows you to use standard Kubernetes network policies to regulate how containers communicate with each other. In addition, Terway is compatible with Calico network policies.

In a cluster that has Terway installed, each pod has a separate network stack and is assigned a separate IP address. Pods on the same Elastic Compute Service (ECS) instance communicate with each other by forwarding packets inside the ECS instance. Pods on different ECS instances communicate with each other through ENIs in the VPC where the ECS instances are deployed. This improves communication efficiency because no tunneling technologies such as Virtual Extensible Local Area Network (VXLAN) are required to encapsulate packets.

How Terway works



Comparison between Flannel and Terway

When you create a Kubernetes cluster, you can choose one of the following network plug-ins:

- **Terway:** a network plug-in developed by Alibaba Cloud for Container Service. Terway allows you to assign ENIs to containers and use standard Kubernetes network policies to regulate how containers communicate with each other. Terway also supports bandwidth throttling on individual containers. Terway uses flexible IP Address Management (IPAM) policies to allocate IP addresses to containers. This avoids IP address waste. If you do not want to use network policies, you can select Flannel as the network plug-in. Otherwise, we recommend that you select Terway.
- **Flannel:** an open source CNI plug-in, which is simple and stable. You can use Flannel with VPC of Alibaba Cloud. This allows your clusters and containers to run in high-performance and stable networks. However, Flannel provides only basic features. It does not support standard Kubernetes network policies. For more information, see [Flannel](#).

| Item | Terway | Flannel |
|-----------------------|--|--|
| Performance | The IP address of each pod in a Kubernetes cluster is assigned from the CIDR block of the VPC where the cluster is deployed. Therefore, you do not need to use the Network Address Translation (NAT) service to translate IP addresses. This avoids IP address waste. In addition, each pod in the cluster can use an exclusive ENI. | Flannel works with VPC of Alibaba Cloud. The CIDR block of pods that you specify must be different from that of the VPC where the cluster is deployed. Therefore, the NAT service is required and some IP addresses may be wasted. |
| Security | Terway supports network policies. | Flannel does not support network policies. |
| IP address management | Terway allows you to assign IP addresses on demand. You do not have to assign CIDR blocks by node. This avoids IP address waste. | You can only assign CIDR blocks by node. In large-scale clusters, a great number of IP addresses may be wasted. |
| SLB | Server Load Balancer (SLB) directly forwards network traffic to pods. You can upgrade the pods without service interruption. | SLB forwards network traffic to the NodePort Service. Then, the NodePort Service routes the network traffic to pods. |

Considerations

- To use the Terway plug-in, we recommend that you use ECS instances of higher specifications and newer types, such as ECS instances that belong to the g5 or g6 instance family with at least eight CPU cores. For more information, see the *Instance families* chapter of *ECS User Guide*.
- The maximum number of pods that each node supports is based on the number of ENIs assigned to the node.
 - Maximum number of pods supported by each shared ENI = (Number of ENIs supported by each ECS instance - 1) × Number of private IP addresses supported by each ENI
 - Maximum number of pods supported by each exclusive ENI = Number of ENIs supported by each ECS instance - 1

Step 1: Plan CIDR blocks

When you create a Kubernetes cluster, you must specify a VPC, vSwitches, the CIDR block of pods, and the CIDR block of Services. If you want to install the Terway plug-in, you must first create a VPC and two vSwitches in the VPC. The two vSwitches must be created in the same zone.

You can assign the following CIDR blocks for a cluster that uses Terway.

- VPC CIDR block: 192.168.0.0/16
- vSwitch CIDR block: 192.168.0.0/19

- CIDR block of pod vSwitch: 192.168.32.0/19
- Service CIDR block: 172.21.0.0/20

 **Note**

- IP addresses within the CIDR block of the vSwitch are assigned to nodes.
- IP addresses within the CIDR block of the pod vSwitch are assigned to pods.

The following example describes how to create a VPC and two vSwitches. The CIDR blocks in the preceding section are assigned in this example.

1. Log on to the VPC console.

 **Note** For more information, see the *Log on to the VPC console* chapter of *VPC*.

2. Create a VPC.
 - i. In the top navigation bar, select the region where you want to create the VPC.
 - ii. On the VPCs page, click **Create VPC**.

iii. On the **Create VPC** page, set the following parameters and click **Submit**.

| Parameter | Description |
|------------------------|---|
| Organization | Select the organization to which the VPC belongs. |
| Resource Set | Select the resource set to which the VPC belongs. |
| Region | Select the region where you want to deploy the VPC. |
| Sharing Scope | <p>Select the sharing scope of the VPC.</p> <ul style="list-style-type: none"> ▪ Current Resource Set: Only the administrator of the current resource set can use the VPC to create resources. ▪ Current Organization and Subordinate Organization: Only the administrators of the current organization and its subordinate organization can create resources for the shared VPC. ▪ Current Organization: Only the administrator of the current organization can use create resources for the shared VPC. |
| VPC Name | <p>Enter a name for the VPC. In this example, vpc_192_168_0_0_16 is used.</p> <p>The name must be 2 to 128 characters in length, and can contain letters, digits, underscores (_), and hyphens (-). The name must start with a letter and cannot start with <code>http://</code> or <code>https://</code>.</p> |
| IPv4 CIDR Block | <p>Select an IPv4 CIDR block for the VPC. In this example, 192.168.0.0/16 is specified.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 5px;"> <p> Note After a VPC is created, the IPv4 CIDR block of the VPC cannot be modified.</p> </div> |
| IPv6 CIDR Block | <p>Specify whether to assign an IPv6 CIDR block.</p> <ul style="list-style-type: none"> ▪ Do Not Assign: The system does not assign an IPv6 CIDR block to the VPC. ▪ Assign: An IPv6 CIDR block is automatically assigned to the VPC. <p>If you set this parameter to Assign, the system automatically creates a free IPv6 gateway for this VPC, and assigns an IPv6 CIDR block with the subnet mask /56, such as 2xx1: db8::/56. By default, IPv6 addresses can be used to communicate within only private networks. If you want to allow an instance assigned with an IPv6 address to access the Internet or be accessed by IPv6 clients over the Internet, you must purchase an Internet bandwidth plan for the IPv6 address. For more information, see Enable Internet connectivity for an IPv6 address the Activate IPv6 Internet bandwidth section of the Manage IPv6 Internet bandwidth topic of the <i>IPv6 gateway user guide</i>.</p> |
| Description | <p>Enter a description for the VPC.</p> <p>The description must be 2 to 256 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code>.</p> |

3. Create vSwitches.

- i. In the left-side navigation pane, click **vSwitches**.
- ii. Select the region of the VPC in which you want to create a vSwitch.
- iii. On the **vSwitches** page, click **Create vSwitch**.
- iv. On the **Create vSwitch** page, set the following parameters and click **Submit**.

 **Note** Make sure that the two vSwitches are in the same zone.

| Parameter | Description |
|---|---|
| Organization | Select the organization to which the vSwitch belongs. |
| Resource Set | Select the resource set to which the vSwitch belongs. |
| Region | Select the region where you want to deploy the vSwitch. |
| Zone | <p>Select the zone to which the vSwitch belongs.</p> <p>In a VPC, each vSwitch can be deployed in only one zone. You cannot deploy a vSwitch across zones. However, you can deploy cloud resources in vSwitches that belong to different zones to achieve cross-zone disaster recovery.</p> <p> Note A cloud instance can be deployed in only one vSwitch.</p> |
| Sharing Scope | <p>Select the sharing scope of the vSwitch.</p> <ul style="list-style-type: none"> ▪ Current Resource Set: Only the administrator of the current resource set can create resources in the shared vSwitch. ▪ Current Organization and Subordinate Organization: Only the administrators of the current organization and its subordinate organizations can create resources in the shared vSwitch. ▪ Current Organization: Only the administrator of the current organization can create resources in the shared vSwitch. |
| VPC | Select the VPC for which you want to create the vSwitch. |
| Dedicated for Out-of-cloud Physical Machines | <p>Specify whether the vSwitch to be created is dedicated for bare-metal instances.</p> <p>For more information about bare-metal instances, see the VPC bare-metal instance features topic in <i>BMS User Guide</i>.</p> |
| vSwitch Name | <p>Enter a name for the vSwitch. In this example, node_switch_192_168_0_0_19 is used.</p> <p>The name must be 2 to 128 characters in length and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code>.</p> |
| IPv4 CIDR Block | Specify an IPv4 CIDR block for the vSwitch. In this example, 192.168.0.0/19 is used. |
| IPv6 CIDR Block | <p>Specify an IPv6 CIDR block for the vSwitch.</p> <ul style="list-style-type: none"> ▪ You must check whether IPv6 is enabled for the specified VPC. If IPv6 is disabled, you cannot assign an IPv6 CIDR block to the vSwitch. ▪ If IPv6 is enabled, you can enter a decimal number ranging from 0 to 255 to define the last 8 bits of the IPv6 CIDR block of the vSwitch. <p>For example, if the IPv6 CIDR block of the VPC is 2xx1:db8::/64, specify 255 to define the last 8 bits of the IPv6 CIDR block. In this case, the IPv6 CIDR block of the vSwitch is 2xx1:db8:ff::/64. ff is the hexadecimal value of 255.</p> |

| Parameter | Description |
|-------------|---|
| Description | Enter a description for the vSwitch. The description must be 2 to 256 characters in length and can contain letters, digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). The name must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . |

- Repeat Step to create a pod vSwitch. Set the name of the pod vSwitch to `pod_switch_192_168_32_0_19` and IPv4 CIDR Block to `192.168.32.0/19`.

Step 2: Set up networks for a cluster that uses Terway

To install Terway in a cluster and set up networks for the cluster, set the following parameters.

Note A Kubernetes cluster is used as an example to show how to set up networks for a cluster that uses Terway as the network plug-in. For more information about how to create a Kubernetes cluster, see [Create a Kubernetes cluster](#).

- VPC:** Select the VPC created in [Step 1: Plan CIDR blocks](#).
- VSwitch:** Select the vSwitch created in [Step 1: Plan CIDR blocks](#).
- Network Plug-in:** Select **Terway**.
- Pod VSwitch:** Select the pod vSwitch created in [Step 1: Plan CIDR blocks](#).
- Service CIDR:** Use the default settings.

3.6.6. Namespaces

3.6.6.1. Create a namespace

You can create a namespace in the Container Service console.

Prerequisites

A Kubernetes cluster is created.

Context

In a Kubernetes cluster, you can use namespaces to create multiple virtual spaces. When a large number of users share a cluster, you can use namespaces to divide different workspaces and allocate cluster resources to different tasks. You can also use [resource quotas](#) to allocate resources to each namespace.

Procedure

- [Log on to the Container Service console](#).
- In the left-side navigation pane, click **Clusters**.
- On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
- In the left-side navigation pane of the details page, click **Namespaces and Quotas**.
- On the **Namespace** page, click **Create** in the upper-right corner.
- In the **Create Namespace** dialog box, configure the namespace.

Create a namespace

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|--------------|--|
| Name | Enter a name for the namespace. In this example, test is entered. The name must be 1 to 63 characters in length and can contain digits, letters, and hyphens (-). It must start and end with a letter or digit. |
| Label | <p>Label: Add one or more labels to the namespace. Labels are used to identify namespaces. For example, you can label a namespace as one used in the test environment.</p> <p>To add a label, enter a key and a value and click Add in the Actions column.</p> |

7. Click **OK**.

You can find the namespace that you created on the Namespace page.

3.6.6.2. Set resource quotas and limits

You can set resource quotas and limits for a namespace in the Container Service console.

Prerequisites

- A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).
- A sample namespace named `test` is created. For more information, see [Create a namespace](#).
- You are connected to a master node of the cluster. For more information, see [Connect to a Kubernetes cluster through kubectl](#).

Context

By default, a running pod uses the CPU and memory resources of a node without limit. This means that pods can compete for computing resources of a cluster. As a result, the pods in a namespace may exhaust all of the computing resources.

Namespaces can be used as virtual clusters to serve multiple purposes. We recommend that you set resource quotas for namespaces.

For a namespace, you can set quotas on resources such as CPU, memory, and pod quantity. For more information, see [Resource Quotas](#).

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, click **Namespaces and Quotas**.
5. Find the namespace that you want to manage and click **Resource Quotas and Limits** in the **Actions** column.
6. In the dialog box that appears, set resource quotas and default resource limits.

Note After you set CPU and memory quotas for a namespace, you must specify CPU and memory limits when you create a pod. You can also set the default resource limits for the namespace. For more information, see [Resource Quotas](#).

- i. Set resource quotas for the namespace.

The screenshot shows the 'Resource Quotas and Limits' dialog box with the 'Resource Quota' tab selected. It is divided into three sections: 'Compute Resource Quota', 'Storage Resource Quota', and 'Other Limits'. Each section contains a list of resources with checkboxes and input fields for their total limits.

| Resource | Limit |
|------------------------|---------|
| CPU Limit | 2 Cores |
| Memory Limit | 4Gi |
| Storage Capacity | 1024Gi |
| PVCs | 50 |
| ConfigMaps | 100 |
| Pods | 50 |
| Services | 20 |
| Load Balancer Services | 5 |
| Secrets | 10 |

- ii. You can set resource limits and resource requests for containers in the namespace. This enables you to control the amount of resources consumed by the containers. For more information, see <https://kubernetes.io/memory-default-namespace/>.

The screenshot shows the 'Resource Quotas and Limits' dialog box with the 'LimitRange' tab selected. It displays a table for setting limits and requests for CPU and Memory.

| | CPU | Memory |
|---------|-----------|--------|
| Limit | 0.5 Cores | 512Mi |
| Request | 0.1 Cores | 256Mi |

- 7. After you set resource quotas and limits, connect to a master node of the cluster and run the following commands to query the resource configurations of the namespace.

```
# kubectl get limitrange,ResourceQuota -n test
NAME AGE
limitrange/limits 8m
NAME AGE
resourcequota/quota 8m
# kubectl describe limitrange/limits resourcequota/quota -n test
Name: limits
Namespace: test
Type Resource Min Max Default Request Default Limit Max Limit/Request Ratio
-----
Container cpu -- 100m 500m -
Container memory -- 256Mi 512Mi -
Name: quota
Namespace: test
Resource Used Hard
-----
configmaps 0 100
limits.cpu 0 2
limits.memory 0 4Gi
persistentvolumeclaims 0 50
pods 0 50
requests.storage 0 1Ti
secrets 1 10
services 0 20
services.loadbalancers 0 5
```

3.6.6.3. Modify a namespace

You can modify an existing namespace.

Prerequisites

- A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).
- A sample namespace named `test` is created. For more information, see [Create a namespace](#).

Context

When you modify a namespace, you can add, delete, or modify the labels that are added to the namespace based on your requirements.

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, click **Namespaces and Quotas**.
5. Find the namespace that you want to modify and click **Edit** in the **Actions** column.
6. In the **Edit Namespace** dialog box, find the label that you want to modify and click **Edit** to modify the label. For example, you can change the key-value pair of the label to `env:test-V2`. Then, click **Save**.
7. Click **OK**.

You can find that the labels of the namespace on the Namespace page are updated.

3.6.6.4. Delete a namespace

You can delete namespaces that are no longer in use.

Prerequisites

- A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).
- A sample namespace named `test` is created. For more information, see [Create a namespace](#).

Context

 **Note** When you delete a namespace, all resource objects under the namespace are deleted. Proceed with caution.

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, click **Namespaces and Quotas**.
5. Find the namespace that you want to delete and click **Delete** in the **Actions** column.
6. In the message that appears, click **Confirm**.

Return to the Namespace page. You can find that the namespace is deleted. Resource objects in the namespace are also deleted.

3.6.7. Applications

3.6.7.1. Create an application from an image

You can use an image to create an NGINX application that is accessible over the Internet.

Prerequisites

- A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).
- Your Kubernetes cluster is accessible over the Internet.

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Workloads > Deployments**.
5. On the **Deployments** page, select the namespace and click **Create from Image** in the upper-right corner.
6. On the page that appears, set the following parameters: **Name**, **Cluster**, **Namespace**, **Replicas**, and **Type**. Then, select **Synchronize Timezone**. Click **Next**.

If you do not set the **Namespace** parameter, the default namespace is used.

The screenshot shows the 'Create Application' wizard with the 'Basic Information' step selected. The form contains the following fields:

- Name:** A text input field containing 'nginx'. Below it is a note: 'The name must be 1 to 64 characters in length and can contain numbers lower case letters and hyphens (-). The name cannot start with a hyphen (-).'.
- Cluster:** A dropdown menu with 'k8s-cluster01' selected.
- Namespace:** A dropdown menu with 'default' selected.
- Replicas:** A text input field containing '2'.
- Type:** A dropdown menu with 'Deployment' selected.

At the bottom right, there are 'Back' and 'Next' buttons.

7. Configure containers.

Note You can configure a pod that contains one or more containers for the application.

i. Configure basic settings.

Basic settings of the container

| Parameter | Description |
|----------------------------------|---|
| Image Name | You can click Select Image to select an image in the dialog box that appears and click OK . In this example, NGINX is selected. You can also specify an image in a private registry. Use the following format to specify an image in a private registry: <code>domainname/namespace/imagename</code> . |
| Image Version | You can click Select Image Version to select a version. If you do not specify the image version, the latest version is used. |
| Always | Container Service caches images to improve the efficiency of application deployment. When Container Service deploys the application, if the specified image version is the same as the cached image version, the cached image is used. Therefore, when you update the application code, if you do not change the image version for reasons such as to support the upper-layer workloads, the cached image is used to deploy the application. After you select the Always option, Container Service always pulls the latest image from the repository. This ensures that the latest image and code are used to deploy the application. |
| Set Image Pull Secret | Click Set Image Pull Secret to set the image Secret. The Secret is required if you need to access a private repository. |
| Resource Limit | The upper limits of CPU and memory resources that can be used by this application. This prevents the application from occupying excessive resources. The unit of CPU resources is Core. The unit of memory is MiB. |
| Required Resources | The amount of CPU and memory resources that are reserved for this application. These resources are exclusive to the container. This prevents other services or processes from occupying the resources of the application. |
| Container Start Parameter | stdin: Pass stdin to the container. tty: Stdin is a TTY. Typically, these two check boxes are both selected. This allows you to associate the terminal (tty) with the standard inputs (stdin) of the container. For example, an interactive program can obtain standard inputs from users and then display the inputs on the terminal. |
| Init Container | When this check box is selected, the system creates an Init Container that contains useful tools. For more information, see . |

ii. (Optional)Set environment variables.

You can set environment variables in key-value pairs for pods. Environment variables are used to apply pod configurations to containers. For more information, see [Pod variable](#).

iii. (Optional)Configure the **Health Check** settings.

Health check settings include liveness and readiness probes. Liveness probes determine when to restart the container. Readiness probes indicate whether the container is ready to accept network traffic. For more information about health checks, see .

The screenshot displays the 'Health Check' configuration panel, which is divided into two sections: 'Liveness' and 'Readiness'. Both sections have an 'Enable' checkbox checked. Each section features a dropdown menu for 'Request type' (set to 'HTTP Request'), a dropdown for 'Protocol' (set to 'HTTP'), and a dropdown for 'Type' (set to 'TCP'). Below these are input fields for 'Path', 'Port', and 'Command'. An 'HTTP Header' section contains 'name' and 'value' input fields. The 'Liveness' section includes fields for 'Initial Delay (s)' (3), 'Period (s)' (10), 'Timeout (s)' (1), 'Success Threshold' (1), and 'Failure Threshold' (3). The 'Readiness' section includes fields for 'Initial Delay (s)' (3), 'Period (s)' (10), 'Timeout (s)' (1), 'Success Threshold' (1), and 'Failure Threshold' (3).

| Request type | Description |
|--------------|-------------|
|--------------|-------------|

| Request type | Description |
|--------------|---|
| HTTP | <p>Sends an HTTP GET request to the container. Supported parameters:</p> <ul style="list-style-type: none"> ▪ Protocol: HTTP or HTTPS. ▪ Path: the requested path on the server. ▪ Port: the container port that you want to open. Enter a port number from 1 to 65535. ▪ HTTP Header: the custom headers in the HTTP request. Duplicate headers are allowed. Key-value pairs are supported. ▪ Initial Delay (s): The initialDelaySeconds field. The time (in seconds) to wait before performing the first probe after the container is started. Default value: 3. ▪ Period (s): the periodSeconds field. The intervals (in seconds) between two adjacent probes. Default value: 10. Minimum value: 1. ▪ Timeout (s): the timeoutSeconds field. The time (in seconds) after which a probe times out. Default value: 1. Minimum value: 1. ▪ Healthy Threshold: the minimum number of times that an unhealthy container must consecutively pass health checks before it is considered healthy. Default value: 1. Minimum value: 1. For liveness probes, this parameter must be set to 1. ▪ Unhealthy Threshold: the minimum number of times that a healthy container must consecutively fail health checks before it is considered unhealthy. Default value: 3. Minimum value: 1. |
| TCP | <p>Sends a TCP socket to the container. Kubelet attempts to open the socket on the specified port. If the connection can be established, the container is considered healthy. Otherwise, the container is considered unhealthy. You can set the following parameters:</p> <ul style="list-style-type: none"> ▪ Port: the container port that you want to open. Enter a port number from 1 to 65535. ▪ Initial Delay (s): the initialDelaySeconds field. The time (in seconds) to wait before performing the first probe after the container is started. Default value: 15. ▪ Period (s): the periodSeconds field. The intervals (in seconds) between two adjacent probes. Default value: 10. Minimum value: 1. ▪ Timeout (s): the timeoutSeconds field. The time (in seconds) after which a probe times out. Default value: 1. Minimum value: 1. ▪ Healthy Threshold: the minimum number of times that an unhealthy container must consecutively pass health checks before it is considered healthy. Default value: 1. Minimum value: 1. For liveness probes, this parameter must be set to 1. ▪ Unhealthy Threshold: the minimum number of times that a healthy container must consecutively fail health checks before it is considered unhealthy. Default value: 3. Minimum value: 1. |

| Request type | Description |
|--------------|--|
| Command | <p>Runs a probe command in the container to check the health status of the container. You can set the following parameters:</p> <ul style="list-style-type: none"> ■ Command: the probe command that is run to check the health status of the container. ■ Initial Delay (s): the <code>initialDelaySeconds</code> field. The time (in seconds) to wait before performing the first probe after the container is started. Default value: 5. ■ Period (s): the <code>periodSeconds</code> field. The intervals (in seconds) between two adjacent probes. Default value: 10. Minimum value: 1. ■ Timeout (s): the <code>timeoutSeconds</code> field. The time (in seconds) after which a probe times out. Default value: 1. Minimum value: 1. ■ Healthy Threshold: the minimum number of times that an unhealthy container must consecutively pass health checks before it is considered healthy. Default value: 1. Minimum value: 1. For liveness probes, this parameter must be set to 1. ■ Unhealthy Threshold: the minimum number of times that a healthy container must consecutively fail health checks before it is considered unhealthy. Default value: 3. Minimum value: 1. |

iv. Configure the lifecycle of the container.

You can set the following parameters to configure the lifecycle of the container in the Start, Post Start, and Pre Stop fields. For more information, see <https://kubernetes.io/docs/tasks/configure-pod-container/attach-handler-lifecycle-event/>.

- **Start:** the pre-start command and parameter.
- **Post Start:** the post-start command.
- **Pre Stop:** the pre-stop command.

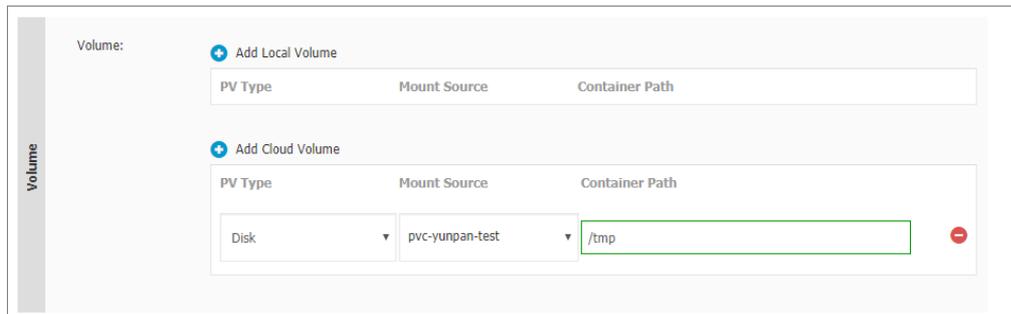
| | | |
|-----------|-------------|---|
| Lifecycle | Start: | Command <input]<="" td="" type="text" value='["/bin/sh", "-c", "echo Hello > /user/share/message"]'/> |
| | | Parameter <input type="text"/> |
| | Post Start: | Command <input type="text"/> |
| | Pre Stop: | Command <input]<="" td="" type="text" value='["/user/sbin/nginx", "-s", "quit"]'/> |

v. (Optional)Configure volumes.

Local storage and cloud storage are supported.

- **Local Storage:** supports hostPath, ConfigMap, Secret, and emptyDir. The specified volume is mounted to a path in the container. For more information, see [Volumes](#).
- **Cloud Storage:** supports disks, Apsara File Storage NAS (NAS) volumes, and Object Storage Service (OSS) volumes.

In this example, a persistent volume (PV) is created from a disk and mounted to the `/tmp` path in the container. Data generated in this path is stored in the disk.



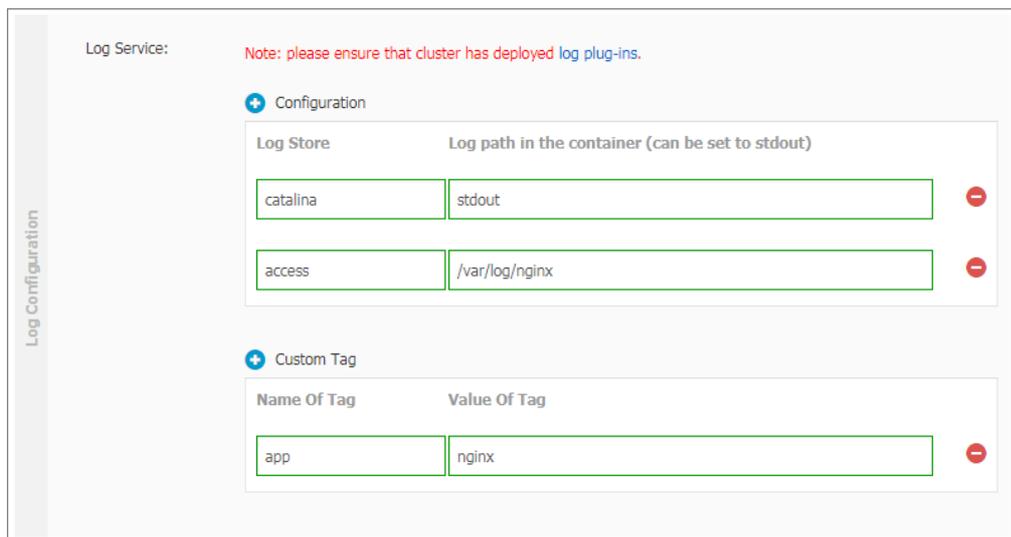
vi. (Optional)Configure **Log Service**. You can specify log collection configurations and custom tags.

? **Note** Make sure that the Log Service agent is installed in the cluster.

You can configure the following log collection settings:

- **Logstore:** Create a Logstore in Log Service to store the collected log data.
- **Log Path in Container:** Specify stdout or a path to collect log data.
 - **stdout:** specifies that the stdout files are collected.
 - **Text Logs:** specifies that the log data in the specified path of the container is collected. In this example, the log data in the `/var/log/nginx` path is collected. Wildcard characters can be used in the path.

You can also add custom tags. Tags are added to the log data of the container when the log data is collected. Log data with tags is easier to aggregate and filter.



8. After you complete the configurations of the container, click **Next**.
9. Configure advanced settings. Configure the **Access Control** settings.

You can configure the method to expose pods and click **Create**. In this example, a ClusterIP Service and an Ingress are created to enable Internet access to the NGINX application.

Note

You can configure the following access control settings based on your business requirements:

- Internal applications: For applications that run inside the cluster, you can create a ClusterIP Service or a NodePort Service to enable internal communication.
- External applications: For applications that need to be accessed over the Internet, you can configure access control settings by using one of the following methods:
 - Create a LoadBalancer Service: When you create a LoadBalancer Service, a Server Load Balancer (SLB) instance is associated with the Service and is created to expose applications to the Internet.
 - Create a ClusterIP Service or a NodePort Service, and then create an Ingress: When you use this method, an Ingress is created to expose applications to the Internet. For more information, see .

i. To create **Services**, click **Create** in the Access Control section. In the dialog box that appears, set the parameters, and then click **Create**.

| Parameter | Description |
|-------------|---|
| Name | Enter a name for the Service. Default value: <code>applicationname-svc</code> . |

| Parameter | Description |
|---------------------|--|
| Type | <p>Select one from the following three Service types:</p> <ul style="list-style-type: none"> ■ Cluster IP: creates a ClusterIP Service. This type of Service exposes applications through an internal IP address of the cluster. If you select this type, applications can be accessed only within the cluster. ■ Node Port: creates a NodePort Service. This type of Service exposes applications through the IP address and static port on each node. A NodePort Service can be used to route requests to a ClusterIP Service. The system automatically creates the ClusterIP Service. You can access a NodePort Service from outside the cluster by requesting <code><NodeIP>:<NodePort></code> . ■ Server Load Balancer: creates a LoadBalancer Service. This type of Service exposes applications through an SLB instance, which supports Internet access or internal access. The SLB instance can route requests to NodePort and ClusterIP Services. |
| Port Mapping | Specify a Service port and a container port. If the Type parameter is set to Node Port, you must specify a node port to avoid port conflicts. TCP and UDP protocols are supported. |
| Annotations | Add annotations to the Service. SLB parameters are supported. For more information, see Access services by using SLB . |
| Label | Add one or more labels to the Service. The labels are used to identify the Service. |

- ii. To create **Ingresses**, click **Create** in the Access Control section. Configure Ingress rules in the dialog box that appears, and then click **Create**. For more information about Ingress configuration, see [Ingress configurations](#).

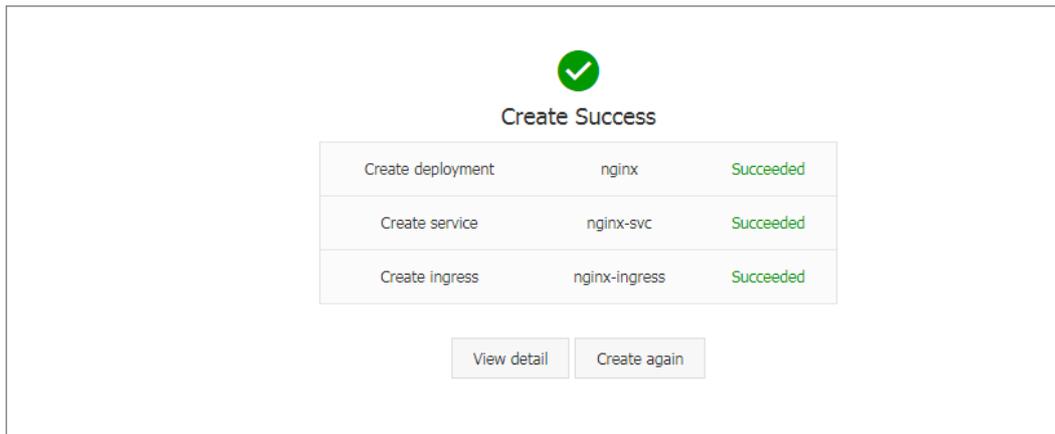
When you create an application from an image, you can create an Ingress for only one Service. In this example, a virtual hostname is specified as the test domain name. You must add a mapping rule to the hosts file for this domain name, as shown in the following code block. In actual scenarios, use a domain name that has obtained an Internet Content Provider (ICP) number.

```
101.37.224.146 foo.bar.com #The IP address of the Ingress.
```

- iii. You can find the newly created Service and Ingress in the Access Control section. You can click **Update** or **Delete** to make changes.

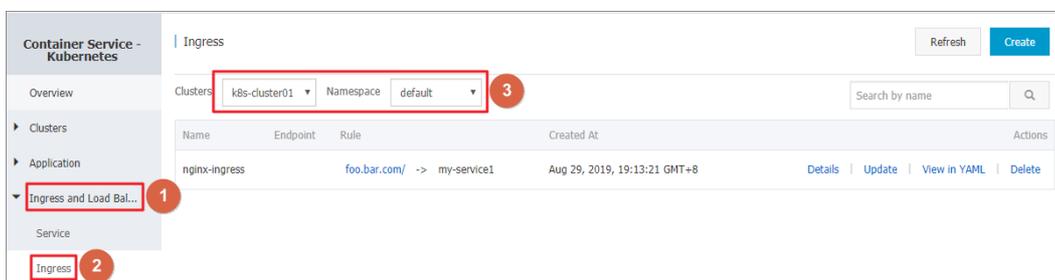
| Create Application | | | | |
|--------------------|------------------------|------------------------|------------------------|------|
| Basic Information | | Container | Advanced | Done |
| Access Control | Services(Service) | Update | Delete | |
| | service port | Container Port | Protocol | |
| | 8080 | 8080 | TCP | |
| Ingresses(Ingress) | Update | Delete | | |
| Domain | path | Name | service port | |
| foo.bar.com | | nginx-svc | 8080 | |

- 10. Click **Create**.
- 11. After the application is created, you are redirected to the Complete page, which displays the resource objects under the application. You can click **View Details** to view application details.



The nginx-deployment details page appears.

- In the left-side navigation pane, choose **Services and Ingresses > Ingresses**. On the Ingresses page, you can find the created Ingress.



- Enter the test domain in the address bar of your browser and press Enter. The NGINX welcome page appears.



3.6.7.2. Create an application from an orchestration template

Container Service provides orchestration templates that you can use to create applications. You can also modify the templates based on YAML syntax to customize applications.

Prerequisites

A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).

Context

This topic describes how to use an orchestration template to create an NGINX application that consists of a Deployment and a Service. The Deployment provisions pods for the application and the Service manages access to the pods at the backend.

Procedure

- Log on to the [Container Service console](#).
- In the left-side navigation pane, click **Clusters**.

3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Workloads > Deployments**.
5. On the **Deployments** page, select the namespace and click **Create from YAML** in the upper-right corner.
6. Configure the template and click **Create**.
 - o **Sample Template**: Container Service provides YAML templates of various resource types to help you quickly deploy resource objects. You can also create a custom template based on YAML syntax to describe the resource that you want to define.
 - o **Add Deployment**: This feature allows you to quickly define a YAML template.
 - o **Use Existing Template**: You can import an existing template to the configuration page.

Based on an orchestration template provided by Container Service, the following sample template is used to create a Deployment of an NGINX application.

 **Note** Container Service supports YAML syntax. You can use the `---` symbol to separate multiple resource objects. This allows you to create multiple resource objects in a single template.

```
apiVersion: apps/v1beta2 # for versions before 1.8.0 use apps/v1beta1
kind: Deployment
metadata:
  name: nginx-deployment
  labels:
    app: nginx
spec:
  replicas: 2
  selector:
    matchLabels:
      app: nginx
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
        - name: nginx
          image: nginx:1.7.9 # replace it with your exactly <image_name:tags>
          ports:
            - containerPort: 80
---
apiVersion: v1 # for versions before 1.8.0 use apps/v1beta1
kind: Service
metadata:
  name: my-service1 #TODO: to specify your service name
  labels:
    app: nginx
spec:
  selector:
    app: nginx #TODO: change label selector to match your backend pod
  ports:
    - protocol: TCP
      name: http
      port: 30080 #TODO: choose an unique port on each node to avoid port conflict
      targetPort: 80
  type: LoadBalancer ##In this example, the type is changed from Nodeport to LoadBalancer.
```

7. Click **Create**. A message that indicates the deployment status appears.

After the application is created, choose **Services and Ingresses > Services** in the left-side navigation pane. On the Services page, you can find that a Service named my-service1 is created for the application. The external endpoint of the Service is also displayed on the page. Click the endpoint in the **External Endpoint** column.

8. You can visit the NGINX welcome page in the browser.



3.6.7.3. Use commands to manage applications

You can use commands to create applications or view application containers.

Prerequisites

Before you use commands on your local host, you have connected to a Kubernetes cluster through kubectl. For more information, see [Connect to a Kubernetes cluster through kubectl](#).

Run a command to create an application

You can use the following command to run a simple container (an NGINX Web server in this example):

```
# kubectl run -it nginx --image=registry.aliyuncs.com/spacexnice/netdia:latest
```

This command creates a service portal for this container. After you specify `--type=LoadBalancer`, an SLB route to the NGINX container is created.

```
# kubectl expose deployment nginx --port=80 --target-port=80 --type=LoadBalancer
```

Run a command to view container information

Run the following command to list all running containers in the default namespace:

```
root@master # kubectl get pods
NAME                READY  STATUS   RESTARTS  AGE
nginx-2721357637-dvwq3  1/1    Running  1         9h
```

3.6.7.4. Create a Service

You can create a Service for your application in the Container Service console. The Service provides access to the application.

In Kubernetes, a Service is an abstraction which defines a logical set of pods and a policy by which to access the pods. This pattern is also known as a microservice. A label selector determines whether the set of pods can be accessed by the Service.

Each pod in Kubernetes clusters has its own IP address. However, pods are frequently created and deleted. Therefore, if you directly expose pods to external access, high availability is not ensured. Services decouple the frontend from the backend. The frontend clients do not need to be aware of which backend pods are used. This provides a loosely-decoupled microservice architecture.

For more information, see [Kubernetes Services](#).

Prerequisites

A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).

Step 1: Create a Deployment

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Workloads > Deployments**.
5. On the **Deployments** page, select the namespace and click **Create from Template** in the upper-right corner.
6. Select a sample template or customize a template, and click **Create**.

In this example, the template of an NGINX Deployment is used.

```
apiVersion: apps/v1 # for versions before 1.8.0 use apps/v1beta1 kind: Deployment metadata: name: nginx-deployment-basic labels: app: nginx spec: replicas: 2 selector: matchLabels: app: nginx template: metadata: labels: app: nginx spec: # nodeSelector: # env: test-team containers: - name: nginx image: nginx:1.7.9 # replace it with your exactly <image_name:tags> ports: - containerPort: 80
```

Query the state of the Deployment.

Step 2: Create a Service

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Services and Ingresses > Services**.
5. On the Services page, select the **Namespace** and click **Create** in the upper-right corner.
6. On the Create Service dialog box, set the parameters.
 - **Name**: Enter a name for the Service. In this example, nginx-svc is used.
 - **Type**: Select the type of the Service. This parameter specifies how to access the Service. Valid values:
 - **Cluster IP**: a ClusterIP Service. This type of Service exposes pods through an internal IP address of the cluster. If you select this option, pods can be accessed only within the cluster. This is the default value.
 - **Node Port**: a NodePort Service. This type of Service exposes pods by using the IP address and a static port of each node. A NodePort Service can be used to route requests to a ClusterIP Service, which is automatically created by the system. You can access a NodePort Service from outside the cluster by sending requests to `<NodeIP>:<NodePort>`.
 - **Server Load Balancer**: a LoadBalancer Service. This type of Service exposes pods by using Server Load Balancer (SLB) instances, which support Internet access or internal access. SLB instances can be used to route requests to NodePort and ClusterIP Services.
 - **Backend**: the backend object that you want to associate with the Service. In this example, select nginx-deployment-basic that was created in the preceding step. If you do not associate the Service with a backend object, no Endpoint object is created. You can also manually associate the Service with an Endpoint object. For more information, see [Create a Service without selectors](#).
 - **Port Mapping**: Set the Service port and container port. The container port must be the same as the one that is exposed in the backend pod.
 - **Annotations**: Add one or more annotations to the Service to configure SLB parameters. For example, set the name to `service.beta.kubernetes.io` and the value to `20`. This means that the maximum bandwidth of

the Service is 20 Mbit/s. For more information, see [Access services by using SLB](#).

- **Label:** Add one or more labels to the Service. The labels are used to identify the Service.
7. Click **Create**. After the nginx-svc Service is created, it appears on the Services page.
 8. On the **Services** page, you can view basic information about the Service. You can also access its external endpoint by using a browser.

3.6.7.5. View a Service

You can view details about a Service in the Container Service console.

Context

Each pod in Kubernetes clusters has its own IP address. However, pods are frequently created and deleted. Therefore, it is not practical to directly expose pods to external access. Services decouple the frontend from the backend, which provides a loosely-coupled microservice architecture.

Procedure

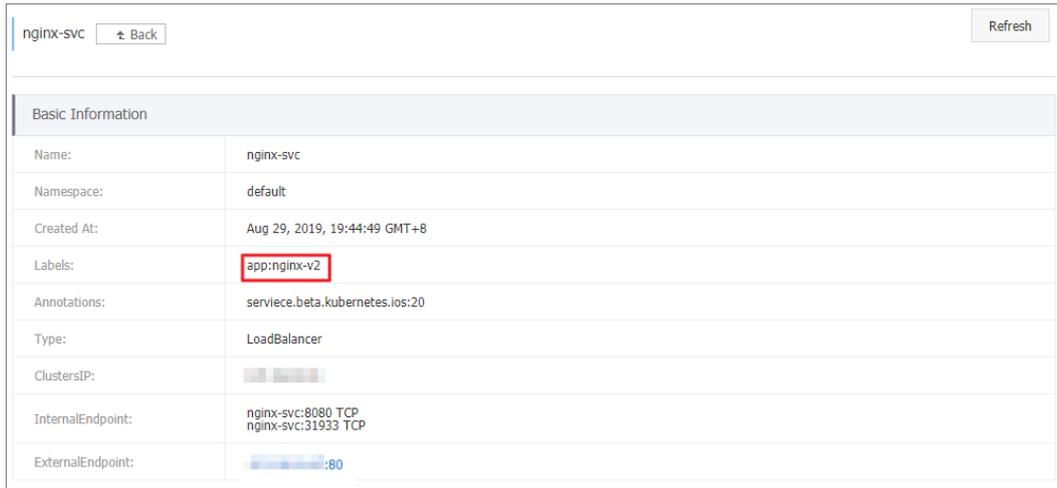
1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane, choose **Services and Ingresses > Services**.
5. On the **Services** page, select the namespace, find the Service that you want to view, and then click **Details** in the **Actions** column.
On the details page, you can view detailed information about the Service.

3.6.7.6. Update a Service

You can update a Service through the Container Service console. This topic describes how to update a Service.

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Services and Ingresses > Services**.
5. On the **Services** page, select the namespace, find the Service that you want to update, and then click **Update** in the **Actions** column.
6. In the **Update Service** dialog box, modify the configurations and click **Update**.
7. In the Service list, find the Service that you updated and click **Details** in the **Actions** column to view configuration changes. In this example, the labels of the Service are modified.



The screenshot shows the details of a Service named 'nginx-svc' in the Container Service console. The service is located in the 'default' namespace and was created on August 29, 2019, at 19:44:49 GMT+8. It has a label 'app:nginx-v2' highlighted with a red box. The service type is 'LoadBalancer'. The internal endpoints are 'nginx-svc:8080 TCP' and 'nginx-svc:31933 TCP'. The external endpoint is '...:80'. There are 'Back' and 'Refresh' buttons at the top of the page.

| Basic Information | |
|-------------------|---|
| Name: | nginx-svc |
| Namespace: | default |
| Created At: | Aug 29, 2019, 19:44:49 GMT+8 |
| Labels: | app:nginx-v2 |
| Annotations: | service.beta.kubernetes.io:20 |
| Type: | LoadBalancer |
| ClustersIP: | ... |
| InternalEndpoint: | nginx-svc:8080 TCP nginx-svc:31933 TCP |
| ExternalEndpoint: | ...:80 |

3.6.7.7. Delete a Service

This topic describes how to delete a Service.

Prerequisites

- A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).
- A Service is created. For more information, see [Create Services](#).

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Services and Ingresses > Services**.
5. On the Services page, select the namespace, find the Service that you want to delete, and then click **Delete** in the **Actions** column.
6. In the message that appears, click **Confirm**.

3.6.7.8. Use a trigger to redeploy an application

You can create a trigger and use it to redeploy an application. This topic describes how to use a trigger.

Prerequisites

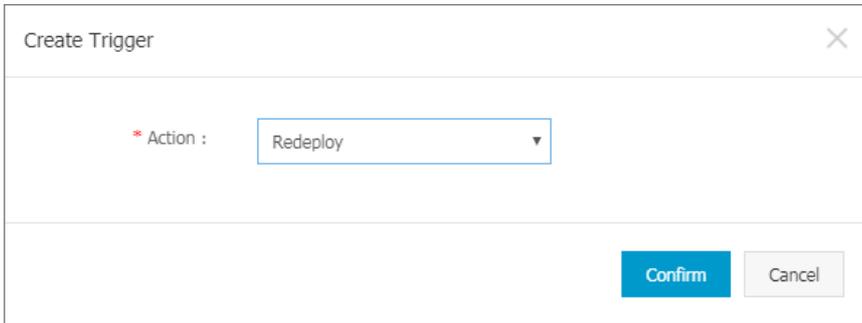
- A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).
- An application is created. Then, a trigger is created for the application and the application is used to test the trigger. In this example, an NGINX application is created.

Procedure

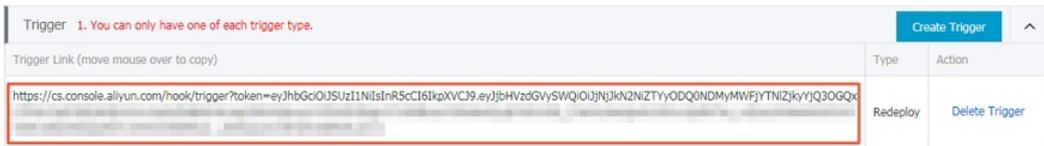
1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Manage** in the **Actions** column of the cluster.
4. In the left-side navigation pane, click **Deployments**. On the **Deployments** page, find the NGINX application and click **Details** in the **Actions** column.
5. On the details page of the application, click the **Triggers** tab. Then, click **Create Trigger**.

6. In the dialog box that appears, set Action to **Redeploy** and click **Confirm**.

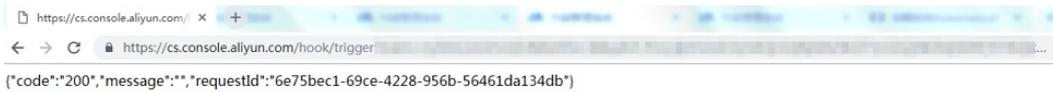
Note You can create triggers only to redeploy applications.



After the trigger is created, the webhook URL of the trigger appears in the Trigger Link Address column on the **nginx - deployment** page.



7. Copy the webhook URL, enter it into the address bar of your browser, and press Enter. A message that indicates specific information such as the request ID appears.



8. Go to the **nginx - deployment** page. A new pod appears on the **Pods** tab.



After the new pod is deployed, the original pod is automatically deleted.

What's next

You can run triggers by sending GET or POST requests from an external system. For example, you can use the curl command-line tool to run triggers.

To run the redeploy trigger, run the following command:

```
curl https://cs.console.aliyun.com/hook/trigger?token=xxxxxxx
```

3.6.7.9. View pods

You can view pods in the Container Service console.

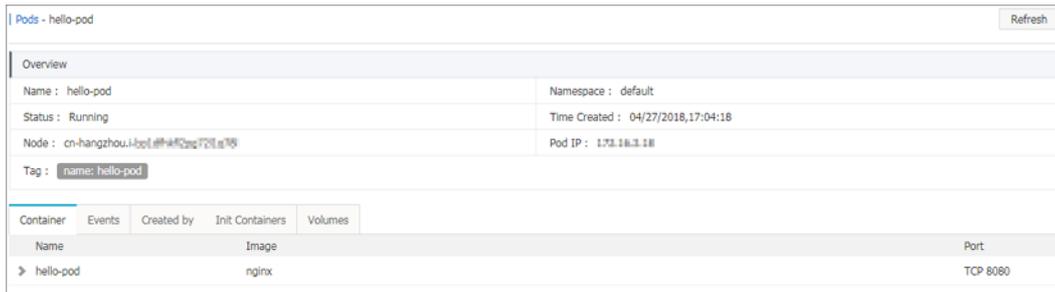
Procedure

1. Log on to the Container Service console.
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.

4. In the left-side navigation pane of the details page, choose **Workloads > Pods**.
5. In the left-side navigation pane, choose **Applications > Pods**. The Pods page appears.
6. On the **Pods** page, select the namespace, find the pod that you want to view, and then click **View Details** in the Actions column.

Note You can update or delete pods on the Pods page. We recommend that you use Deployments to manage pods if the pods are created by Deployments.

On the details page of the pod, you can view detailed information about the pod.



3.6.7.10. Manage pods

Pods are the smallest deployable units in Kubernetes. A pod runs an instance of an independent application in Kubernetes. Each pod contains one or more containers that are tightly coupled. You can modify pods, view pods, and manually scale the number of pods for an application in the Container Service console.

View pods

View pod details

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click the name or click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Workloads > Pods**.
5. Find the pod that you want to view and click **View Details** in the Actions column.

You can view details of pods by using one of the following methods:

- o Method 1: In the left-side navigation pane of the details page, choose **Workloads > Deployments**. Find the application that you want to manage and click the name of the application. On the **Pods** tab, click the name of the pod to view details.
- o Method 2: In the left-side navigation pane of the details page, choose **Services and Ingresses > Services**. Click the name of the Service that you want to manage. On the page that appears, find and click the name of the application that you want to manage. On the **Pods** tab, click the name of the pod to view details.

Note On the Pods page, you can modify and delete pods. For pods that are created by using a Deployment, we recommend that you use the Deployment to manage the pods.

View pod log

You can view a pod log by using the following methods:

Navigate to the Pods tab, find the pod that you want to manage, and then click **Logs** on the right side to view the log data.

Modify pod configurations

1. Log on to the [Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click the name or click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Workloads > Pods**.
5. On the **Pods** page, find the pod that you want to manage and click **Edit** in the **Actions** column.
6. In the dialog box that appears, modify the configuration of the pod and click **Update**.

Edit YAML

✕

```
1  apiVersion: v1
2  kind: Pod
3  metadata:
4    annotations:
5      kubernetes.io/psp: ack.privileged
6    creationTimestamp: '2021-04-27T06:18:19Z'
7    generateName: ack-node-problem-detector-daemonset-
8    labels:
9      app: ack-node-problem-detector
10     controller-revision-hash: 67db699848
11     pod-template-generation: '2'
12  managedFields:
13    - apiVersion: v1
14      fieldType: FieldsV1
15      fieldsV1:
16        'f:metadata': {}
17        'f:generateName': {}
18        'f:labels': {}
19        'f:app': {}
20        'f:controller-revision-hash': {}
21        'f:pod-template-generation': {}
22        'f:ownerReferences': {}
23        'f:kubernetes.io/psp': {}
24        'f:uid': {}
25        'f:apiVersion': {}
26        'f:blockOwnerDeletion': {}
27        'f:controller': {}
28        'f:kind': {}
29        'f:name': {}
30        'f:uid': {}
31        'f:spec': {}
32        'f:affinity': {}
33        'f:nodeAffinity': {}
34        'f:requiredDuringSchedulingIgnoredDuringExecution': {}
```

[Update](#) [Download](#) [Save As](#) [Cancel](#)

Manually scale the number of pods for an application

After an application is created, you can scale in or out the pods that are provisioned for the application.

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click the name or click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Workloads > Deployments**.
5. Select the namespace where the Deployment is deployed, find the Deployment, and then click **Scale** in the **Actions** column.
6. In the dialog box that appears, set **Desired Number of Pods** to 4 and click **OK**.

Note By default, the resources created by a Deployment are updated based on the rollingUpdate strategy. This ensures that a minimum number of pods are running during the update. You can specify the minimum number of pods that must run in Pod Template of the YAML file.

3.6.7.11. Schedule pods to specific nodes

You can add labels to nodes and schedule pods to the nodes with specified labels. This topic describes how to schedule a pod to a node with specified labels.

You can add labels to nodes and then configure `nodeSelector` to schedule pods to nodes with specified labels. For more information about how `nodeSelector` works, see [nodeSelector](#).

To meet business requirements, you may want to deploy a controller service on a master node, or deploy services on nodes that use standard SSDs. You can use the following method to schedule pods to nodes with specified labels.

Prerequisites

A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).

Step 1: Add a label to a node

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Nodes > Nodes**.
5. On the **Nodes** page, click **Manage Labels and Taints** in the upper-right corner.
6. Select one or more nodes and then click **Add Label**. In this example, a worker node is selected.
7. In the dialog box that appears, enter the name and value of the label and click **OK**.

On the Labels tab, you can find the `group:worker` label next to the selected node.

You can also run the following command to add a label to a node: `kubectl label nodes <node-name> <label-key>=<label-value>` .

Step 2: Schedule a pod to the node

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.

3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Workloads > Deployments**.
5. On the **Deployments** page, select the namespace and click **Create from Template** in the upper-right corner.
6. On the **Create** page, select a template from the Sample Template drop-down list. In this example, a custom template is selected. Copy the following content to the custom template and click **Create**.

The following template is used as an example:

```
apiVersion: v1
kind: Pod
metadata:
  labels:
    name: hello-pod
spec:
  containers:
    - image: nginx
      imagePullPolicy: IfNotPresent
      name: hello-pod
      ports:
        - containerPort: 8080
          protocol: TCP
  resources: {}
  securityContext:
    capabilities: {}
    privileged: false
    terminationMessagePath: /dev/termination-log
  dnsPolicy: ClusterFirst
  restartPolicy: Always
  nodeSelector:
    group: worker          ## This value must be the same as the node label that is added in Step 1.
  status: {}
```

3.6.7.12. Simplify application deployment by using Helm

This topic introduces the basic terms and components of Helm and describes how to use Helm to deploy the sample applications WordPress and Spark in a Kubernetes cluster.

Prerequisites

- A Kubernetes cluster is created in the Container Service console. For more information, see [Create a Kubernetes cluster](#).

Tiller is automatically deployed to the cluster when the Kubernetes cluster is created. The Helm CLI is automatically installed on each master node. An Alibaba Cloud chart repository is added to Helm.

- A Kubernetes version that supports Helm is used.

Only Kubernetes 1.8.4 and later support Helm. If the Kubernetes version of your cluster is 1.8.1, you can **upgrade** the cluster on the Clusters page of the Container Service console.

Context

Application management is the most challenging task in Kubernetes. The Helm project provides a unified method to package software and manage software versions. You can use Helm to simplify application distribution and deployment. App Catalog is integrated with Helm in the Container Service console and provides extended features based on Helm. App Catalog also supports Alibaba Cloud chart repositories to help you accelerate application deployments. You can deploy applications in the Container Service console or by using the Helm CLI.

This topic introduces the basic terms and components of Helm and describes how to use Helm to deploy the sample applications WordPress and Spark in a Kubernetes cluster.

Basic terms

Helm is an open source project initiated by Deis. Helm can be used to simplify the deployment and management of Kubernetes applications.

Helm serves as a package manager for Kubernetes and allows you to find, share, and use applications built by Kubernetes. Before you use Helm, you must familiarize yourself with the following basic terms:

- **Chart:** a packaging format used by Helm. Each chart contains the images, dependencies, and resource definitions that are required to run an application. A chart may contain service definitions in a Kubernetes cluster. A Helm chart is similar to a Homebrew formula, an Advanced Package Tool (APT) dpkg, or a Yum rpm.
- **Release:** an instance of a chart that runs in a Kubernetes cluster. A chart can be installed multiple times into a Kubernetes cluster. After a chart is installed, a new release is created. For example, you can install a MySQL chart. If you want to run two databases in your cluster, you can install the MySQL chart twice. Each time a chart is installed, a release is created with a different name.
- **Repository:** the storage of charts. Charts are published and stored in repositories.

Helm components

Helm uses a client-server architecture and consists of the following components:

- The Helm CLI is the Helm client that runs on your on-premises machine or on the master nodes of a Kubernetes cluster.
- Tiller is the server-side component and runs in a Kubernetes cluster. Tiller manages the lifecycles of Kubernetes applications.
- A repository is used to store charts. The Helm client can access the index file and packaged charts in a chart repository over HTTP.

Deploy an application in the Container Service console

1. [Log on to the Container Service console](#)
2. In the left-side navigation pane, choose **Marketplace > App Catalog**.
3. On the App Catalog page, find and click the WordPress chart to go to the details page of the chart.
4. Click the **Parameters** tab and modify the configurations.

In this example, a dynamically provisioned disk volume is associated with a persistent volume claim (PVC). For more information, see [Use Apsara Stack disks](#).

 **Note** You must first provision a disk as a persistent volume (PV). The capacity of the PV cannot be less than the capacity specified in the PVC.

5. In the **Deploy** section, select the cluster in which you want to deploy the application and click **Create**. After the application is deployed, you are redirected to the release page of the application.
6. In the left-side navigation pane, click **Clusters**.
7. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
8. In the left-side navigation pane of the details page, choose **Services and Ingresses > Services**. On the Services page, select the namespace, find the Service that is created for the application, and then check its HTTP and HTTPS external endpoints.
9. Click one of the endpoints to go to the WordPress website where you can publish blog posts.

Deploy an application by using the Helm CLI

After the Helm CLI is automatically installed on the master nodes of the Kubernetes cluster and the required chart repository is added to Helm, you can log on to the master nodes by using SSH. Then, you can deploy applications by using the Helm CLI. For more information, see [Connect to a master node through SSH](#). You can also install and configure the Helm CLI and kubectl on your on-premises machine.

In this example, the Helm CLI and kubectl are installed and configured on your on-premises machine, and then an Apache Spark-based WordPress application is deployed.

1. Install and configure the Helm CLI and kubectl.
 - i. Install and configure kubectl on your on-premises machine.
For more information, see [Connect to a Kubernetes cluster through kubectl](#).
To view the details of a Kubernetes cluster, run the `kubectl cluster-info` command.
 - ii. Install Helm on your on-premises machine.
For more information, see [Install Helm](#).
2. Deploy the WordPress application.

In the following example, a WordPress blog website is deployed by using Helm.

- i. Run the following command:

```
helm install --name wordpress-test stable/wordpress
```

 **Note** Container Service allows you to mount disks as dynamically provisioned volumes. Before you deploy the application, you must create a disk volume.

The following output is returned:

```
NAME: wordpress-testLAST DEPLOYED: Mon Nov 20 19:01:55 2017NAMESPACE: defaultSTATUS: DEPLOYED...
```

- ii. Run the following command to query the release and Service created for the WordPress application:

```
helm listkubectl get svc
```

- iii. Run the following command to view the pods provisioned for the WordPress application. You may need to wait until the pods change to the Running state.

```
kubectl get pod
```

- iv. Run the following command to obtain the endpoint of the WordPress application:

```
echo http://$(kubectl get svc wordpress-test-wordpress -o jsonpath='{.status.loadBalancer.ingress[0].ip}')
```

You can enter the preceding URL into the address bar of your browser to access the WordPress application.

You can also run the following command based on the chart description to obtain the username and password of the administrator account for the WordPress application:

```
echo Username: userecho Password: $(kubectl get secret --namespace default wordpress-test-wordpress -o jsonpath="{.data.wordpress-password}" | base64 --decode)
```

- v. To delete the WordPress application, run the following command:

```
helm delete --purge wordpress-test
```

Use a third-party chart repository

You can use the default Alibaba Cloud chart repository. If a third-party chart repository is accessible from your cluster, you can also use the third-party chart repository. Run the following command to add a third-party chart repository to Helm:

```
helm repo add Repository name Repository URLhelm repo update
```

For more information about Helm commands, see [Helm documentation](#).

References

Helm contributes to the development of Kubernetes. A growing number of software suppliers, such as Bitnami, have provided high-quality charts. For more information about available charts, visit <https://kubernetes.io/docs/concepts/containers/helm/>.

3.6.8. SLB and Ingress

3.6.8.1. Overview

Container Service allows you to flexibly manage load balancing and customize load balancing policies for Kubernetes clusters. Kubernetes clusters provide you with a variety of methods to access containerized applications. They also allow you to use SLB or Ingress to access internal services and implement load balancing.

3.6.8.2. Use SLB to access Services

You can access a Service through Server Load Balancer (SLB).

Use the CLI

1. Create an NGINX application by using the command-line interface (CLI).

```
root@master # kubectl run nginx --image=registry.aliyuncs.com/acs/netdia:latest
root@master # kubectl get po
NAME                READY   STATUS    RESTARTS   AGE
nginx-2721357637-dvwq3    1/1     Running   1          6s
```

2. Create a Service for the NGINX application and set `type=LoadBalancer` to expose the NGINX Service to the Internet through an SLB instance.

```
root@master # kubectl expose deployment nginx --port=80 --target-port=80 --type=LoadBalancer
root@master # kubectl get svc
NAME                CLUSTER-IP   EXTERNAL-IP   PORT(S)          AGE
nginx               172.19.XX.XX  101.37.XX.XX  80:31891/TCP    4s
```

3. Copy `http://101.37.XX.XX` into the address bar of your browser and press Enter to access the NGINX Service.

SLB parameters

SLB provides a variety of parameters that you can use to configure features and services such as health check, billing method, and SLB instance type. For more information, see [SLB parameters](#).

Annotations

You can add annotations to use the load balancing features provided by SLB.

Use an existing internal-facing SLB instance

You need to add two annotations. You must replace "yourloadbalancer-id" with your SLB instance ID.

```

apiVersion: v1
kind: Service
metadata:
  annotations:
    service.beta.kubernetes.io/alibabacloud-loadbalancer-address-type: intranet
    service.beta.kubernetes.io/alibabacloud-loadbalancer-id: your-loadbalancer-id
  labels:
    run: nginx
  name: nginx
  namespace: default
spec:
  ports:
    - name: web
      port: 80
      protocol: TCP
      targetPort: 80
  selector:
    run: nginx
  sessionAffinity: None
  type: LoadBalancer

```

Save the preceding code as an `slb.svc` file and run the following command: `kubectl apply -f slb.svc`.

Create an HTTPS-based Service of the Loadbalancer type

You must first create a certificate in the SLB console. Then, you can use the certificate ID (`cert-id`) and the following template to create a LoadBalancer Service and an HTTPS-based SLB instance.

```

apiVersion: v1
kind: Service
metadata:
  annotations:
    service.beta.kubernetes.io/alibabacloud-loadbalancer-cert-id: your-cert-id
    service.beta.kubernetes.io/alibabacloud-loadbalancer-protocol-port: "https:443"
  labels:
    run: nginx
  name: nginx
  namespace: default
spec:
  ports:
    - name: web
      port: 443
      protocol: TCP
      targetPort: 443
  selector:
    run: nginx
  sessionAffinity: None
  type: LoadBalancer

```

 **Note** Annotations are case sensitive.

SLB parameters

| Annotation | Description | Default value |
|---|--|---------------|
| <code>service.beta.kubernetes.io/alibabacloud-loadbalancer-protocol-port</code> | The listening port. Separate multiple ports with commas (.). Example: <code>https:443,http:80</code> . | N/A |

| Annotation | Description | Default value |
|--|---|--|
| service.beta.kubernetes.io/alibabacloud-loadbalancer-address-type | The type of the SLB instance. Valid values: internet and intranet. | internet |
| service.beta.kubernetes.io/alibabacloud-loadbalancer-slb-network-type | The network type of the SLB instance. Valid values: classic and vpc. | classic |
| service.beta.kubernetes.io/alibabacloud-loadbalancer-charge-type | The billing method of the SLB instance. Valid values: paybytraffic and paybybandwidth. | paybybandwidth |
| service.beta.kubernetes.io/alibabacloud-loadbalancer-id | The ID of the SLB instance. You can set the loadbalancer-id parameter to specify an existing SLB instance and its existing listeners will be overwritten. The SLB instance will not be deleted if the Service is deleted. | N/A |
| service.beta.kubernetes.io/alibabacloud-loadbalancer-backend-label | The labels that specify the nodes to be added as backend servers of the SLB instance. | N/A |
| service.beta.kubernetes.io/alibabacloud-loadbalancer-region | The region where the SLB instance is deployed. | N/A |
| service.beta.kubernetes.io/alibabacloud-loadbalancer-bandwidth | The bandwidth of the SLB instance. | 50 |
| service.beta.kubernetes.io/alibabacloud-loadbalancer-cert-id | The certificate ID. You must upload the certificate first. | "" |
| service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-flag | Valid values: on and off. | Default value: off. If TCP is used, do not modify this parameter. The health check feature is enabled for TCP listeners by default and cannot be disabled. |
| service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-type | For more information, see <i>CreateLoadBalancerTCPListener</i> in the <i>SLB Developers Guide</i> . | N/A |
| service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-uri | For more information, see <i>CreateLoadBalancerTCPListener</i> in the <i>SLB Developers Guide</i> . | N/A |
| service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-connect-port | For more information, see <i>CreateLoadBalancerTCPListener</i> in the <i>SLB Developers Guide</i> . | N/A |
| service.beta.kubernetes.io/alibabacloud-loadbalancer-healthy-threshold | For more information, see <i>CreateLoadBalancerTCPListener</i> in the <i>SLB Developers Guide</i> . | N/A |
| service.beta.kubernetes.io/alibabacloud-loadbalancer-unhealthy-threshold | For more information, see <i>CreateLoadBalancerTCPListener</i> in the <i>SLB Developers Guide</i> . | N/A |

| Annotation | Description | Default value |
|---|---|---------------|
| service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-interval | For more information, see <i>CreateLoadBalancerTCPListener</i> in the <i>SLB Developers Guide</i> . | N/A |
| service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-connect-timeout | For more information, see <i>CreateLoadBalancerTCPListener</i> in the <i>SLB Developers Guide</i> . | N/A |
| service.beta.kubernetes.io/alibabacloud-loadbalancer-health-check-timeout | For more information, see <i>CreateLoadBalancerTCPListener</i> in the <i>SLB Developers Guide</i> . | N/A |

3.6.8.3. Configure Ingress monitoring

You can enable the virtual host traffic status (VTS) dashboard to view Ingress monitoring data.

Enable the VTS dashboard by using the CLI

1. Add the following configuration item to the Ingress ConfigMap: `enable-vts-status: "true"`.

```
root@master # kubectl edit configmap nginx-configuration -n kube-system
configmap "nginx-configuration" edited
```

The following template shows the modified Ingress ConfigMap:

```
apiVersion: v1
data:
  enable-vts-status: "true"# Enables the VTS dashboard.
  proxy-body-size: 20m
kind: ConfigMap
metadata:
  annotations:
    kubectl.kubernetes.io/last-applied-configuration: |
      {"apiVersion":"v1","data":{"proxy-body-size":"20m"},"kind":"ConfigMap","metadata":{"annotations":{},"labels":{"app":"ingress-nginx"},"name":"nginx-configuration","namespace":"kube-system"}}
  creationTimestamp: 2018-03-20T07:10:18Z
  labels:
    app: ingress-nginx
    name: nginx-configuration
    namespace: kube-system
  selfLink: /api/v1/namespaces/kube-system/configmaps/nginx-configuration
```

2. Verify that the VTS dashboard is enabled for the NGINX Ingress controller.

```
root@master # kubectl get pods --selector=app=ingress-nginx -n kube-system
NAME                READY  STATUS  RESTARTS  AGE
nginx-ingress-controller-79877595c8-78gq8  1/1   Running  0         1h
root@master # kubectl exec -it nginx-ingress-controller-79877595c8-78gq8 -n kube-system -- cat /etc/nginx/nginx.conf | grep vhost_traffic_status_display
vhost_traffic_status_display;
vhost_traffic_status_display_format html;
```

3. Access the VTS dashboard from an on-premises machine.

 **Note** By default, the VTS port is not exposed due to security concerns. In the following example, port forwarding is used to access the VTS dashboard.

```
root@master # kubectl port-forward nginx-ingress-controller-79877595c8-78gq8 -n kube-system 18080
Forwarding from 127.0.0.1:18080 -> 18080
Handling connection for 18080
```

4. Visit `http://localhost:18080/nginx_status` to access the VTS dashboard.

Nginx Vhost Traffic Status

Server main

| Host | Version | Uptime | Connections | | | | Requests | | | Shared memory | | | | |
|---|---------|---------|-------------|---------|---------|---------|----------|---------|-------|---------------|----------------------|----------|----------|----------|
| | | | active | reading | writing | waiting | accepted | handled | Total | Req/s | name | maxSize | usedSize | usedNode |
| nginx-ingress-controller-79877595c8-78gq8 | 1.13.7 | 32m 41s | 7 | 0 | 1 | 6 | 93566 | 93566 | 1428 | 1 | vhost_traffic_status | 10.0 MiB | 2.4 KiB | 1 |

Server zones

| Zone | Requests | | | Responses | | | | | Traffic | | | | | Cache | | | | | | | | | |
|------|----------|-------|------|-----------|-----|-----|-----|-----|---------|---------|-----------|---------|--------|-------|--------|---------|-------|----------|-------------|-----|-------|-------|---|
| | Total | Req/s | Time | 1xx | 2xx | 3xx | 4xx | 5xx | Total | Sent | Rcvd | Sent/s | Rcvd/s | Miss | Bypass | Expired | Stale | Updating | Revalidated | Hit | Scare | Total | |
| - | 660 | 1 | 0ms | 0 | 660 | 0 | 0 | 0 | 660 | 1.7 MiB | 145.4 KiB | 1.1 KiB | 503 B | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| * | 660 | 1 | 0ms | 0 | 660 | 0 | 0 | 0 | 660 | 1.7 MiB | 145.4 KiB | 1.1 KiB | 503 B | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

Upstreams

upstream-default-backend

| Server | State | Response Time | Weight | MaxFails | FailTimeout | Requests | | | Responses | | | | | Traffic | | | | | | | | | | |
|-----------------|-------|---------------|--------|----------|-------------|----------|-------|------|-----------|-----|-----|-----|-----|---------|------|------|--------|--------|---|---|---|---|---|---|
| | | | | | | Total | Req/s | Time | 1xx | 2xx | 3xx | 4xx | 5xx | Total | Sent | Rcvd | Sent/s | Rcvd/s | | | | | | |
| 172.16.3.6:8080 | up | 0ms | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |

update interval: 1 sec

[JSON](#) | [GITHUB](#)

3.6.8.4. Ingress support

Kubernetes clusters support Ingress rules. You can define Ingress rules based on your needs to implement load balancing.

In Kubernetes clusters, an Ingress is a collection of routing rules that authorize external access to cluster Services. You can use an Ingress to enable Layer-7 load balancing. You can configure an Ingress to provide Kubernetes Services with externally reachable URLs, SLB instances, SSL connections, and name-based virtual hosting.

Prerequisites

To test a complex routing scenario, an NGINX application that consists of multiple Services is created in this example. You need to create a Deployment and multiple Services for the NGINX application in advance. In practice, replace Service names with the actual values.

```
kubectl run nginx --image=registry.cn-hangzhou.aliyuncs.com/acs/netdia:latest
kubectl expose deploy nginx --name=http-svc --port=80 --target-port=80
kubectl expose deploy nginx --name=http-svc1 --port=80 --target-port=80
kubectl expose deploy nginx --name=http-svc2 --port=80 --target-port=80
kubectl expose deploy nginx --name=http-svc3 --port=80 --target-port=80
```

Create a simple Ingress

Run the following commands to create a simple Ingress that redirects traffic to the `/svc` path to a `http-svc` Service. `nginx.ingress.kubernetes.io/rewrite-target: /` redirects traffic destined for the `/svc` path to the `/path` that can be recognized by the backend application.

```
cat <<EOF | kubectl create -f -
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: simple
  annotations:
    nginx.ingress.kubernetes.io/rewrite-target: /
spec:
  rules:
  - http:
      paths:
      - path: /svc
        backend:
          serviceName: http-svc
          servicePort: 80
EOF
```

```
kubectl get ing
NAME      HOSTS      ADDRESS      PORTS      AGE
simple    *          101.37.19*.*** 80         11s
```

You can visit `http://101.37.19*.***/svc` to access the NGINX application.

Create a simple fanout Ingress that uses multiple domain names

You can create a simple fanout Ingress to route traffic from an external IP address to multiple Services with different domain names. The following example shows the configuration of a simple fanout Ingress:

```
cat <<EOF | kubectl create -f -
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: simple-fanout
spec:
  rules:
  - host: foo.bar.com
    http:
      paths:
      - path: /foo
        backend:
          serviceName: http-svc1
          servicePort: 80
      - path: /bar
        backend:
          serviceName: http-svc2
          servicePort: 80
  - host: foo.example.com
    http:
      paths:
      - path: /film
        backend:
          serviceName: http-svc3
          servicePort: 80
EOF
```

```
kubectl get ing
NAME          HOSTS          ADDRESS          PORTS  AGE
simple-fanout *  101.37.19*.*** 80    11s
```

You can visit `http://foo.bar.com/foo` to access Service `http-svc1` , visit `http://foo.bar.com/bar` to access Service `http-svc2` , and visit `http://foo.example.com/film` to access Service `http-svc3` .

Note

- In a production environment, you need to point the domain name to the returned address `101.37.19*.***` .
- In a test environment, you need to add the following mapping rules to the `hosts` file.

```
101.37.19*.*** foo.bar.com
101.37.19*.*** foo.example.com
```

Create a simple Ingress that uses the default domain name

If you do not have a domain name, you can use the default domain name associated with the Ingress to access the Service. The default domain name is in the following format: `*.[cluster-id].[region-id].alicontainer.com` . You can find the default domain name in the basic information of the Kubernetes cluster in the Container Service console.

The following example shows the configuration of a simple Ingress that allows external access to the Service through the default domain name:

```
cat <<EOF | kubectl create -f -
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: shared-dns
spec:
  rules:
  - host: foo.[cluster-id].[region-id].alicontainer.com ## Replace with the default domain name of your cluster.
    http:
      paths:
      - path: /
        backend:
          serviceName: http-svc1
          servicePort: 80
  - host: bar.[cluster-id].[region-id].alicontainer.com ## Replace with the default domain name of your cluster.
    http:
      paths:
      - path: /
        backend:
          serviceName: http-svc2
          servicePort: 80
EOF
```

```
kubectl get ing
NAME          HOSTS          ADDRESS          PORTS  AGE
shared-dns  foo.[cluster-id].[region-id].alicontainer.com,bar.[cluster-id].[region-id].alicontainer.com  47.95.16*.***
80    40m
```

You can visit `http://foo.[cluster-id].[region-id].alicontainer.com/` to access Service `http-svc1` and visit `http://bar.[cluster-id].[region-id].alicontainer.com` to access Service `http-svc2` .

Create an Ingress to secure data transmission

Container Service allows you to use multiple types of certificates to reinforce the security of your applications.

1. Prepare your certificate.

If you do not have a certificate, perform the following steps to generate a test certificate:

 **Note** The domain name must be the same as the one specified in your Ingress configuration.

```
openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout tls.key -out tls.crt -subj "/CN=foo.bar.com/O=foo.bar.com"
```

After you run the preceding command, a certificate file `tls.crt` and a private key file `tls.key` are generated.

Use the certificate and private key to create a Kubernetes Secret named `foo.bar`. You need to reference the Secret when you create the Ingress.

```
kubectl create secret tls foo.bar --key tls.key --cert tls.crt
```

2. Create an Ingress to secure data transmission.

```
cat <<EOF | kubectl create -f -
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: tls-fanout
spec:
  tls:
  - hosts:
    - foo.bar.com
    secretName: foo.bar
  rules:
  - host: foo.bar.com
    http:
      paths:
      - path: /foo
        backend:
          serviceName: http-svc1
          servicePort: 80
      - path: /bar
        backend:
          serviceName: http-svc2
          servicePort: 80
EOF
```

```
kubectl get ing
NAME      HOSTS      ADDRESS      PORTS      AGE
tls-fanout *          101.37.19*.*** 80         11s
```

3. You need to configure the `hosts` file or set a domain to access the `tls` Ingress, as described in [Simple fanout based on domains](#).

You can visit `http://foo.bar.com/foo` to access Service `http-svc1` and visit `http://foo.bar.com/bar` to access Service `http-svc2`.

You can also access the HTTPS Service by using HTTP. By default, the Ingress redirects HTTP traffic to the HTTPS address. Traffic to `http://foo.bar.com/foo` is automatically redirected to `https://foo.bar.com/foo`.

Create an Ingress

1. [Log on to the Container Service console](#).

2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Services and Ingresses > Ingresses**.
5. On the **Ingresses** page, select a namespace and click **Create Resources in YAML** to create an Ingress.
6. On the **Create** page, select **Custom** from the **Sample Template** drop-down list, copy the following content to the template, and then click **Create**.

```
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: simple
spec:
  rules:
  - http:
      paths:
      - path: /svc
        backend:
          serviceName: http-svc
          servicePort: 80
```

This way, an Ingress that routes Layer-7 traffic for Service `http-svc` is created.

3.6.8.5. Ingress configurations

Container Service provides Ingress controller components. Integrated with Apsara Server Load Balancer, these components provide Kubernetes clusters with flexible and reliable Ingress service.

An Ingress orchestration template is provided below. When you configure an Ingress through the console, you need to configure annotations and may need to create dependencies. For more information, see [Create an ingress through the console](#), [Ingress support](#), and [Kubernetes Ingress](#). You can also create ConfigMaps to configure Ingresses. For more information, see <https://kubernetes.github.io/ingress-nginx/user-guide/nginx-configuration/configmap/>.

```

apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  annotations:
    nginx.ingress.kubernetes.io/service-match: 'new-nginx: header("foo", /^bar$/)' #Canary release rule. In this example, the request header is used.
    Nginx.ingress.kubernetes.io/service-weight: 'New-nginx: 50, old-nginx: 50' #The route weight.
  creationTimestamp: null
  generation: 1
  name: nginx-ingress
  selfLink: /apis/extensions/v1beta1/namespaces/default/ingresses/nginx-ingress
spec:
  rules: ##The Ingress rule.
  - host: foo.bar.com
    http:
      paths:
      - backend:
          serviceName: new-nginx
          servicePort: 80
        path: /
      - backend:
          serviceName: old-nginx
          servicePort: 80
        path: /
  tls: ## Enable TLS for secure routing.
  - hosts:
    - *.xxxxxx.cn-hangzhou.alicontainer.com
    - foo.bar.com
    secretName: nginx-ingress-secret ##The Secret name.
status:
  loadBalancer: {}

```

Annotations

For each Ingress, you can configure its annotations, Ingress controller, and rules, such as the route weight, canary release rule, and rewrite rules. For more information about annotations, see <https://kubernetes.github.io/ingress-nginx/user-guide/nginx-configuration/annotations/>.

For example, the following rewrite annotation, `nginx.ingress.kubernetes.io/rewrite-target: /`, indicates that `/path` is redirected to the root path `/`, which can be recognized by the backend service.

Rules

Ingress rules are used to manage external access to the services in the cluster and can be HTTP or HTTPS rules. You can configure the following items in rules: domain name (virtual hostname), URL path, service name, and port.

For each rule, you need to set the following parameters:

- Domain: The test domain or virtual hostname of your service, such as `foo.bar.com`.
- Path: The URL path of your service. Each path is associated with a backend service. Server Load Balancer only forwards traffic to the backend if the incoming request matches the domain and path.
- Service: Specify the service in the form of `service:port`. You also need to specify a route weight for each service. The Ingress routes traffic to the matching service based on the route weight.
 - Name: The name of the backend service.
 - Port: The port of the service.

- **Weight:** The route weight of the service in the service group.

 **Note**

- a. The weight is a percentage value. For example, you can set two services to the same weight of 50%.
- b. A service group includes services that have the same domain and path defined in the Ingress configuration. If no weight is set for a service, the default value, 100, is used.

Canary release

Container Service supports multiple traffic splitting approaches to suit scenarios such as canary release and A/B testing.

 **Note** Currently, only Ingress controllers of 0.12.0-5 and later versions support traffic splitting.

1. Traffic splitting based on request header
2. Traffic splitting based on cookie
3. Traffic splitting based on query parameter

After canary release is configured, only requests that match certain rules are routed to the corresponding service. If the weight of the corresponding service is lower than 100%, requests that match certain rules are routed to one of the services in the service group based on the weight.

TLS

You can use a Secret that contains a TLS private key and certificate to encrypt the Ingress. This ensures secure routing. The TLS Secret must contain a certificate named `tls.crt` and a private key named `tls.key`. For more information about how TLS works, see [TLS](#). For how to create a Secret, see [Configure a secure Ingress](#).

Labels

You can add labels to the Ingress.

3.6.8.6. Create an Ingress in the console

The Container Service console is integrated with the Ingress service. You can create an Ingress in the console and manage inbound traffic that is forwarded to different Services to meet your business requirements.

Prerequisites

- A Kubernetes cluster is created and an Ingress controller runs as normal in the cluster. For more information, see [Create a Kubernetes cluster](#).
- You are connected to a master node by using `kubectl`. For more information, see [Connect to a Kubernetes cluster through kubectl](#).
- The image address used in this example requires Internet access. You can use the image address of your own cluster to replace it. You can also build and push the image used in this example to a repository and then pull it from the repository when you use it.

Step 1: Create a Deployment and a Service

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Workloads > Deployments**.
5. On the **Deployments** page, select a namespace and click **Create from Template** in the upper-right corner.

6. On the **Create** page, select a sample template or customize a template and click **Create**.

In this example, two NGINX applications are created: old-nginx and new-nginx.

The following template is used to create the old-nginx application:

```
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: old-nginx
spec:
  replicas: 2
  selector:
    matchLabels:
      run: old-nginx
  template:
    metadata:
      labels:
        run: old-nginx
    spec:
      containers:
        - image: registry.cn-hangzhou.aliyuncs.com/xianlu/old-nginx
          imagePullPolicy: Always
          name: old-nginx
          ports:
            - containerPort: 80
              protocol: TCP
          restartPolicy: Always
---
apiVersion: v1
kind: Service
metadata:
  name: old-nginx
spec:
  ports:
    - port: 80
      protocol: TCP
      targetPort: 80
  selector:
    run: old-nginx
  sessionAffinity: None
  type: NodePort
```

The following template is used to create the new-nginx application:

```
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: new-nginx
spec:
  replicas: 1
  selector:
    matchLabels:
      run: new-nginx
  template:
    metadata:
      labels:
        run: new-nginx
    spec:
      containers:
        - image: registry.cn-hangzhou.aliyuncs.com/xianlu/new-nginx
          imagePullPolicy: Always
          name: new-nginx
          ports:
            - containerPort: 80
              protocol: TCP
          restartPolicy: Always
---
apiVersion: v1
kind: Service
metadata:
  name: new-nginx
spec:
  ports:
    - port: 80
      protocol: TCP
      targetPort: 80
  selector:
    run: new-nginx
  sessionAffinity: None
  type: NodePort
```

7. In the left-side navigation pane of the details page, choose **Services and Ingresses** > **Services**.
After the Services are created, you can find these Services on the Services page.

Step 2: Create an Ingress

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Services and Ingresses** > **Ingresses**.
5. On the **Ingresses** page, select a namespace and click **Create** in the upper-right corner.
6. In the dialog box that appears, enter a name for the Ingress. In this example, the Ingress is named `nginx-ingress`.
7. Configure Ingress rules.

Ingress rules are used to manage external access to Services in the cluster. Ingress rules can be HTTP or HTTPS rules. You can configure the following items in rules: domain name (virtual host name), URL path, Service name, port, and weight. For more information, see [Ingress configurations](#).

In this example, a complex rule is added to configure Services for the default domain name and virtual hostname of the cluster. Traffic is routed based on domain names.

Rule: + Add

Domain:

Select *:

path:

Service + Add

| Name | Port | Weight | Percent of Weight |
|--|---------------------------------|----------------------------------|-------------------|
| <input type="text" value="new-nginx"/> | <input type="text" value="80"/> | <input type="text" value="100"/> | 50.0% |
| <input type="text" value="old-nginx"/> | <input type="text" value="80"/> | <input type="text" value="100"/> | 50.0% |

Create a simple fanout Ingress that uses multiple domain names

In this example, a virtual host name is used as the test domain name for external access. Route weights are specified for two backend Services and canary release settings are configured for one of the Services. In your production environment, you can use a domain name that has obtained an Internet Content Provider (ICP) number for external access.

- Domain: Enter the test domain name. In this example, the test domain name is `foo.bar.com`.

You must add the following domain name mapping to the hosts file:

```
118.178.XX.XX foo.bar.com #The IP address of the Ingress.
```

- Services: Set the names, paths, port numbers, and weights of the backend Services that you want to access.
 - Path: Enter the URL of the backend Service. In this example, the root path `/` is used.
 - Name: In this example, both the `old-nginx` and `new-nginx` Services are specified.
 - Port: In this example, port `80` is open.
 - Weight: Set a weight for each backend Service. The weight is a percentage value. The default value is `100`. In this example, the weight of each backend Service is `50`. This means that the two backend Services have the same weight.
8. Configure Transport Layer Security (TLS). Select **Enable TLS** to enable TLS and configure a secure Ingress. For more information, see [Configure a secure Ingress](#).
- You can use an existing Secret.

TLS: Enable Exist secret Create secret

- Log on to a master node. Create a file named `tls.key` and another file named `tls.crt`.

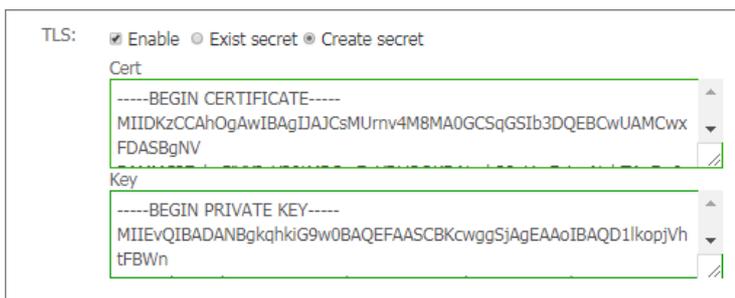
```
openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout tls.key -out tls.crt -subj "/CN=foo.bar.com/O=foo.bar.com"
```

- Create a Secret.

```
kubectl create secret tls foo.bar --key tls.key --cert tls.crt
```

- Run the `kubectl get secret` command and verify that the Secret is created. Then, you can select the newly created Secret `foo.bar`.

- o You can also use the TLS private key and certificate to create a Secret.



- Log on to a master node, and then create a file named `tls.key` and another file named `tls.crt`.

```
openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout tls.key -out tls.crt -subj "/CN=foo.bar.com/O=foo.bar.com"
```

- Run the `vim tls.key` and `vim tls.crt` commands to obtain the private key and certificate that are generated.
- Paste the certificate to the Cert field and the private key to the Key field.

9. Configure canary release settings.

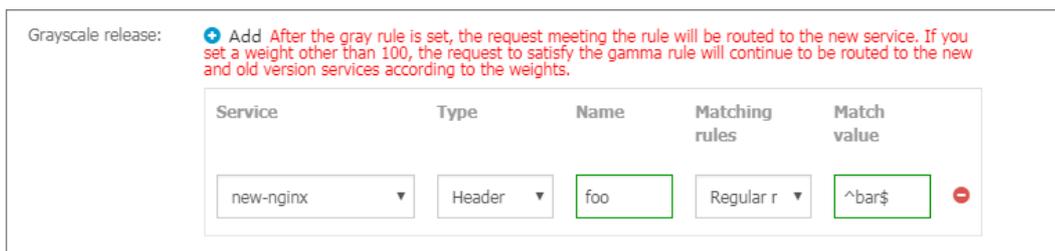
Note Only Ingress controllers of 0.12.0-5 and later versions support traffic splitting.

Container Service supports multiple traffic splitting methods. This allows you to select suitable solutions for specific scenarios, such as canary releases and A/B testing:

- Traffic splitting based on request headers.
- Traffic splitting based on cookies
- Traffic splitting based on query parameters.

After canary release settings are configured, only requests that match the specified rules are routed to the new-nginx Service. If the weight of new-nginx is lower than 100%, requests that match certain rules are routed to this Service based on the weight.

In this example, the rule is added to specify a request header that matches the regular expression `foo=^bar$`. Only requests with headers that match the regular expression are forwarded to new-nginx.



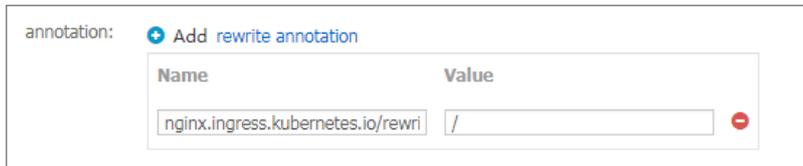
- o **Services:** Specify the Services to be accessed.
- o **Type:** Select the type of matching rule. Valid values: Header, Cookie, and Query.
- o **Name and Match Value:** Specify custom request fields. Each field is a key-value pair.
- o **Matching Rule:** Regular expressions and exact matches are supported.

10. Configure annotations.

Click **Rewrite Annotation** and add an annotation to redirect inbound traffic for the Ingress. For example, `nginx.ingress.kubernetes.io/rewrite-target: /` specifies that `/path` is redirected to the root path `/`. The root path can be recognized by the backend Services.

Note In this example, no path is configured for the backend Services. Therefore, you do not need to configure rewrite annotations. Rewrite annotations allow the Ingress to forward traffic through root paths to the backend Services. This avoids the 404 error that is caused by invalid paths.

You can also click **Add** to enter annotation names and values in key-value pairs. For more information about Ingress annotations, visit <https://kubernetes.github.io/ingress-nginx/user-guide/nginx-configuration/annotations/>.

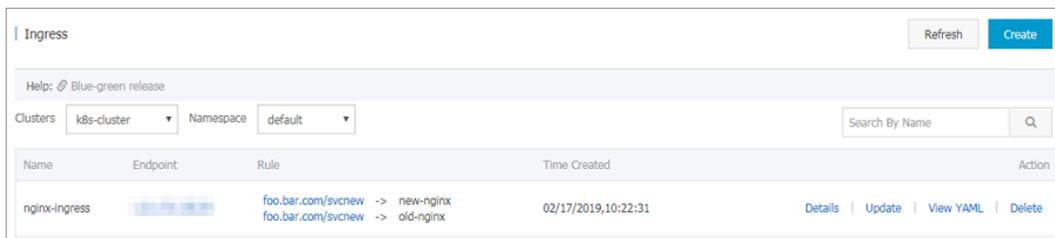


11. Add labels.

Add labels to describe the Ingress.

12. Click **Create**. You are redirected to the Ingresses page.

After the Ingress is created, you can find the nginx-ingress Ingress on the Ingresses page.



13. Click **foo.bar.com** to view the NGINX welcome page.

Click the domain name that points to new-nginx, the old-nginx Service page appears.

Note By default, when you enter the domain name in the browser, request headers do not match the `foo=^bar$` regular expression. Therefore, requests are directed to old-nginx.



14. Log on to a master node by using SSH. Run the following commands to simulate requests with specific headers and check the results:

```
curl -H "Host: foo.bar.com" http://47.107.XX.XX
old
curl -H "Host: foo.bar.com" http://47.107.XX.XX
old
curl -H "Host: foo.bar.com" http://47.107.XX.XX      #Similar to a browser request.
old
curl -H "Host: foo.bar.com" -H "foo: bar" http://47.107.XX.XX      #Simulate a request with a specific header. The results are returned based on the weight.
new
curl -H "Host: foo.bar.com" -H "foo: bar" http://47.107.XX.XX
old
curl -H "Host: foo.bar.com" -H "foo: bar" http://47.107.XX.XX
old
curl -H "Host: foo.bar.com" -H "foo: bar" http://47.107.XX.XX
new
```

3.6.8.7. Update an Ingress

You can update Ingresses in the Container Service console.

Prerequisites

- A Kubernetes cluster is created and an Ingress controller is running as normal in the cluster. For more information, see [Create a Kubernetes cluster](#).
- An Ingress is created. For more information, see [Create an ingress through the console](#).

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Services and Ingresses > Ingresses**.
5. On the **Ingresses** page, select the namespace, find the Ingress that you want to update, and then click **Update** in the **Actions** column.
6. In the dialog box that appears, modify the parameters and click **Update**. In this example, change `foo.bar.com` to `test.bar.com`.
On the Ingresses page, you can find the changed Ingress rule.

3.6.8.8. Delete an Ingress

This topic describes how to delete an Ingress.

Prerequisites

- A Kubernetes cluster is created and an Ingress controller is running as normal in the cluster. For more information about how to create a cluster, see [Create a Kubernetes cluster](#).
- An Ingress is created. For more information, see [Create an ingress through the console](#).

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Services and Ingresses > Ingresses**.
5. On the **Ingresses** page, select the namespace, find the Ingress that you want to delete, and then click **Delete**

in the Actions column.

6. In the message that appears, click **Confirm**.

3.6.9. Config maps and secrets

3.6.9.1. Create a ConfigMap

In the Container Service console, you can create a ConfigMap on the ConfigMap page or by using a template. This topic describes how to create a ConfigMap.

Create a ConfigMap on the ConfigMap page

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Configurations > ConfigMaps**.
5. On the **ConfigMap** page, select a namespace and click **Create**.
6. Set the parameters and click **OK**. Create a ConfigMap on the ConfigMap page

| Parameter | Description |
|-----------------------|--|
| Clusters | The ID of the cluster that you have selected. |
| Namespaces | The namespace that you have selected. A ConfigMap is a Kubernetes resource object and must be scoped to a namespace. |
| ConfigMap Name | The name of the ConfigMap. The name can contain lowercase letters, digits, hyphens (-), and periods (.). This parameter is required. Other resource objects must reference the ConfigMap name to obtain the configuration information. |
| ConfigMap | Specify Name and Value , and then click Add to add the key-value pair. You can also click Edit YAML file , modify the parameters in the dialog box that appears, and then click OK . |

In this example, two variables named `enemies` and `lives` are created. Their values are set to `aliens` and `3` separately.

7. Click **OK**. You can find the test-config ConfigMap on the ConfigMap page.

You can also click **Browse** and upload a configuration file to create a ConfigMap.

Create a ConfigMap from a template

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Workloads > Deployments**.
5. On the **Deployments** page, select the namespace and click **Create from Template** in the upper-right corner.
6. On the page that appears, set the parameters and click **Create**. Create a ConfigMap from a template

| Parameter | Description |
|-----------------|--|
| Sample Template | Container Service provides YAML templates of various resource types to help you quickly deploy resource objects. You can select <i>Custom</i> from the drop-down list and configure your ConfigMap based on YAML syntax. You can also select the <i>Resource-ConfigMap</i> template to create a ConfigMap. In the sample template, the ConfigMap is named aliyun-config and contains two variable files: <i>game.properties</i> and <i>ui.properties</i> . You can modify the ConfigMap based on your needs. |
| Template | Enter the template content based on YAML syntax. The template can contain multiple resource objects that are separated by <code>---</code> . |
| Add Deployment | This feature allows you to quickly define a YAML template. You can click Use Existing Template to import an existing template. |

After the deployment is completed, you can find the ConfigMap named *aliyun-config* on the ConfigMap page.

3.6.9.2. Use a ConfigMap in a pod

This topic describes how to use a ConfigMap in a pod.

You can use a ConfigMap in a pod in the following scenarios:

- Use a ConfigMap to define environment variables for a pod.
- Use a ConfigMap to set command line parameters.
- Use a ConfigMap in a volume.

For more information, see [Configure a pod to use a ConfigMap](#).

Limits

To use a ConfigMap in a pod, make sure that the ConfigMap and the pod are in the same cluster and namespace.

Create a ConfigMap

In this example, a ConfigMap named `special-config` is created. This ConfigMap consists of two key-value pairs:

```
SPECIAL_LEVEL: very and SPECIAL_TYPE: charm .
```

You can use the following YAML template to create the ConfigMap:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: special-config
  namespace: default
data:
  SPECIAL_LEVEL: very
  SPECIAL_TYPE: charm
```

You can also log on to the Container Service console and choose **Configuration > ConfigMaps** in the left-side navigation pane. You can then click **Create** to create the ConfigMap.

Clusters: [dropdown]

Namespace: default

* ConfigMap Name:

The name must be 1 to 253 characters in length and can contain only lower-case letters numbers hyphens (-) and periods (.).

ConfigMap:

| Name | Value |
|--|------------------------------------|
| <input type="text" value="SPECIAL_TYPE"/> | <input type="text" value="charm"/> |
| <input type="text" value="SPECIAL_LEVEL"/> | <input type="text" value="very"/> |

A name can contain only numbers letters underscores (_) hyphens (-) and periods (.).

Use a ConfigMap to define one or multiple environment variables for a pod

Use a key-value pair of a ConfigMap to define one environment variable

You can log on to the Container Service console. In the left-side navigation pane, click **Clusters**. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column. In the left-side navigation pane of the details page, choose **Workloads > Deployments**. Click **Create from Template**, select a pod template from the Sample Template drop-down list, and then start the deployment.

You can use the following sample template to create a pod and defines environment variables in the pod. `valueFrom` is used to reference the value of `SPECIAL_LEVEL` to define an environment variable.

```

apiVersion: v1
kind: Pod
metadata:
  name: config-pod-1
spec:
  containers:
  - name: test-container
    image: busybox
    command: [ "/bin/sh", "-c", "env" ]
    env:
    - name: SPECIAL_LEVEL_KEY
      valueFrom:          ##valueFrom is used to reference the value of the ConfigMap to define an environment variable.
        configMapKeyRef:
          name: special-config    ##The name of the referenced ConfigMap.
          key: SPECIAL_LEVEL     ##The key of the referenced key-value pair.
    restartPolicy: Never

```

To use the values of multiple ConfigMaps to define multiple environment variables, add multiple env parameters to the pod configuration file.

Use all the key-value pairs of a ConfigMap to define multiple environment variables

To define the key-value pairs of a ConfigMap as pod environment variables, you can use the envFrom parameter. The keys in a ConfigMap are used as the names of the environment variables.

The following sample template is used to create a pod:

```

apiVersion: v1
kind: Pod
metadata:
  name: config-pod-2
spec:
  containers:
  - name: test-container
    image: busybox
    command: [ "/bin/sh", "-c", "env" ]
    envFrom:          ##Reference all the key-value pairs of the special-config ConfigMap.
    - configMapRef:
      name: special-config
    restartPolicy: Never

```

Use a ConfigMap to set command line parameters

You can use ConfigMaps to define the commands or parameter values for a container by using the environment variable replacement syntax `$(VAR_NAME)`. The following template is used as an example:

```
apiVersion: v1
kind: Pod
metadata:
  name: config-pod-3
spec:
  containers:
    - name: test-container
      image: busybox
      command: [ "/bin/sh", "-c", "echo ${SPECIAL_LEVEL_KEY} ${SPECIAL_TYPE_KEY}" ]
      env:
        - name: SPECIAL_LEVEL_KEY
          valueFrom:
            configMapKeyRef:
              name: special-config
              key: SPECIAL_LEVEL
        - name: SPECIAL_TYPE_KEY
          valueFrom:
            configMapKeyRef:
              name: special-config
              key: SPECIAL_TYPE
      restartPolicy: Never
```

After you run the pod, the following output is returned:

```
very charm
```

Use a ConfigMap in a volume

You can use a ConfigMap to define volumes. The following sample template specifies a ConfigMap name under volumes. This stores the key-value pairs of the ConfigMap to the path that you specified in the mountPath field. In this example, the path is /etc/config. This generates configuration files that are named after the keys of the ConfigMap. The corresponding values of the ConfigMap are stored in these files.

```
apiVersion: v1
kind: Pod
metadata:
  name: config-pod-4
spec:
  containers:
    - name: test-container
      image: busybox
      command: [ "/bin/sh", "-c", "ls /etc/config/" ] ##List the files under the directory.
      volumeMounts:
        - name: config-volume
          mountPath: /etc/config
      volumes:
        - name: config-volume
          configMap:
            name: special-config
      restartPolicy: Never
```

After you run the pod, the following output is returned:

```
SPECIAL_TYPE
SPECIAL_LEVEL
```

3.6.9.3. Update a ConfigMap

You can use multiple methods to update a ConfigMap.

Considerations

If you update a ConfigMap, the applications that use the ConfigMap are affected.

Procedure

1. [Log on to the Container Service console.](#)
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Configurations > ConfigMaps**.
5. On the **ConfigMap** page, select the namespace, find the ConfigMap that you want to update, and then click **Edit** in the **Actions** column.
6. In the dialog box that appears, modify the configurations and click **OK**.

3.6.9.4. Delete a ConfigMap

You can use multiple methods to delete a ConfigMap.

Considerations

If you delete a ConfigMap, the applications that use this ConfigMap are affected.

Procedure

1. [Log on to the Container Service console.](#)
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Configurations > ConfigMaps**.
5. On the **ConfigMap** page, select the namespace, find the ConfigMap that you want to delete, and then click **Delete** in the **Actions** column.
6. In the message that appears, click **OK**.

3.6.9.5. Create a Secret

You can create Secrets for applications in the Container Service console.

Prerequisites

A Kubernetes cluster is created.

Context

We recommend that you use Secrets to store sensitive information in Kubernetes clusters, such as passwords and certificates.

Secrets are classified into the following types:

- **Service account:** A service account is automatically created by Kubernetes and automatically mounted to the `/run/secrets/kubernetes.io/serviceaccount` directory of a pod. The service account provides an identity for the pod to interact with the API server.
- **Opaque:** This type of secret is encoded in Base64 and used to store sensitive information, such as passwords and certificates.

By default, you can create only Opaque Secrets in the Container Service console. Opaque Secrets store map type data. Therefore, values must be encoded in Base64. You can create Secrets in the Container Service console with a few clicks. Plaintext is automatically encoded in Base64.

You can also create Secrets by using the CLI. For more information, see [Kubernetes Secrets](#).

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Configurations > Secrets**.
5. On the **Secrets** page, select the namespace and click **Create** in the upper-right corner.
6. Configure the Secret and click **OK**.

 **Note** To enter Secret data in plaintext, select **Encode Data Values Using Base64**.

| Parameter | Description |
|---------------------------------|--|
| Name | Enter a name for the Secret. The name must be 1 to 253 characters in length, and can contain only lowercase letters, digits, hyphens (-), and periods (.). |
| Namespace | Select the namespace of the Secret. |
| Type | You can select Opaque, Private Repository Logon Secret, or TLS Certificate. |
| Opaque | If you set Type to Opaque, configure the following parameters: <ul style="list-style-type: none"> ◦ (Optional) To enter Secret data in plaintext, select Encode Data Values Using Base64. ◦ Configure the Secret in key-value pairs. Click + Add. Enter the keys and values for the Secret in the Name and Value fields. |
| Private Repository Logon Secret | If you set Type to Private Repository Logon Secret, configure the following parameters: <ul style="list-style-type: none"> ◦ Docker Registry URL: Enter the address of the Docker registry where your Secret is stored. ◦ Username: Enter the username that is used to log on to the Docker registry. ◦ Password: Enter the password that is used to log on to the Docker registry. |
| TLS Certificate | If you set Type to TLS Certificate, configure the following parameters: <ul style="list-style-type: none"> ◦ Cert: Enter a TLS certificate. ◦ Key: Enter the key for the TLS certificate. |

You can view the newly created Secret on the Secrets page.

3.6.9.6. Modify a Secret

This topic describes how to modify a Secret in the Container Service console.

Prerequisites

- A Kubernetes cluster is created.
- A Secret is created. For more information, see [Create a secret](#).

Procedure

1. [Log on to the Container Service console.](#)
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Configurations > Secrets**.
5. On the **Secrets** page, select the namespace, find the Secret that you want to modify, and then click **Edit** in the **Actions** column.
6. In the dialog box that appears, modify the Secret and click **OK**.

3.6.9.7. Delete a Secret

This topic describes how to delete a Secret in the Container Service console.

Prerequisites

- A Kubernetes cluster is created.
- A Secret is created. For more information, see [Create a secret](#).

Context

 **Note** Do not delete Secrets that are generated when the cluster is created.

Procedure

1. [Log on to the Container Service console.](#)
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Configurations > Secrets**.
5. On the **Secrets** page, select the namespace, find the Secret that you want to delete, and then click **Delete** in the **Actions** column.
6. In the message that appears, click **OK**.

3.6.10. Templates

3.6.10.1. Create an orchestration template

This topic describes how to use multiple methods to create orchestration templates through the Container Service console.

Procedure

1. [Log on to the Container Service console.](#)
2. In the left-side navigation pane, choose **Marketplace > Orchestration Templates** and click **Create** in the upper-right corner.
3. In the dialog box that appears, configure the template, and then click **Save**. This example demonstrates how to create a Tomcat application template that contains a deployment and a service.
 - **Name**: The name of the template.
 - **Description**: Optional. The description of the template.
 - **Template**: Enter the template content based on YAML syntax. The template can contain multiple resource objects that are separated by `---`.

Create

Name:
The name should be 1-64 characters long, and can contain numbers, English letters, Chinese characters and hyphens.

Description:

Template:

```
1 apiVersion: apps/v1beta2 # for versions before 1.8.0 use
2   apps/v1beta1
3   kind: Deployment
4   metadata:
5     name: tomcat-deployment
6     labels:
7       app: tomcat
8   spec:
9     replicas: 1
10    selector:
11      matchLabels:
12        app: tomcat
13    template:
14      metadata:
15        labels:
16          app: tomcat
17      spec:
18        containers:
19          - name: tomcat
20            image: tomcat # replace it with your exactly
21              <image_name:tags>
22            ports:
23              - containerPort: 8080
```

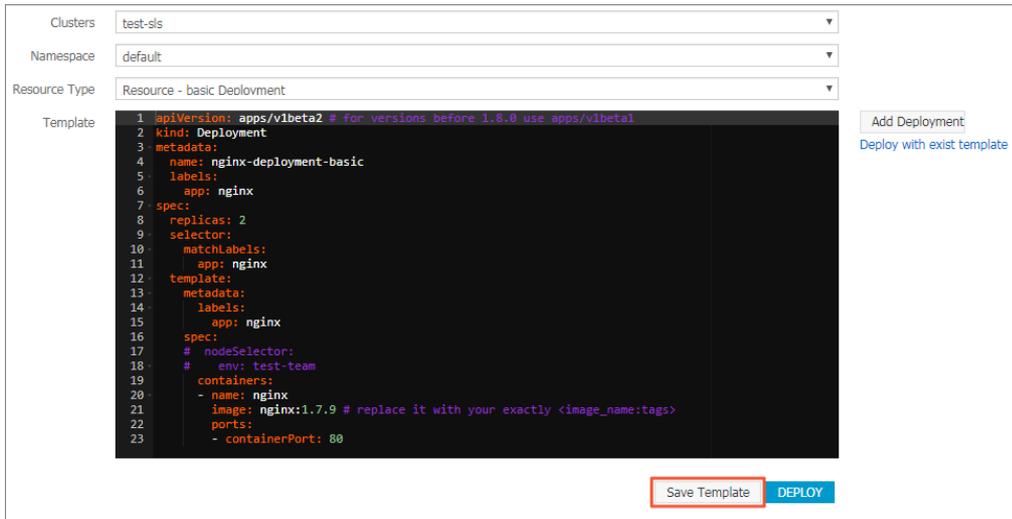
Save Cancel

4. After the template is created, you are redirected to the **Templates** page by default. You can find the template on the **My Templates** tab.



5. (Optional) You can also choose **Applications > Deployments** in the left-side navigation pane, and click **Create from Template** to go to the **Create from Template** page. You can modify a built-in template provided by Container Service and save it as a custom template.

i. Select a built-in template and click **Save Template**.



ii. In the dialog box that appears, specify the name, description, and content. Click **Save** to save the template.

Note You can modify the built-in template based on your needs.

iii. In the left-side navigation pane, choose **Market place > Orchestration Templates**. You can find the newly created template on the **My Templates** tab.



What's next

You can use the orchestration templates on the **My Templates** tab to quickly create applications.

3.6.10.2. Update an orchestration template

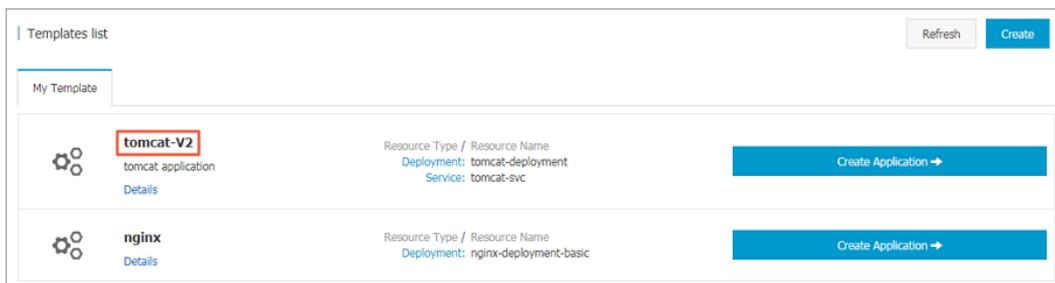
This topic describes how to edit and update an orchestration template.

Prerequisites

You have created an orchestration template. For more information, see [Create orchestration templates](#).

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, choose **Market place > Orchestration Templates**. The **Templates** page appears. You can view existing templates on the **My Templates** tab.
3. Select the target template and click **Details**.
4. On the template details page, click **Edit** in the upper-right corner.
5. In the dialog box that appears, edit the name, description, and template content, and click **Save**.
6. Go to the **Templates** page. You can view the template that you have updated on the **My Templates** tab.



3.6.10.3. Save an orchestration template as a new one

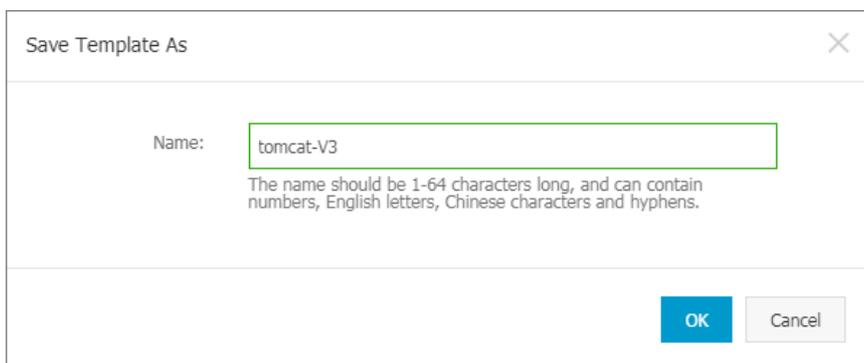
This topic describes how to save an orchestration template as a new one.

Prerequisites

You have created an orchestration template. For more information, see [Create orchestration templates](#).

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, choose **Market place > Orchestration Templates**. The **Templates** page appears. You can view existing templates on the **My Templates** tab.
3. Select the target template and click **Details**.
4. On the template details page, modify the template and click **Save As** in the upper-right corner.
5. In the dialog box that appears, enter the template name and click **OK**.



6. Go to the **Templates** page. The newly saved template is displayed on the **My Templates** tab.



3.6.10.4. Download an orchestration template

This topic describes how to download an orchestration template.

Prerequisites

You have created an orchestration template. For more information, see [Create orchestration templates](#).

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, choose **Market place > Orchestration Templates**. The **Templates** page appears. You can view existing templates on the **My Templates** tab.
3. Select the target template and click **Details**.
4. On the template details page, click **Download** in the upper-right corner to download the template as a YAML file.

3.6.10.5. Delete an orchestration template

Prerequisites

You have created an orchestration template. For more information, see [Create orchestration templates](#).

Procedure

1. [Log on to the Container Service console](#)
2. In the left-side navigation pane, choose **Market place > Orchestration Templates**. The **Templates** page appears. You can view existing templates on the **My Templates** tab.
3. Select the target template and click **Details**.
4. On the template details page, click **Delete** in the upper-right corner.
5. In the dialog box that appears, click **OK**.

3.6.11. Auto scaling

3.6.11.1. Auto scaling of nodes

Container Service provides the auto scaling component to automatically scale the number of nodes. Regular instances and GPU-accelerated instances can be automatically added to or removed from a Container Service cluster to meet your business requirements. This component supports multiple scaling modes, various instance types, and instances that are deployed across zones. This component is applicable to diverse scenarios.

How it works

The auto scaling model of Kubernetes is different from the traditional scaling model that is based on the resource usage threshold. Developers must understand the differences between the two scaling models before they migrate workloads from traditional data centers or other orchestration systems such as Swarm to Kubernetes.

The traditional scaling model is based on resource usage. For example, if a cluster contains three nodes and the CPU utilization or memory usage of the nodes exceeds the scaling threshold, new nodes are added to the cluster. However, you must consider the following issues when you use the traditional scaling model:

- How do you set a proper scaling threshold and how does the system check whether the threshold is exceeded?
In a Kubernetes cluster, the resource usage of hot nodes is higher than that of other nodes. If you specify the average resource usage as the scaling threshold, scaling activities may not be promptly triggered. If you specify the lowest node resource usage as the scaling threshold, the newly added nodes may not be used. This causes a waste of resources.
- How are the loads balanced after new nodes are added?
In Kubernetes, pods are the smallest deployable units for applications. Pods are deployed on different nodes in a Kubernetes cluster. When auto scaling is triggered for a cluster or a node in the cluster, pods with high resource usage are not replicated and the resource limits of these pods are not changed. As a result, the loads cannot be balanced to newly added nodes.
- How do you determine whether scaling activities must be triggered and how are scaling activities performed?

If scale-in activities are triggered based on resource usage, pods that request large amounts of resources but have low resource usage may be evicted. If the number of these pods is large within a Kubernetes cluster, resources may be exhausted and some pods may fail to be scheduled.

How does the auto scaling model of Kubernetes fix these issues? Kubernetes provides a two-layer scaling model that decouples pod scheduling from resource scaling.

In simple terms, pods are scaled based on resource usage. When pods enter the Pending state due to insufficient resources, a scale-out activity is triggered. After new nodes are added to the cluster, the pending pods are automatically scheduled to the newly added nodes. This way, the loads of the application are balanced. The following section describes the auto scaling model of Kubernetes in detail:

- How is a scale-out activity triggered?

The cluster-autoscaler component scans for pending pods and then triggers scaling activities. When pods enter the Pending state due to insufficient resources, cluster-autoscaler simulates pod scheduling to decide the scaling group that can provide new nodes to accept the pending pods. If a scaling group meets the requirement, nodes from this scaling group are added to the cluster. In simple terms, a scaling group is treated as a node during the simulation. The instance type of the scaling group specifies the CPU, memory, and GPU resources of the node. The labels and taints of the scaling group are also applied to the node. The node is used to simulate the scheduling of the pending pods. If the pending pods can be scheduled to the node, cluster-autoscaler calculates the number of nodes that need to be added from the scaling group.

- How is a scale-in activity triggered?

Only nodes that are added by scaling activities can be removed by cluster-autoscaler. This component cannot manage static nodes. Each node is separately evaluated to determine whether the node needs to be removed. If the resource usage of a node drops below the scale-in threshold, a scale-in activity is triggered for the node. In this case, cluster-autoscaler simulates the eviction of all workloads on the node to determine whether the node can be completely drained. cluster-autoscaler does not drain the nodes that contain specific pods, such as non-DaemonSet pods in the kube-system namespace and pods that are controlled by PodDisruptionBudgets (PDBs). A node is drained before it is removed. After pods on the node are evicted to other nodes, the node can be removed.

- How is a scaling group selected from multiple scaling groups that meet the requirements?

Different scaling groups are treated as nodes in different specifications. cluster-autoscaler selects a scaling group based on a scoring policy that is similar to a scheduling policy. Nodes are first filtered by the scheduling policy. Among the filtered nodes, the nodes that conform to policies, such as affinity settings, are selected. If no scheduling policy or affinity settings are configured, cluster-autoscaler selects a scaling group based on the least-waste policy. The least-waste policy selects the scaling group that has the fewest idle resources after simulation. If a scaling group of CPU-accelerated nodes and a scaling group of GPU-accelerated nodes both meet the requirements, the scaling group of CPU-accelerated nodes is selected by default.

- How can the success rate of auto scaling be increased? The success rate of auto scaling depends on the following factors:

- Whether the scheduling policy is met

After you configure a scaling group, you must be aware of the pod scheduling policies that the scaling group supports. If you are unaware of the pod scheduling policies, you can simulate a scaling activity by using the node selectors of pending pods and the labels of the scaling group.

- Whether resources are sufficient

After the scaling simulation is complete, a scaling group is selected. However, the scaling activity fails if the specified types of Elastic Compute Service (ECS) instances in the scaling group are out of stock. To increase the success rate of auto scaling, you can select different types of instances in more than one zone.

- How can auto scaling be accelerated?

- Method 1: Perform auto scaling in swift mode. After a scaling group experiences a scale-in activity and a scale-out activity, the swift mode is enabled for the scaling group.

- Method 2: Use custom images that are created from the base image of Alibaba Cloud Linux 2 (formerly known as Aliyun Linux 2). This ensures that the resources of Infrastructure as a Service (IaaS) are delivered 50% faster.

Considerations

- For each account, the default CPU quota for pay-as-you-go instances in each region is 50 vCPUs. You can add at most 48 custom route entries to each route table of a virtual private cloud (VPC). To request a quota increase, submit a ticket.
- The stock of ECS instances may be insufficient for auto scaling if you specify only one ECS instance type for a scaling group. We recommend that you specify multiple ECS instance types with the same specification for a scaling group. This increases the success rate of auto scaling.
- In swift mode, when a node is shut down and reclaimed, the node stops running and enters the *NotReady* state. When a scale-out activity is triggered, the state of the node changes to *Ready*.
- If a node is shut down and reclaimed in swift mode, you are charged only for the disks. This rule does not apply to nodes that use local disks, such as the instance type of `ecs.d1ne.2xlarge`, for which you are also charged a computing fee. If the stock of nodes is sufficient, nodes can be launched within a short period of time.
- If elastic IP addresses (EIPs) are associated with pods, we recommend that you do not delete the scaling group or remove ECS instances from the scaling group in the ECS console. Otherwise, these EIPs cannot be automatically released.

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the Clusters page, find the cluster that you want to manage and choose **More > Auto Scaling** in the **Actions** column.
4. On the **Configure Auto Scaling** page, set the following parameters and click **Submit**.

| Parameter | Description |
|------------------------|---|
| Cluster | The name of the cluster for which you want to enable auto scaling. |
| Scale-in Threshold | For a scaling group that is managed by cluster-autoscaler, set the value to the ratio of the requested resources per node to the total resources per node. If the actual value is lower than the threshold, the node is removed from the cluster. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p>Note In auto scaling, a scale-out activity is automatically triggered based on node scheduling. Therefore, you need to set only scale-in parameters.</p> </div> |
| GPU Scale-in Threshold | The scale-in threshold for GPU-accelerated instances. If the actual value is lower than the threshold, the node is removed from the cluster. |
| Defer Scale-in For | The amount of time that the cluster must wait before the cluster scales in. The default value is 10 minutes. |
| Cooldown | The cooldown period after a scale-in activity is triggered. No scale-in activity is triggered during the cooldown period. The default value is 10 minutes. |

5. Click **Create Scaling Group** and specify the type of resource for auto scaling based on your business requirements. Regular instances and GPU-accelerated instances are supported.
6. In the **Auto Scaling Group Configuration** dialog box, set the following parameters.

| Parameter | Description |
|-----------|---|
| Region | The region where you want to deploy the scaling group. The scaling group and the Kubernetes cluster must be deployed in the same region. You cannot change the region after the scaling group is created. |
| VPC | The scaling group and the Kubernetes cluster must be deployed in the same VPC. |
| VSwitch | The vSwitches of the scaling group. You can specify vSwitches of different zones. The vSwitches allocate pod CIDR blocks to the scaling group. |

7. Configure worker nodes.

| Parameter | Description |
|-----------------|---|
| Instance Type | The instance types in the scaling group. |
| Selected Types | The instance types that you select. You can select at most 10 instance types. |
| System Disk | The system disk of the scaling group. |
| Mount Data Disk | Specify whether to mount data disks to the scaling group. By default, no data disk is mounted. |
| Instances | <p>The number of instances contained in the scaling group.</p> <div style="background-color: #e6f2ff; padding: 10px; border: 1px solid #d9e1f2;"> <p> Note</p> <ul style="list-style-type: none"> ◦ Existing instances in the cluster are excluded. ◦ By default, the minimum number of instances is 0. If you specify one or more instances, the system adds the instances to the scaling group. When a scale-out activity is triggered, the instances in the scaling group are added to the cluster to which the scaling group is bound. </div> |
| Password | <p>Use a password.</p> <ul style="list-style-type: none"> ◦ Password: Enter the password that is used to log on to the nodes. ◦ Confirm Password: Enter the password again. |
| Scaling Mode | You can select Standard or Swift . |
| RDS Whitelist | The ApsaraDB RDS instances that can be accessed by the nodes in the scaling group after a scaling activity is triggered. |
| Label | Labels are automatically added to nodes that are added to the cluster by scale-out activities. |
| ECS Label | You can add labels to the selected ECS instances. |
| Taints | After you add taints to a node, Container Service no longer schedules pods to the node. |

| Parameter | Description |
|------------|--|
| CPU Policy | Specify the CPU policy. Valid values: <ul style="list-style-type: none"> None: indicates that the default CPU affinity is used. This is the default policy. Static: allows pods with specific resource characteristics on the node to be granted with enhanced CPU affinity and exclusivity. |

8. Set advanced options.

| Parameter | Description |
|-----------------------|--|
| Custom Security Group | Set a custom security group. |
| Custom Image | You can select a custom image. Then, all nodes in the Kubernetes cluster are deployed based on the image. |
| User Data | Customize the startup behaviors of ECS instances and import data to the ECS instances. The user data can be used to perform the following operations: <ul style="list-style-type: none"> Run scripts during instance startup. Import user data as normal data to an ECS instance for future reference. |

9. Click OK.

Check the results

In the left-side navigation pane, choose **Applications > Deployments**, select the kube-system namespace. You can find the cluster-autoscaler component. This indicates that the scaling group is created.

FAQ

- Why does the auto scaling component fail to add nodes after a scale-out activity is triggered?

Check whether the following situations exist:

- The instance types in the scaling group cannot fulfill the resource request from pods. By default, system components are installed for each node. Therefore, the requested pod resources must be less than the resource capacity of the instance type.
- The Resource Access Management (RAM) role does not have the permissions to manage the Kubernetes cluster. You must complete the authorization for each Kubernetes cluster that is involved in the scale-out activity.
- The Kubernetes cluster cannot connect to the Internet. The auto scaling component must call Alibaba Cloud API operations. Therefore, the nodes must have access to the Internet.

- Why does the auto scaling component fail to remove nodes after a scale-in activity is triggered?

Check whether the following situations exist:

- The requested resource threshold of each pod is higher than the configured scale-in threshold.
- Pods that belong to the *kube-system* namespace are running on the node.
- A scheduling policy forces the pods to run on the current node. Therefore, the pods cannot be scheduled to other nodes.
- `PodDisruptionBudget` is set for the pods on the node and the minimum value of `PodDisruptionBudget` is reached.

For more information about FAQ, see [open source component](#).

- How does the system choose a scaling group for a scaling activity?

When pods cannot be scheduled to nodes, the auto scaling component simulates the scheduling of the pods based on the configuration of scaling groups. The configuration includes labels, taints, and instance specifications. If a scaling group can simulate the scheduling of the pods, this scaling group is selected for the scale-out activity. If more than one scaling groups meet the requirements, the system selects the scaling group that has the fewest idle resources after simulation.

3.6.11.2. Horizontal pod autoscaling

You can create an application that has Horizontal Pod Autoscaling (HPA) enabled in the Container Service console. HPA can automatically scale container resources for your application. You can also use a YAML file to describe HPA settings.

Create an application that has HPA enabled in the Container Service console

Container Service provided by Alibaba Cloud is integrated with HPA. You can create an application that has HPA enabled in the Container Service console. You can enable HPA when you create an application or after the application is created.

Enable HPA when you create an application

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click the name or click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Workloads > Deployments**.
5. On the **Deployments** tab, click **Create from Image**.
6. On the **Basic Information** wizard page, enter a name for your application, set other required parameters, and then click **Next**.

| Parameter | Description |
|----------------------|--|
| Name | Enter a name for the application. |
| Replicas | The number of pods that are provisioned for the application. Default value: 2. |
| Type | The type of the application. You can select Deployments , StatefulSets , Jobs , Cron Jobs , or DaemonSets . |
| Label | Add a label to the application. The label is used to identify the application. |
| Annotations | Add an annotation to the application. |
| Synchronize Timezone | Specify whether to synchronize the time zone between nodes and containers. |

7. On the **Container** wizard page, set the container parameters, select an image, and then configure the required computing resources. Click **Next**. For more information, see [Configure the containers](#).

 **Note** You must configure the required computing resources for the Deployment. Otherwise, you cannot enable HPA.

8. On the **Advanced** wizard page, find the **Access Control** section, click **Create** on the right side of Services, and then set the parameters. For more information, see [Create an application from an image](#).

9. On the **Advanced** wizard page, select **Enable** for **HPA** and configure the scaling threshold and related settings.
 - o **Metric**: Select CPU Usage or Memory Usage. The selected resource type must be the same as the one that you have specified in the Required Resources field.
 - o **Condition**: Specify the resource usage threshold. HPA triggers scaling activities when the threshold is exceeded.
 - o **Max. Replicas**: Specify the maximum number of pods to which the Deployment can be scaled.
 - o **Min. Replicas**: Specify the minimum number of pods that must run for the Deployment.
10. In the lower-right corner of the Advanced wizard page, click **Create**. The application is created with HPA enabled.

Verify the result

- i. Click **View Details** or choose **Workloads > Deployments**. On the page that appears, click the **name of the created application** or click **Details** in the **Actions** column. Then, click the **Horizontal Pod Autoscaler** tab to view information about the scaling group of the application.
- ii. After the application starts to run, container resources are automatically scaled based on the CPU utilization. You can also check whether HPA is enabled in the staging environment by performing a CPU stress test on the pods of the application. Verify that the pods are automatically scaled within 30 seconds.

Enable HPA after an application is created

This example describes how to enable HPA for a stateless application.

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage, and click the name of the cluster or click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Workloads > Deployments**.
5. On the **Deployments** page, click the name of the application that you want to manage.
6. Click the **Pod Scaling** tab and click **Create**.
7. In the **Create** dialog box, configure the HPA settings. For more information about how to set the parameters, see [HPA settings](#) in Step 9.
8. Click **OK**.

Create an application that has HPA enabled by using kubectl

You can also create a Horizontal Pod Autoscaler by using an orchestration template and associate the Horizontal Pod Autoscaler with the Deployment for which you want to enable HPA. Then, you can run **kubectl** commands to enable HPA.

In the following example, HPA is enabled for an NGINX application.

1. Create a file named *nginx.yml* and copy the following content into the file.

The following code block is a YAML template that is used to create a Deployment:

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: nginx
  labels:
    app: nginx
spec:
  replicas: 2
  selector:
    matchLabels:
      app: nginx
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
      - name: nginx
        image: nginx:1.7.9 # replace it with your exactly <image_name:tags>
        ports:
        - containerPort: 80
      resources:
        requests:          ##To enable HPA, you must set this parameter.
        cpu: 500m

```

2. Run the following command to create an NGINX application:

```
kubectl create -f nginx.yml
```

3. Create a Horizontal Pod Autoscaler.

Use `scaleTargetRef` to associate the Horizontal Pod Autoscaler with the Deployment named `nginx`.

```

apiVersion: autoscaling/v2beta1
kind: HorizontalPodAutoscaler
metadata:
  name: nginx-hpa
  namespace: default
spec:
  scaleTargetRef:          ##Associate the Horizontal Pod Autoscaler with the nginx Deployment.
    apiVersion: apps/v1
    kind: Deployment
    name: nginx
  minReplicas: 1
  maxReplicas: 10
  metrics:
  - type: Resource
    resource:
      name: cpu
      targetAverageUtilization: 50

```

 **Note** You must configure the requested resources for the pods of the application. Otherwise, the Horizontal Pod Autoscaler cannot be started.

4. Run the `kubectl describe hpa name` command. The following output is an example of a warning that is returned:

```
Warning FailedGetResourceMetric 2m (x6 over 4m) horizontal-pod-autoscaler missing request for cpu on container nginx in pod default/nginx-deployment-basic-75675f5897-mqzs7
Warning FailedComputeMetricsReplicas 2m (x6 over 4m) horizontal-pod-autoscaler failed to get cpu utilization: missing request for cpu on container nginx in pod default/nginx-deployment-basic-75675f5
```

- 5. After the Horizontal Pod Autoscaler is created, run the `kubectl describe hpa name` command. If the following output is returned, it indicates that the Horizontal Pod Autoscaler is running as expected:

```
Normal SuccessfulRescale 39s horizontal-pod-autoscaler New size: 1; reason: All metrics below target
```

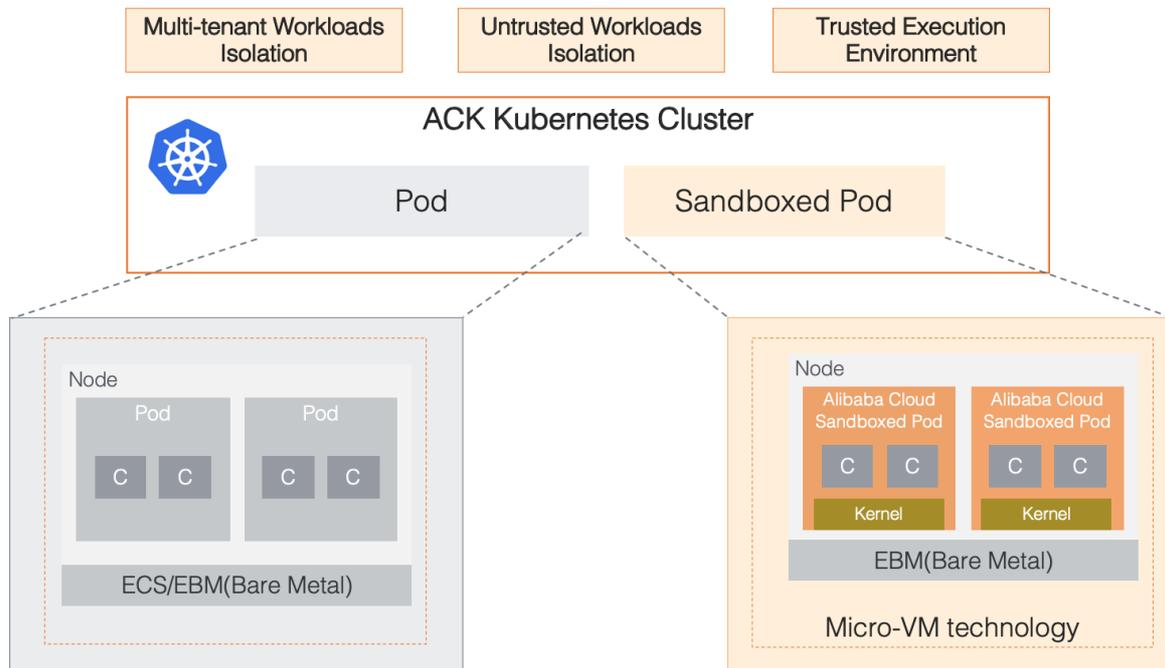
If the CPU utilization of the NGINX application pod exceeds 50% as specified in the HPA settings, the Horizontal Pod Autoscaler automatically adds pods. If the CPU utilization of the NGINX application pod drops below 50%, the Horizontal Pod Autoscaler automatically removes pods.

3.6.12. Sandboxed-containers

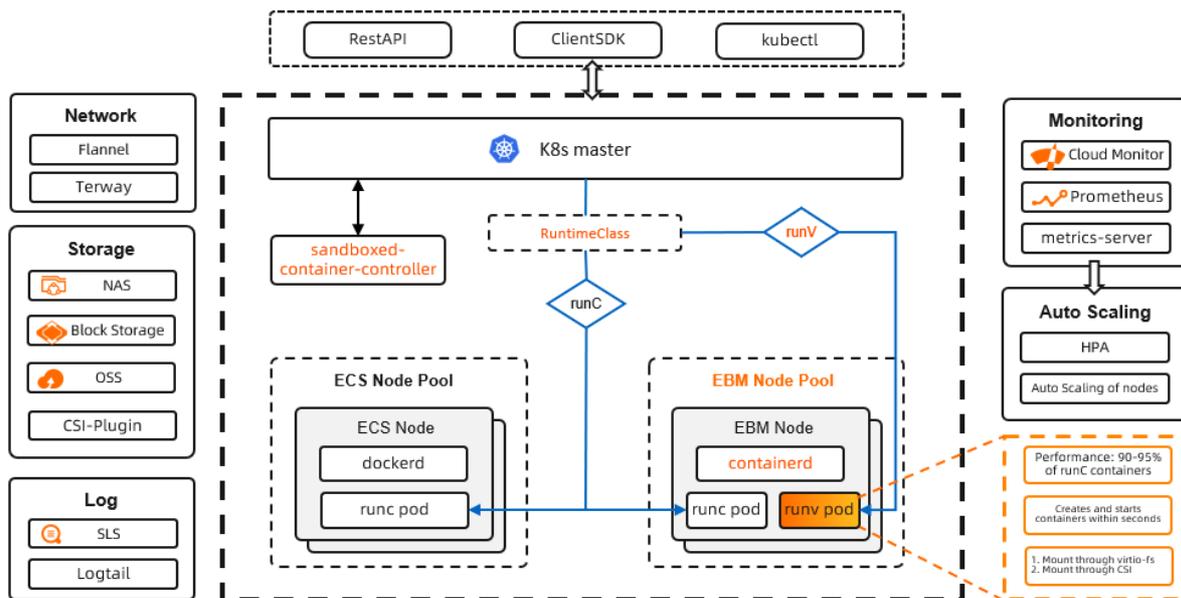
3.6.12.1. Overview

Sandboxed-Container is an alternative to the Docker runtime. Sandboxed-Container allows you to run applications in a sandboxed and lightweight virtual machine that has a dedicated kernel. This enhances resource isolation and improves security.

Sandboxed-Container is applicable to scenarios such as untrusted application isolation, fault isolation, performance isolation, and workload isolation among multiple users. Sandboxed-Container provides higher security. Sandboxed-Container has minor impacts on application performance and offers the same user experience as Docker in terms of logging, monitoring, and elastic scaling.



Architecture



Features

Sandboxed-Container is a container-securing runtime that is developed by Alibaba Cloud based on sandboxed and lightweight virtual machines. Compared with Sandboxed-Container V1, Sandboxed-Container V2 maintains the same isolation performance and reduces the pod overhead by 90%. It also allows you to start sandboxed containers 3 times faster and increases the maximum number of pods that can be deployed on a host by 10 times. Sandboxed-Container V2 provides the following key features:

- Strong isolation based on sandboxed and lightweight virtual machines.
- Good compatibility with runC in terms of application management.
- Network Attached Storage (NAS) file systems, disks, and Object Storage Service (OSS) buckets can be mounted both directly and through virtio-fs.
- The same user experience as runC in terms of logging, monitoring, and storage.
- Supports RuntimeClass (runC and runV). For more information, see [RuntimeClass](#).
- Easy to use with minimum requirements on technical skills.
- Higher stability than Kata Containers. For more information about Kata Containers, see [Kata Containers](#).

3.6.12.2. Create a Kubernetes cluster that runs sandboxed containers

This topic describes how to create a Kubernetes cluster that runs sandboxed containers in the Container Service console.

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane of the Container Service console, click **Clusters**. On the Clusters page that appears, click **Create Kubernetes Cluster** in the upper-right corner.
3. On the **Create Cluster** page, set basic configurations for the cluster.

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|-----------------------|---|
| Cluster Name | <p>Enter a name for the cluster. The name must be 1 to 63 characters in length, and can contain digits, letters, and hyphens (-).</p> <p>Note The cluster name must be unique among clusters that belong to the same Alibaba Cloud account.</p> |
| Region | Select the region where you want to deploy the Kubernetes cluster. |
| VPC | <p>You can select a virtual private cloud (VPC) from the drop-down list.</p> <ul style="list-style-type: none"> ◦ If the specified VPC is already associated with a NAT gateway, the cluster uses this NAT gateway. ◦ Otherwise, the system automatically creates a NAT gateway. If you do not want the system to create a NAT gateway, clear Configure SNAT for VPC. <p>Note If you disallow the system to automatically create a NAT gateway and want the VPC to access the Internet, you must manually associate the VPC with a NAT gateway or create Source Network Address Translation (SNAT) rules for the VPC.</p> |
| VSwitch | <p>Select one or more vSwitches for the cluster.</p> <p>You can select up to three vSwitches that are deployed in different zones.</p> |
| Kubernetes Version | Select a Kubernetes version. |
| Container Runtime | Select a runtime for the Kubernetes cluster. |
| Billing Method | Only pay-as-you-go nodes are supported. |
| Master Configurations | <p>Set the Instance Type and System Disk parameters:</p> <ul style="list-style-type: none"> ◦ Master Node Quantity: You can add up to three master nodes. ◦ Instance Type: You can select one or more instance types. For more information, see the <i>Instance types</i> chapter of <i>ECS User Guide</i>. ◦ System Disk: Standard SSDs and ultra disks are supported. <p>Note You can select Enable Backup to back up disk data.</p> |
| Worker Instance | By default, Create Instance is selected. |

| Parameter | Description |
|---------------------------------|--|
| Worker Configurations | <p>If Worker Instance is set to Create Instance, set the following parameters:</p> <ul style="list-style-type: none"> Instance Type: Select Elastic Compute Service (ECS) bare metal instance types. Selected Types: The selected instance types are displayed. Quantity: Set the number of worker nodes. System Disk: Standard SSDs and ultra disks are supported. <p>Note You can select Enable Backup to back up disk data.</p> <ul style="list-style-type: none"> Mount Data Disk: Standard SSDs and ultra disks are supported. <p>Note You can enable disk encryption and data backup when you mount disks. The disks are used to store the root file systems of containers on the nodes. Therefore, you must mount a disk of at least 200 GiB. We recommend that you mount a disk of at least 1 TiB.</p> |
| Password | <p>Set a password that is used to log on to the nodes.</p> <p>Note The password must be 8 to 30 characters in length, and must contain at least three of the following types of character: uppercase letters, lowercase letters, digits, and special characters.</p> |
| Confirm Password | Enter the password again. |
| Network Plug-in | Flannel and Terway are supported. By default, Flannel is selected. |
| Pod CIDR Block and Service CIDR | <p>These parameters are optional. For more information, see <i>Network planning</i> in <i>VPC User Guide</i>.</p> <p>Note These parameters are available only when you select an existing VPC.</p> |
| Configure SNAT | This parameter is optional. If you clear Configure SNAT for VPC, you must create a NAT gateway or configure SNAT rules for the VPC. |
| Access to the Internet | <p>Specify whether to expose the API server with an elastic IP address (EIP). The Kubernetes API server provides multiple HTTP-based RESTful APIs that can be used to create, delete, modify, query, and watch resource objects such as pods and Services.</p> <ul style="list-style-type: none"> If you select this check box, an EIP is created and attached to an internal-facing Server Load Balancer (SLB) instance. Port 6443 used by the API server is exposed on the master nodes. You can connect to and manage the cluster by using kubeconfig over the Internet. If you clear this check box, no EIP is created. You can connect to and manage the cluster only by using kubeconfig from within the VPC. |
| Ingress | Specify whether to install Ingress controllers. By default, Install Ingress Controller is selected. |

| Parameter | Description |
|---------------------|--|
| Log Service | If you enable Log Service, you can select an existing project or create a project. If you select Enable Log Service , the Log Service plug-in is automatically installed in the cluster. If you select Create Ingress Dashboard , Ingress access logs are collected and displayed on dashboards. |
| Volume Plug-in | By default, CSI is selected. |
| Deletion Protection | If you select this check box, the cluster cannot be deleted in the console or by calling API operations. |
| RDS Whitelist | Add the IP addresses of the nodes to the whitelist of the ApsaraDB RDS instance that is allowed to access the Kubernetes cluster. Note To enable an ApsaraDB RDS instance to access the Kubernetes cluster, you must deploy the ApsaraDB RDS instance in the same VPC as the Kubernetes cluster. |
| Node Protection | This check box is selected by default to prevent nodes from being deleted in the console or by calling API operations. |
| Label | Add labels to the cluster. |

4. Complete the advanced settings of the cluster.

| Parameter | Description |
|-----------------------|--|
| IP Addresses per Node | The number of IP addresses that is assigned to a node. |
| Kube-proxy Mode | iptables and IPVS are supported. <ul style="list-style-type: none"> iptables is a mature and stable kube-proxy mode. It uses iptables rules to conduct service discovery and load balancing. The performance of this mode is restricted by the size of the Kubernetes cluster. This mode is suitable for Kubernetes clusters that manage a small number of Services. IPVS is a high-performance kube-proxy mode. It uses Linux Virtual Server (LVS) to conduct service discovery and load balancing. This mode is suitable for Kubernetes clusters that manage a large number of Services. We recommend that you use this mode in scenarios where high-performance load balancing is required. |
| Custom Node Name | Specify whether to use a custom node name. |
| Node Port Range | Specify the node port range. |
| Taints | Add taints to all worker nodes in the Kubernetes cluster. |
| CPU Policy | Specify the CPU policy. Valid values: <ul style="list-style-type: none"> None: indicates that the default CPU affinity is used. This is the default policy. Static: allows pods with specific resource characteristics on the node to be granted with enhanced CPU affinity and exclusivity. |
| Cluster Domain | The default domain name of the cluster is cluster.local. You can specify a custom domain name. |

| Parameter | Description |
|------------|---|
| Cluster CA | Specify whether to enable the cluster CA certificate. |
| User Data | <p>You can customize the startup behaviors of ECS instances and import data to the ECS instances. The user data can be used to perform the following operations:</p> <ul style="list-style-type: none"> ◦ Run user data scripts during instance startup. ◦ Import user data as common data to an ECS instance for future reference. |

5. Click **Create Cluster** in the upper-right corner of the page.
6. On the **Confirm** page, click **OK** to start the deployment.

Result

After the cluster is created, you can find the cluster on the **Clusters** page in the Container Service console.

3.6.12.3. Expand a Container Service cluster that runs sandboxed containers

This topic describes how to scale out the worker nodes in a Container Service cluster that runs sandboxed containers in the Container Service console.

Prerequisites

You cannot scale out the master nodes in a Container Service cluster that runs sandboxed containers.

To expand a Container Service cluster that runs sandboxed containers, you must set the parameters as required in the following table. Otherwise, the added nodes cannot run sandboxed containers.

1. [Log on to the Container Service console.](#)
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Expand** in the **Actions** column.
4. Go to the **Expand** page and set the required parameters.

In this example, the number of worker nodes in the Container Service cluster is increased from three to five. The following table describes the required parameters.

| Parameter | Description |
|-------------------|--|
| Cluster Name | By default, the name of the Container Service cluster appears. |
| Region | The region where the Container Service cluster is deployed. |
| Container Runtime | By default, Sandboxed-Container appears. |

| Parameter | Description |
|----------------------------|---|
| VPC | <p>You can select a virtual private cloud (VPC) from the drop-down list.</p> <ul style="list-style-type: none"> ◦ If the specified VPC is already associated with a NAT gateway, the cluster uses this NAT gateway. ◦ Otherwise, the system automatically creates a NAT gateway. If you do not want the system to create a NAT gateway, clear Configure SNAT for VPC. <p> Note If you disallow the system to automatically create a NAT gateway and want the VPC to access the Internet, you must manually associate the VPC with a NAT gateway or create Source Network Address Translation (SNAT) rules for the VPC.</p> |
| VSwitch | <p>Select one or more vSwitches for the cluster.</p> <p>You can select up to three vSwitches that are deployed in different zones.</p> |
| Billing Method | Only pay-as-you-go nodes are supported. |
| Existing Worker Nodes | The number of existing workers in the Container Service cluster. |
| Nodes to Add | Set the number of worker nodes to add. |
| Worker Nodes After Scaling | The number of worker nodes after the scaling. |
| Instance Type | Select ECS Bare Metal Instance. |
| Selected Types | The selected instance types are displayed. |
| System Disk | <p>Standard SSDs and ultra disks are supported.</p> <p> Note You can select Enable Backup to back up disk data.</p> |
| Mount Data Disk | <p>Standard SSDs and ultra disks are supported.</p> <p> Note You can enable disk encryption and data backup when you mount disks. The disks are used to store the root file systems of containers on the nodes. Therefore, you must mount a disk of at least 200 GiB. We recommend that you mount a disk of at least 1 TiB.</p> |
| Password | <ul style="list-style-type: none"> ◦ Password: Enter the password that is used to log on to the nodes. ◦ Confirm Password: Enter the password again. |
| RDS Whitelist | Set the Apsara RDS whitelist. Add the IP addresses of the nodes in the cluster to the RDS whitelist. |
| Label | Add labels to the cluster. |
| Taints | Add taints to all worker nodes in the Kubernetes cluster. |

| Parameter | Description |
|------------|--|
| CPU Policy | Specify the CPU policy. Valid values: <ul style="list-style-type: none"> None: indicates that the default CPU affinity is used. This is the default policy. Static: allows pods with specific resource characteristics on the node to be granted with enhanced CPU affinity and exclusivity. |
| User Data | You can customize the startup behaviors of ECS instances and import data to the ECS instances. The user data can be used to perform the following operations: <ul style="list-style-type: none"> Run user data scripts during instance startup. Import user data as common data to an ECS instance for future reference. |

5. Click **Submit**.

What's next

After the Container Service cluster is expanded, go to the details page of the Container Service cluster. In the left-side navigation pane, choose **Clusters > Node Pools**. You can find that the number of worker nodes is increased from 3 to 5.

3.6.12.4. Create an application that runs in sandboxed containers

This topic describes how to use an image to create an NGINX application that runs in sandboxed containers. The NGINX application is accessible over the Internet.

Prerequisites

A cluster that contains sandboxed containers is created. For more information, see [Create a Kubernetes cluster that supports sandboxed containers](#).

Procedure

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Workloads > Deployments**.
5. On the **Deployments** page, select the namespace and click **Create from Image** in the upper-right corner.
6. On the **Basic Information** wizard page, specify the basic information of the application and click **Next**.

Set **Container Runtime** to **runv**. Set the following parameters: **Name**, **Replicas**, **Type**, **Label**, and **Annotations**. Select whether you want to enable **Synchronize Timezone**. The number of replicas specifies the number of pods that are provisioned for in the application.

 **Note**

Deployments is selected in this example.

7. Configure containers.

 **Note** In the upper part of the **Container** wizard page, click **Add Container** to add more containers for the application.

The following table describes the parameters that are required to configure the containers.

o General settings

| Parameter | Description |
|---------------------------|--|
| Image Name | <p>Click Select Image. In the dialog box that appears, select an image and click OK. In this example, an NGINX image is selected.</p> <p>You can also enter the address of an image stored in a private registry. The image address must be in the following format: <code>domainname/namespace/imageName:tag</code>.</p> |
| Image Version | <ul style="list-style-type: none"> ▪ Click Select Image Version and select an image version. If you do not specify an image version, the latest image version is used. ▪ You can select the following image pulling policies: <ul style="list-style-type: none"> ▪ ifNotPresent: If the image that you want to pull is found in the region where the cluster is deployed, Container Service uses the local image. Otherwise, Container Service pulls the image from the corresponding repository. ▪ Always: Container Service pulls the image from the repository each time the application is deployed or expanded. ▪ Never: Container Service uses only images on your on-premise machine. <p> Note If you select Image Pull Policy, no image pulling policy is applied.</p> <ul style="list-style-type: none"> ▪ To pull the image without a Secret, click Set Image Pull Secret to set a Secret for pulling images. |
| Resource Limit | You can specify an upper limit for the CPU, memory, and ephemeral storage space that the container can consume. This prevents the container from occupying an excessive amount of resources. The CPU resource is measured in milicores (one thousandth of one core). The memory resource is measured in MiB. The ephemeral storage resource is measured in GiB. |
| Required Resources | The amount of CPU and memory resources that are reserved for this application. These resources are exclusive to the container. This prevents the application from becoming unavailable if other services or processes compete for computing resources. |
| Container Start Parameter | <ul style="list-style-type: none"> ▪ <code>stdin</code>: Pass stdin to the container. ▪ <code>tty</code>: Stdin is a TeleTYpewriter (TTY). |
| Privileged Container | <ul style="list-style-type: none"> ▪ If you select Privileged Container, <code>privileged=true</code> is set for the container and the privilege mode is enabled. ▪ If you do not select Privileged Container, <code>privileged=false</code> is set for the container and the privilege mode is disabled. |
| Init Container | If you select Init Container, an init container is created. An init container provides tools to manage pods. For more information, see Init Containers . |

◦ (Optional)Ports

Configure container ports.

- Name: Enter a name for the port.
- Container Port: Enter the container port that you want to open. Enter a port number from 1 to 65535.
- Protocol: Select TCP or UDP.

◦ (Optional)Environments

You can configure environment variables for pods in key-value pairs. Environment variables are used to apply pod configurations to containers. For more information, see [Pod variables](#).

- Type: Select the type of the environment variable. You can select **Custom**, **ConfigMaps**, **Secrets**, **Value/ValueFrom**, or **ResourceFieldRef**. If you select ConfigMaps or Secret as the type of the environment variable, all values in the selected ConfigMap or Secret are passed to the container environment variables. In this example, Secret is selected.

Select **Secrets** from the Type drop-down list and select a Secret from the **Value/ValueFrom** drop-down list. All values in the selected Secret are passed to the environment variable.

| Type | Variable Key | Value/ValueFrom |
|--------|--------------|-----------------|
| Secret | e.g. foo | |

In this case, the YAML file that is used to deploy the application contains the settings that reference all data in the selected Secret.

```
envFrom:
  - secretRef:
      name: test
```

- Variable Key: Specify the name of the environment variable.
- Value/ValueFrom: Specify the value that is referenced by the environment variable.

◦ (Optional)Health Check

Health check settings include liveness and readiness probes. Liveness probes determine when to restart the container. Readiness probes determine whether the container is ready to accept network traffic. For more information about health checks, see [Configure Liveness, Readiness, and Startup Probes](#).

| Request type | Description |
|--------------|-------------|
|--------------|-------------|

| Request type | Description |
|--------------|---|
| HTTP | <p>Sends an HTTP GET request to the container. You can configure the following parameters:</p> <ul style="list-style-type: none"> ■ Protocol: HTTP or HTTPS. ■ Path: the requested path on the server. ■ Port: Enter the container port that you want to open. Enter a port number from 1 to 65535. ■ HTTP Header: Enter the custom headers in the HTTP request. Duplicate headers are allowed. Key-value pairs are supported. ■ Initial Delay (s): the initialDelaySeconds field in the YAML file. This field specifies the time (in seconds) that the system must wait before it can send a probe to the container after the container is started. Default value: 3. ■ Period (s): the periodSeconds field in the YAML file. This field specifies the interval (in seconds) at which probes are performed. Default value: 10. Minimum value: 1. ■ Timeout (s): the timeoutSeconds field in the YAML file. This field specifies the time (in seconds) after which a probe times out. Default value: 1. Minimum value: 1. ■ Healthy Threshold: the minimum number of times that an unhealthy container must consecutively pass health checks before it is considered healthy. Default value: 1. Minimum value: 1. For liveness probes, this parameter must be set to 1. ■ Unhealthy Threshold: the minimum number of times that a healthy container must consecutively fail health checks before it is considered unhealthy. Default value: 3. Minimum value: 1. |
| TCP | <p>Sends a TCP socket to the container. kubelet attempts to open the socket on the specified port. If the connection can be established, the container is considered healthy. Otherwise, the container is considered unhealthy. You can set the following parameters:</p> <ul style="list-style-type: none"> ■ Port: Enter the container port that you want to open. Enter a port number from 1 to 65535. ■ Initial Delay (s): the initialDelaySeconds field in the YAML file. This field specifies the time (in seconds) that the system must wait before it can send a probe to the container after the container is started. Default value: 15. ■ Period (s): the periodSeconds field in the YAML file. This field specifies the time interval (in seconds) at which probes are performed. Default value: 10. Minimum value: 1. ■ Timeout (s): the timeoutSeconds field in the YAML file. This field specifies the time (in seconds) after which a probe times out. Default value: 1. Minimum value: 1. ■ Healthy Threshold: the minimum number of times that an unhealthy container must consecutively pass health checks before it is considered healthy. Default value: 1. Minimum value: 1. For liveness probes, this parameter must be set to 1. ■ Unhealthy Threshold: the minimum number of times that a healthy container must consecutively fail health checks before it is considered unhealthy. Default value: 3. Minimum value: 1. |

| Request type | Description |
|--------------|---|
| Command | <p>Runs a probe command in the container to check the health status of the container. You can set the following parameters:</p> <ul style="list-style-type: none"> ■ Command: the probe command that is run to check the health status of the container. ■ Initial Delay (s): the <code>initialDelaySeconds</code> field in the YAML file. This field specifies the time (in seconds) that the system must wait before it can send a probe to the container after the container is started. Default value: 5. ■ Period (s): the <code>periodSeconds</code> field in the YAML file. This field specifies the interval (in seconds) at which probes are performed. Default value: 10. Minimum value: 1. ■ Timeout (s): the <code>timeoutSeconds</code> field in the YAML file. This field specifies the time (in seconds) after which a probe times out. Default value: 1. Minimum value: 1. ■ Healthy Threshold: the minimum number of times that an unhealthy container must consecutively pass health checks before it is considered healthy. Default value: 1. Minimum value: 1. For liveness probes, this parameter must be set to 1. ■ Unhealthy Threshold: the minimum number of times that a healthy container must consecutively fail health checks before it is considered unhealthy. Default value: 3. Minimum value: 1. |

○ Lifecycle

You can set the following parameters to configure the lifecycle of the container: Start, Post Start, and Pre Stop. For more information, see [Configure the lifecycle of a container](#).

- **Start:** Set the command and parameter that take effect before the container starts.
- **Post Start:** Set the command that takes effect after the container starts.
- **Pre Stop:** Set the command that takes effect before the container stops.

○ (Optional)Volume

You can mount local volumes and persistent volume claims (PVCs) to the container.

- **Add Local Storage:** You can select `HostPath`, `ConfigMap`, `Secret`, and `EmptyDir`. The source directory or file is mounted to a path in the container. For more information, see [Volumes](#).
- **Add PVC:** Cloud Storage is supported.

In this example, a PVC named `disk-ssd` is mounted to the `/tmp` path of the container.

○ (Optional)Log

Configure **Log Service**. You can specify log collection configurations and add custom tags.

 **Notice** Make sure that the Log Service agent is installed in the cluster.

| Parameter | Description |
|-----------|--|
| | Logstore: creates a Logstore in Log Service to store collected log data. |

| Parameter Collection Configuration | Description |
|------------------------------------|--|
| | <p>Log Path in Container: specifies stdout or a path to collect log data</p> <ul style="list-style-type: none"> ■ stdout: specifies that the stdout files are collected. ■ Text Logs: specifies that log data in the specified path of the container is collected. In this example, <code>/var/log/nginx</code> is specified as the path. Wildcard characters can be used to specify the path. |
| Custom Tag | You can also add custom tags. Custom tags are added to the log data of the container when the log data is collected. Log data with tags is easier to aggregate and filter. |

8. Configure the parameters based on your business requirements and click **Next**.
9. (Optional) Configure advanced settings.
 - o Access Control

Note

You can configure the following access control settings based on your business requirements:

- **Internal applications**: For applications that run inside the cluster, you can create a ClusterIP or NodePort Service to enable internal communication.
- **External applications**: For applications that are open to the Internet, you can configure access control by using one of the following methods:
 - Create a LoadBalancer Service and enable access to your application over the Internet by using a Server Load Balancer (SLB) instance.
 - Create an Ingress and use the Ingress to expose your application to the Internet. For more information, see [Ingress](#).

You can also specify how the backend pods are exposed to the Internet. In this example, a ClusterIP Service and an Ingress are created to expose the NGINX application to the Internet.

| Parameter | Description |
|-----------|---|
| Services | Click Create on the right side of Services . In the Create Service dialog box, set the parameters. Select Cluster IP . |
| Ingresses | <p>Click Create on the right side of Ingresses. In the Create dialog box, set the parameters.</p> <p>Note When you create an application from an image, you can create an Ingress only for one Service. In this example, a virtual hostname is used as the test domain name. You must add the following entry to the hosts file to map the domain name to the IP address of the Ingress. In actual scenarios, use a domain name that has obtained an Internet Content Provider (ICP) number.</p> <pre>101.37.224.146 foo.bar.com #The IP address of the Ingress.</pre> |

You can find the created Service and Ingress in the **Access Control** section. You can click **Update** or **Delete** to change the configurations.

o Scaling

Specify whether to enable HPA to automatically scale the number of pods based on the CPU and memory usage. This enables the application to run smoothly at different load levels.

Note To enable HPA, you must configure required resources for the container. Otherwise, HPA does not take effect.

- **Metric:** Select CPU Usage or Memory Usage. The selected resource type must be the same as the one you have specified in the Required Resources field.
- **Condition:** Specify the resource usage threshold. HPA triggers scaling events when the threshold is exceeded.
- **Max. Replicas:** Specify the maximum number of replicated pods to which the application can be scaled.
- **Min. Replicas:** Specify the minimum number of replicated pods that must run.

o Scheduling

You can set the following parameters: Update Method, Node Affinity, Pod Affinity, Pod Anti Affinity, and Toleration. For more information, see [Affinity and anti-affinity](#).

Note Node affinity and pod affinity affect pod scheduling based on node labels and pod labels. You can add node labels and pod labels that are provided by Kubernetes to configure node affinity and pod affinity. You can also add custom labels to nodes and pods, and then configure node affinity and pod affinity based on these custom labels.

| Parameter | Description |
|---------------|--|
| Update Method | Select Rolling Update or OnDelete. For more information, see Deployments . |
| Node Affinity | <p>Set Node Affinity by adding labels to worker nodes.</p> <p>Node Affinity supports required and preferred rules, and various operators, such as In, NotIn, Exists, DoesNotExist, Gt, and Lt.</p> <ul style="list-style-type: none"> ■ Required: Specify the rules that must be matched for pod scheduling. In the YAML file, these rules are defined by the <code>requiredDuringSchedulingIgnoredDuringExecution</code> field of the <code>nodeAffinity</code> parameter. These rules have the same effect as the Node Selector parameter. In this example, pods can be scheduled only to nodes with the specified labels. You can create multiple required rules. However, only one of them must be met. ■ Preferred: Specify the rules that are not required to be matched for pod scheduling. Pods are scheduled to a node that matches the preferred rules when multiple nodes match the required rules. In the YAML file, these rules are defined by the <code>preferredDuringSchedulingIgnoredDuringExecution</code> field of the <code>nodeAffinity</code> parameter. In this example, the scheduler attempts to schedule a pod to a node that matches the preferred rules. You can also set weights for preferred rules. If multiple nodes match the rule, the node with the highest weight is preferred. You can create multiple preferred rules. However, all of them must be met before the pod can be scheduled. |

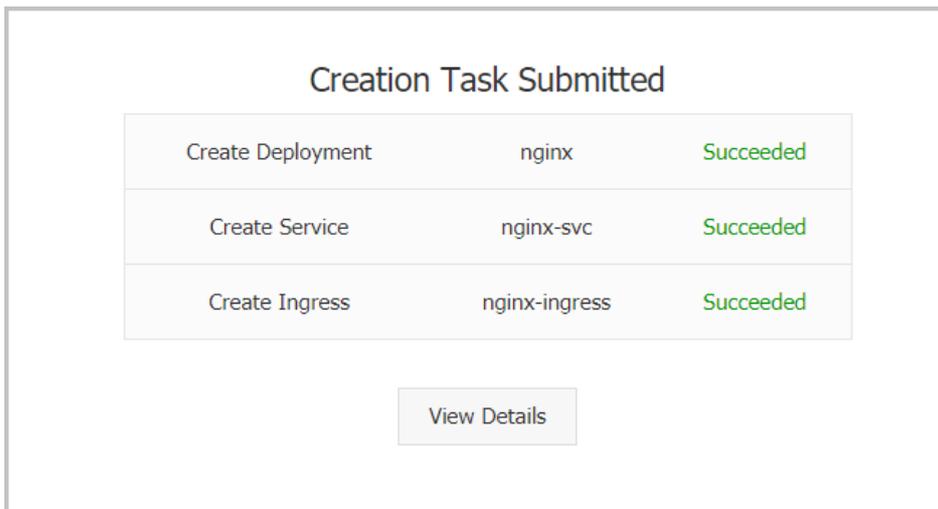
| Parameter | Description |
|--------------|--|
| Pod Affinity | <p>Pod affinity rules specify how pods are deployed relative to other pods in the same topology domain. For example, you can use pod affinity to deploy services that communicate with each other to the same topological domain, such as a host. This reduces the network latency between these services.</p> <p>Pod affinity enables you to select nodes to which pods can be scheduled based on the labels of other running pods. Pod affinity supports required and preferred rules, and the following operators: In, NotIn, Exists, and DoesNotExist.</p> <ul style="list-style-type: none"> ▪ Required: Specify rules that must be matched for pod scheduling. In the YAML file, these rules are defined by the <code>requiredDuringSchedulingIgnoredDuringExecution</code> field of the <code>podAffinity</code> parameter. A node must match the required rules before pods can be scheduled to the node. ▪ Namespace: Specify the namespace to apply the required rule. Pod affinity rules are defined based on the labels that are added to pods and therefore must be scoped to a namespace. ▪ Topological Domain: Set the <code>topologyKey</code>. This specifies the key for the node label that the system uses to denote the topological domain. For example, if you set the parameter to <code>kubernetes.io/hostname</code>, topologies are determined by nodes. If you set the parameter to <code>beta.kubernetes.io/os</code>, topologies are determined by the operating systems of nodes. ▪ Selector: Click Add to add pod labels. ▪ View Applications: Click View Applications and set the namespace and application in the dialog box that appears. You can view the pod labels on the selected application and add the labels as selectors. ▪ Required Rules: Specify labels on existing applications, the operator, and the label value. In this example, the required rule specifies that the application to be created is scheduled to a host that runs applications with the <code>app:nginx</code> label. ▪ Preferred: Specify rules that are not required to be matched for pod scheduling. In the YAML file, preferred rules are defined by the <code>preferredDuringSchedulingIgnoredDuringExecution</code> field of the <code>podAffinity</code> parameter. The scheduler attempts to schedule the pod to a node that matches the preferred rules. You can set weights for preferred rules. The other parameters are the same as those of required rules. <div style="background-color: #e0f2f7; padding: 10px; margin-top: 10px;"> <p> Note Weight: Set the weight of a preferred rule to a value from 1 to 100. The scheduler calculates the weight of each node that meets the preferred rule based on an algorithm, and then schedules the pod to the node with the highest weight.</p> </div> |

| Parameter | Description |
|---------------------------|--|
| Pod Anti Affinity | <p>Pod anti-affinity rules specify that pods are not scheduled to topological domains where pods with matching labels are deployed. Pod anti-affinity rules apply to the following scenarios:</p> <ul style="list-style-type: none"> ▪ Schedule the pods of an application to different topological domains, such as multiple hosts. This allows you to enhance the stability of the service. ▪ Grant a pod exclusive access to a node. This enables resource isolation and ensures that no other pod can share the resources of the specified node. ▪ Schedule pods of an application to different hosts if the pods may interfere with each other. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p>? Note The parameters of pod anti-affinity rules are the same as those of pod affinity rules. You can create the rules for different scenarios.</p> </div> |
| Toleration | Configure toleration rules to allow pods to be scheduled to nodes with matching taints. |
| Schedule to Virtual Nodes | Specify whether to schedule pods to virtual nodes. This option is unavailable if the cluster does not contain a virtual node. |

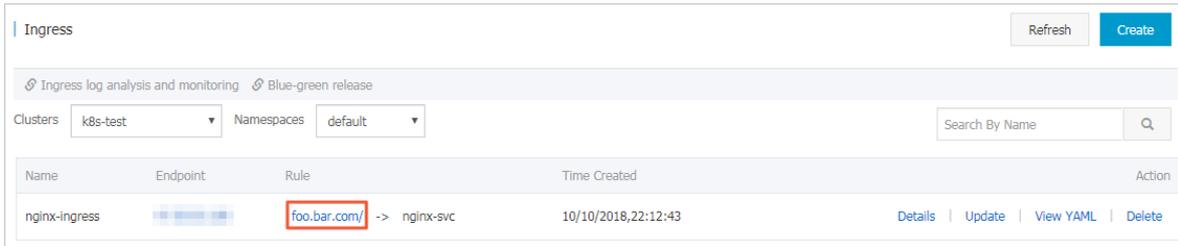
- Labels and Annotations
 - Pod Labels: Add a label to the pod. The label is used to identify the application.
 - Pod Annotations: Add an annotation to the pod.

10. Click **Create**.

After the application is deployed, you are redirected to the Complete wizard page. The resource objects of the application are displayed. You can find the resource objects under the application and click **View Details** to view application details.



11. In the left-side navigation pane, choose **Ingresses and Load Balancing > Ingresses**. The created Ingress rule is displayed on the page.



Result

Enter the test domain in the address bar of your browser and press Enter. The NGINX welcome page appears.



3.6.12.5. Configure a Kubernetes cluster that runs both sandboxed and Docker containers

Node pools support multiple types of container runtime. However, nodes in the same node pool must use the same type of container runtime. Container Service allows you to create node pools of different container runtime types for a cluster. This topic describes how to create a node pool that runs sandboxed containers and a node pool that runs Docker containers for a Kubernetes cluster.

Prerequisites

A Kubernetes cluster is created. For more information, see [Create a Kubernetes cluster](#).

Notice The Kubernetes cluster must meet the following requirements:

- The cluster version must be 1.14.6-aliyun.1 or later.
- The network plug-in must be Flannel or Terway. Terway must run in One ENI for Multi-Pod mode.
- The volume plug-in must be CSI-Plugin 1.14.8.39-0d749258-aliyun or later. Flexvolume is not supported.
- The logtail-ds version must be 0.16.34.2-f6647154-aliyun or later.

Considerations

- By default, a cluster can contain at most 100 nodes.
- Before you add an existing Elastic Compute Service (ECS) instance that is deployed in a virtual private cloud (VPC), make sure that an elastic IP address (EIP) is associated with the ECS instance, or a NAT gateway is created in the VPC. In addition, the nodes that you want to add to the node pool must have access to the Internet. Otherwise, the ECS instance cannot be added.

Create a node pool that runs Docker containers

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.

4. In the left-side navigation pane of the details page, choose **Nodes > Node Pools**.
5. On the **Node Pools** page, click **Create Node Pool** and set the parameters.

For more information, see [Create a Kubernetes cluster](#). The following table describes the parameters.

| Parameter | Description |
|-----------------------|--|
| Name | Enter a name for the node pool. |
| Container Runtime | Select Docker. This specifies that all containers in the node pool are Docker containers. |
| Quantity | Specify the initial number of nodes in the node pool. If you do not need to create nodes, set this parameter to 0. |
| Operating System | The CentOS and Aliyun Linux operating systems are supported. |
| ECS Label | You can add labels to the ECS instances. |
| Node Label | You can add labels to the nodes in the cluster. |
| Custom Resource Group | Specify the resource group of the nodes to be added to the node pool. |
| Custom Security Group | Select a custom security group. |

6. Click **OK**.

Create a node pool that runs sandboxed containers

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Nodes > Node Pools**.
5. On the **Node Pools** page, click **Create Node Pool** and set the parameters.

For more information, see [Create a Kubernetes cluster](#). The following table describes the parameters.

| Parameter | Description |
|-------------------|--|
| Name | Enter a name for the node pool. |
| Container Runtime | Select Sandboxed-Container. This specifies that all containers in the node pool are sandboxed containers. |
| Quantity | Specify the initial number of nodes in the node pool. If you do not need to create nodes, set this parameter to 0. |
| Operating System | Select an operating system. Sandboxed containers support only the Aliyun Linux operating system. |
| Mount Data Disk | You must mount a disk of at least 200 GiB. |
| ECS Label | You can add labels to the ECS instances. |
| Node Label | You can add labels to the nodes in the cluster. |

| Parameter | Description |
|-----------------------|---|
| Custom Resource Group | Specify the resource group of the nodes to be added to the node pool. |
| Custom Security Group | Select a custom security group. |

6. Click OK.

Result

- After you perform the preceding steps, check the states of the node pools on the **Node Pools** page. If the node pools are in the **Activate** state, the node pools are created.
- You can connect to the Kubernetes cluster and view detailed information about the nodes in the node pools.
 - i. On the **Node Pools** page, select a node pool that you have created and click its name. In the **Node Pool Information** section, find and record the **Node Pool ID**.

| Node Pool Information | | | |
|------------------------------------|---------------------------|------------------|--|
| Node Pool ID: <input type="text"/> | Container Runtime: docker | CPU Policy: none | Created At: Jul 13, 2020, 16:04:47 UTC+8 |

- ii. Connect to the Kubernetes cluster by using kubectl. For more information, see [Connect to a cluster through kubectl](#).
- iii. Run the following command to query the names of the nodes in a specified node pool:

```
kubectl get node --show-labels | grep -E "${node pool ID}/${node pool ID}"
```

```
shell@alicloud:~$ kubectl get node --show-labels | grep -E "cn-hangzhou/containers"
cn-hangzhou/containers-ecs-hfc6.xlarge Ready <none> 6m14s v1.16.6-aliyun.1 ack.aliyun.com/cf76e4e6932c49b09d381f4318fe0447,alibabacloud.com/container-runtime-version=1.1.0,alibabacloud.com/container-runtime=Sandboxed-Container,runv,alibabacloud.com/nodepool-id=cn-hangzhou-ecs-hfc6.xlarge,alibabacloud.com/instance-type=ecs.ebmg5a2.4xlarge,beta.kubernetes.io/os=linux,failure-domain=beta.kubernetes.io/region=cn-hangzhou,failure-domain=beta.kubernetes.io/arch=amd64,beta.kubernetes.io/instance-type=ecs.ebmg5a2.4xlarge,beta.kubernetes.io/os=linux,topology.diskplugin.csi.alibabacloud.com/zone=cn-hangzhou-g,beta.kubernetes.io/zone=cn-hangzhou-g,alibabacloud.com/nodepool-id=cn-hangzhou-ecs-hfc6.xlarge
cn-hangzhou/containers-ecs-hfc6.xlarge Ready <none> 49m v1.16.6-aliyun.1 ack.aliyun.com/cf76e4e6932c49b09d381f4318fe0447,alibabacloud.com/nodepool-id=cn-hangzhou-ecs-hfc6.xlarge,beta.kubernetes.io/arch=amd64,beta.kubernetes.io/instance-type=ecs.hfc6.xlarge,beta.kubernetes.io/os=linux,failure-domain=beta.kubernetes.io/region=cn-hangzhou,failure-domain=beta.kubernetes.io/zone=cn-hangzhou-l,beta.kubernetes.io/zone=cn-hangzhou-l,alibabacloud.com/nodepool-id=cn-hangzhou-ecs-hfc6.xlarge,beta.kubernetes.io/zone=cn-hangzhou-l,alibabacloud.com/zone=cn-hangzhou-l
```

- iv. Run the following command to query detailed information about a specified node:

```
kubectl get node -o wide | grep -E "${node name} {node name}"
```

```
shell@alicloud:~$ kubectl get node -o wide | grep -E "cn-hangzhou/containers-ecs-hfc6.xlarge {cn-hangzhou/containers-ecs-hfc6.xlarge}"
cn-hangzhou/containers-ecs-hfc6.xlarge Ready <none> 7m35s v1.16.6-aliyun.1 <none> Aliyun Linux 2.1903 (Hunting Beagle) 4.19.48-14.el7.x86_64 containers/1.2.5
cn-hangzhou/containers-ecs-hfc6.xlarge Ready <none> 50m v1.16.6-aliyun.1 <none> CentOS Linux 7 (Core) 3.10.0-1062.9.1.el7.x86_64 containers/1.2.5
shell@alicloud:~$
```

3.6.12.6. How do I select between Docker and Sandboxed-Container?

Containers and images have become the industry standards for software packaging and delivery. Kubernetes has become a standard platform for building, developing, and managing containerized cloud-native applications. An increasing number of enterprises and customers choose to deploy their applications in Container Service. Container Service supports two types of runtime: Docker and Sandboxed-Container. This topic describes the differences between these runtimes in the following aspects: implementations and limits, commonly used commands provided by Docker Engine and Containerd, and deployment architectures. This provides references for you to select between Docker and Sandboxed-Container based on your requirements.

Differences between Docker and Sandboxed-Container in terms of implementations and limits

| Item | Docker | Sandboxed-Container V2 | Description |
|--------------|-----------|------------------------|-------------|
| Cluster type | All types | All types | N/A |

| Item | Docker | Sandboxed-Container V2 | Description |
|-----------------------------|--|--|--|
| Node type | <ul style="list-style-type: none"> ECS EBM | EBM | N/A |
| Node operating system | <ul style="list-style-type: none"> CentOS Alibaba Cloud Linux2 | Alibaba Cloud Linux2 | <ul style="list-style-type: none"> You cannot deploy both Docker and Sandboxed-Container on the same node. To deploy both Docker and Sandboxed-Container in a cluster, you can create node pools of different runtime types. |
| Container engine | Docker | Containerd | N/A |
| Monitoring | Supported | Supported | N/A |
| Container log collection | Supported | Sidecar: supported. Manual configuration is required. | N/A |
| Container stdout collection | Supported | Supported | N/A |
| RuntimeClass | Not supported | Supported (runV) | N/A |
| Pod scheduling | No configuration is required. | <ul style="list-style-type: none"> For Kubernetes V1.14.x, you must add the following configuration to the nodeSelector field: <ul style="list-style-type: none"> alibabacloud.com/sandboxed-container: Sandboxed-Container.runv For Kubernetes V1.16.x and later, no configuration is required. | N/A |
| HostNetwork | Supported | Not supported | N/A |
| exec/logs | Supported | Supported | N/A |
| Node data disk | N/A | Required. The data disk must be at least 200 GiB. | N/A |
| Network plug-in | <ul style="list-style-type: none"> Flannel Terway | <ul style="list-style-type: none"> Flannel Terway: supports only the One ENI for Multi-Pod mode. | N/A |
| Kube-proxy mode | <ul style="list-style-type: none"> Iptables IPVS | <ul style="list-style-type: none"> Iptables IPVS | N/A |

| Item | Docker | Sandboxed-Container V2 | Description |
|----------------------------|------------|------------------------|-------------|
| Volume plug-in | CSI Plugin | CSI Plugin | N/A |
| Container root file system | OverlayFS | VirtioFS | N/A |

Differences in the commonly used commands provided by Docker Engine and Containerd

Docker uses Docker Engine for container lifecycle management. Sandboxed-Container uses Containerd for container lifecycle management. These tools support different commands that can be used to manage images and containers. The following table lists the commonly used commands.

| Command | Docker | Containerd | |
|---|----------------|----------------------|------------------------|
| | docker | crictl (recommended) | ctr |
| Queries containers | docker ps | crictl ps | ctr -n k8s.io c ls |
| Queries container details | docker inspect | crictl inspect | ctr -n k8s.io c info |
| Queries container logs | docker logs | crictl logs | N/A |
| Runs a command in a container | docker exec | crictl exec | N/A |
| Mounts local standard input, output, and error streams to a running container | docker attach | crictl attach | N/A |
| Queries resource usage statistics | docker stats | crictl stats | N/A |
| Creates a container | docker create | crictl create | ctr -n k8s.io c create |
| Starts one or more containers | docker start | crictl start | ctr -n k8s.io run |
| Stops one or more containers | docker stop | crictl stop | N/A |
| Removes one or more containers | docker rm | crictl rm | ctr -n k8s.io c del |
| Queries images | docker images | crictl images | ctr -n k8s.io i ls |
| Queries image details | docker inspect | crictl inspecti | N/A |
| Pulls an image | docker pull | crictl pull | ctr -n k8s.io i pull |
| Pushes an image | docker push | N/A | ctr -n k8s.io i push |
| Removes one or more images | docker rmi | crictl rmi | ctr -n k8s.io i rm |
| Queries pods | N/A | crictl pods | N/A |
| Queries pod details | N/A | crictl inspectp | N/A |

| Command | Docker | Containerd | |
|-------------------------|--------|----------------------|-----|
| | docker | crictl (recommended) | ctr |
| Starts one or more pods | N/A | crictl runp | N/A |
| Stops one or more pods | N/A | crictl stopp | N/A |

Differences between Docker and Sandboxed-Container in terms of deployment architectures

| Runtime | Deployment architecture |
|------------------------|--|
| Docker | kubelet -> dockerd -> containerd -> containerd-shim -> runC containers |
| Sandboxed-Container V1 | kubelet -> (CRI)containerd -> containerd-shim -> runC containers -> containerd-shim-kata-v2 -> runV sandboxed containers |
| Sandboxed-Container V2 | kubelet -> (CRI)containerd -> containerd-shim -> runC containers -> containerd-shim-rund-v2 -> runV sandboxed containers |

3.6.12.7. Benefits of Sandboxed-Container

This topic describes the advantages and application scenarios of Sandboxed-Container and provides a comparison between Sandboxed-Container and open source Kata Containers. This allows you to learn more about the benefits of Sandboxed-Container.

Context

Sandboxed-Container provides an alternative to the Docker runtime environment. It supports the following features:

- Sandboxed-Container allows your applications to run in a sandboxed and lightweight virtual machine. This virtual machine is equipped with a dedicated kernel and provides better isolation and enhanced security.
- Compared with open source Kata Containers, Sandboxed-Container is optimized for storage, networking, and stability.
- You can use Sandboxed-Container to isolate untrusted applications and applications of different tenants for higher security. You can also use Sandboxed-Container to isolate applications with faults and applications with degraded performance. This minimizes the negative impact on your service. In addition, Sandboxed-Container offers the same user experience as Docker in terms of logging, monitoring, and elastic scaling.

Benefits

Compared with Docker, Sandboxed-Container has the following benefits:

- Strong isolation based on sandboxed and lightweight virtual machines.
- Compatibility with runC in terms of application management.
- High performance that corresponds to 90% performance of applications based on runC.

- The same user experience as runC in terms of logging, monitoring, and storage.
- Support for RuntimeClass.
- Easy to use with limited expertise that is required to use virtual machines.
- Higher stability than that provided by Kata Containers.

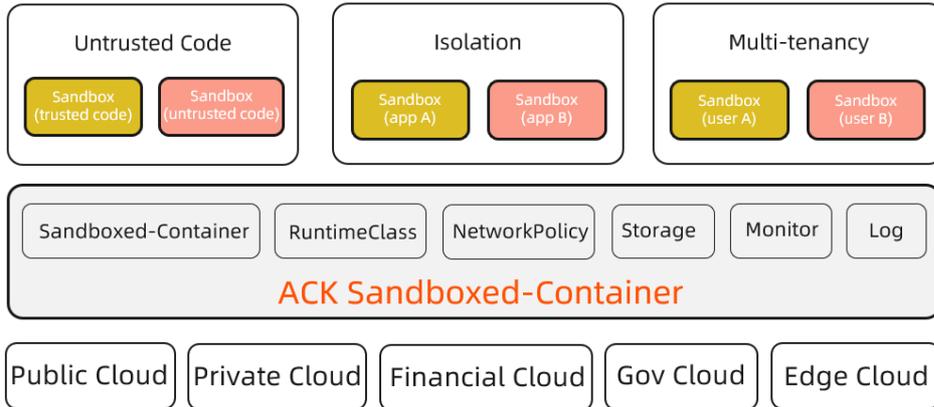
Comparison between Sandboxed-Container and Kata Containers

Sandboxed-Container outperforms Kata Containers in the following aspects.

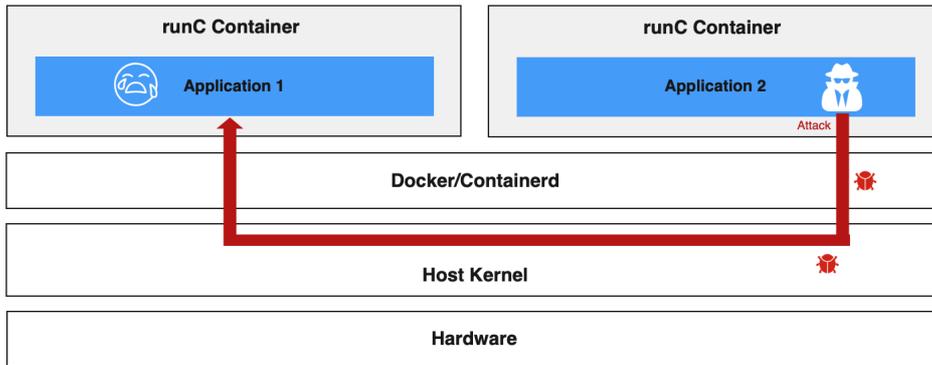
| Item | Category | Sandboxed-Container | Kata Containers |
|----------------------------------|----------|---|---|
| Sandbox startup time consumption | | About 150 ms | About 500 ms |
| Root file system | | OverlayFS over virtio-fs. Performance: ☆☆☆☆ | OverlayFS over 9pfs. Performance: ☆☆ |
| Volume | HostPath | Disks are mounted to Sandboxed-Container over 9pfs. Performance: ☆☆ | Disks are mounted to Kata Containers over 9pfs. Performance: ☆☆ |
| | EmptyDir | over VirtioFS | By default, the volume is mounted to Kata Containers over 9pfs. |
| | Disk | By default, cloud disks are mounted to Sandboxed-Container over virtio-fs. Performance: ☆☆☆☆ | Cloud disks are mounted to Kata Containers over 9pfs. Performance: ☆☆ |
| | NAS | By default, Apsara File Storage NAS (NAS) file systems are mounted to Sandboxed-Container over virtio-fs. Performance: ☆☆☆☆ | NAS file systems are mounted to Kata Containers over 9pfs. Performance: ☆ |
| Network plug-in | | <ul style="list-style-type: none"> • The Terway network plug-in is used. Its network performance is 20% to 30% higher than Flannel. Terway supports features such as NetworkPolicy. This allows you to define the networking policies for pods. • Flannel | Flannel |
| Monitoring and alerting | | <ul style="list-style-type: none"> • Enhanced monitoring of disks and network conditions for pods that host Sandboxed-Container. • Integrated with Cloud Monitor. This facilitates cluster monitoring and alerting. | Monitoring of disks and network conditions is unavailable for pods that host Sandboxed-Container. |
| Stability | | ☆☆☆☆☆ | ☆☆ |

Applicable scenarios of Sandboxed-Container

This section describes the applicable scenarios of Sandboxed-Container.



- Scenario 1: Sandboxed-Container can run untrusted code and applications in isolated containers. This is not supported by containers in runC.
 - Security risks of runC



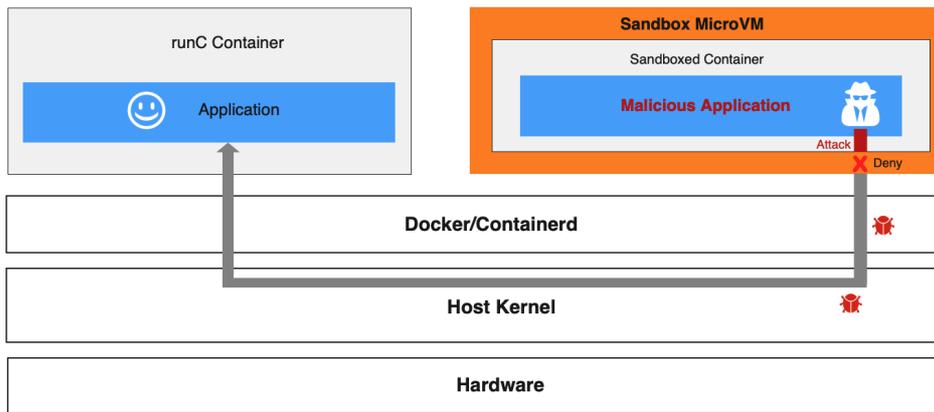
- runC isolates containers by using namespaces and control groups (cgroups). This exposes containers to security threats.
- All containers on a node share the host kernel. If a kernel vulnerability is exposed, malicious code may escape to the host and then infiltrate the backend network. Malicious code execution may cause privilege escalation, compromise sensitive data, and destroy system services and other applications.
- Attackers may also exploit application vulnerabilities to infiltrate the internal network.

You can implement the following measures to reduce security risks of containers in runC.

- Seccomp: filters system calls.
- SELinux: restricts the permissions of container processes, files, and users.
- Capability: limits the capability of container processes.
- dockerd rootless mode: forbids users to use root permissions to run the Docker daemon and containers.

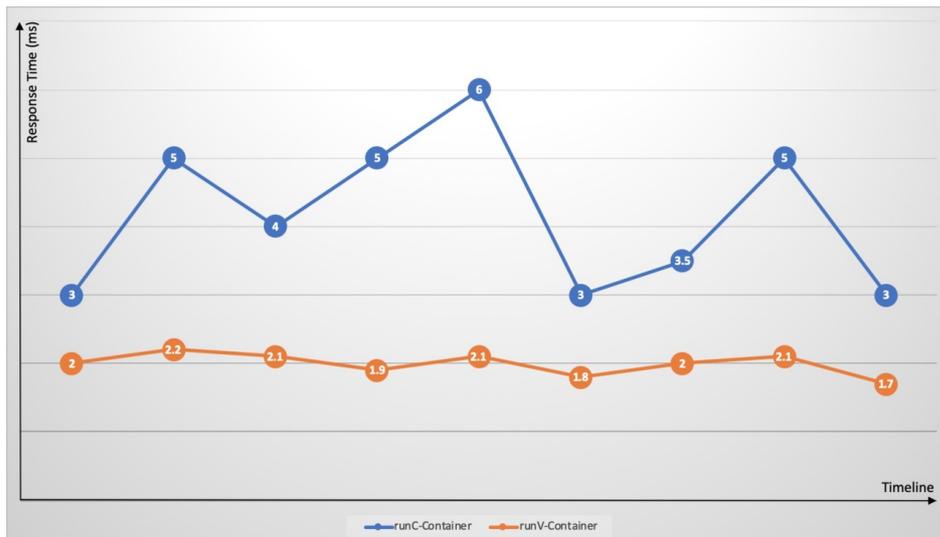
The preceding measures can enhance the security of containers in runC and reduce attacks on the host kernel by malicious containers. However, container escapes and host kernel vulnerabilities remain unresolved.

- o Sandboxed-Container prevents potential risks based on container isolation



In a Sandboxed-Container runtime environment, applications that have potential security risks are deployed on sandboxed and lightweight virtual machines. Each of the virtual machines has a dedicated guest OS kernel. If a security vulnerability is detected on a guest OS kernel, the attack is limited to one sandbox and does not affect the host kernel or other containers. The Terway network plug-in allows you to define networking policies for pods. This enables system isolation, data isolation, and network isolation for Sandboxed-Containers.

- Scenario 2: Sandboxed-Container resolves common issues of runC containers, such as fault spreading, resource contention, and performance interference.

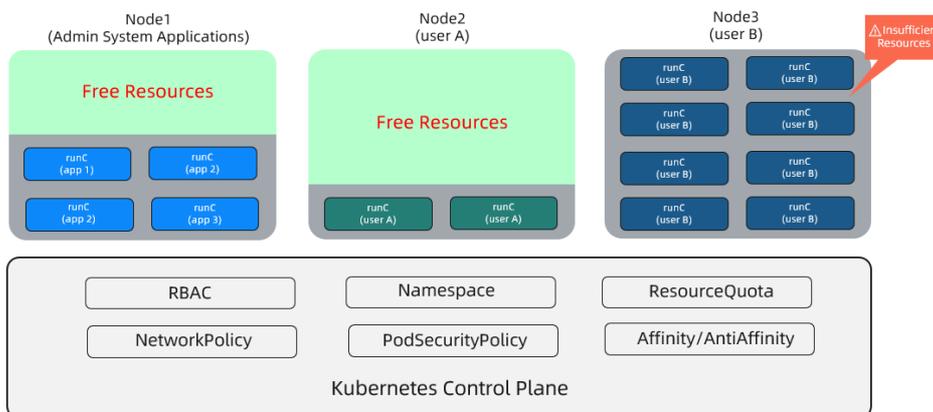


Kubernetes provides easy deployment of different containers on a single node. However, cgroups are not optimized to address resource contention. Resource-intensive applications (such as CPU-intensive, I/O-intensive applications) may compete for the same resources. This causes significant fluctuations in response time and increases the overall response time. Exceptions or faults on an application may spread to the hosting node and disrupt the running of the total cluster. For example, memory leaks and frequent core dumps of an application may overload the node, and exceptions on a container may trigger a host kernel bug that results in complete system failure. Sandboxed-Container addresses the issues that are common with runC containers by using dedicated guest OS kernels and hypervisors. The issues include failure spreading, resource contention, and performance interference.

- Scenario 3: Sandboxed-Container supports multi-tenant services.

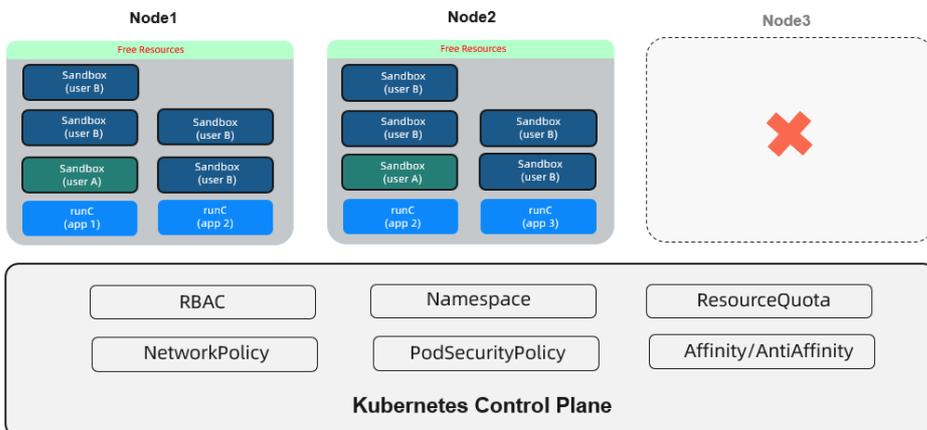
You may need to isolate the applications of an enterprise that consists of multiple business lines or departments. For example, a financial department requires high security applications. However, other non-security-sensitive applications do not have high security requirements. Containers in runC fail to eliminate the potential risks that arise in untrusted applications. In this scenario, you can implement the following counter measures:

- Deploy multiple independent single-tenant clusters. For example, deploy financial business and other non-security-sensitive business in different clusters.
- Deploy a multi-tenant cluster and separate applications of different business lines by namespaces. The resource of a node is exclusive to a single business line. This solution provides data isolation for coordination with the resource quotas and network policies to implement multi-tenant isolation. Compared with multiple single-tenant clusters, this solution focuses on fewer management planes and thus reduces management costs. However, this solution cannot avoid resource waste on nodes. This issue is caused by low resource utilization of some tenants.



Sandboxed-Container allows you to isolate untrusted applications by using sandboxed virtual machines. This prevents the risks of container escapes. This also allows you to deploy different containerized applications on each node. This way, the following benefits are provided:

- Resource scheduling is simplified.
- A node is no longer exclusive to a service. This improves node resource usage and reduces resource fragments and cluster resource costs.
- Sandboxed containers use lightweight virtual machines to provide almost the same performance as containers in runC.



3.6.12.8. Differences between runC and runV

This topic describes the differences between runC and Sandboxed-Container (runV) in terms of their performance and pod creation methods. This allows you to better understand and utilize the benefits of sandboxed containers.

Differences between runC and runV

| Item | runC | runV |
|------------------------------|--|---|
| Container engine | Docker and Containerd | Containerd |
| Node type | Elastic Compute Service (ECS) instances and ECS Bare Metal instances | EBM |
| Container kernel | Share the host kernel | Dedicated kernel |
| Container isolation | Cgroups and namespaces | Lightweight virtual machines (VMs) |
| Rootfs Graph Driver | OverlayFS | DeviceMapper |
| RootFS I/O throttling | Cgroups | DeviceMapper Block IO Limit Note Supported by only Sandboxed-Container V1. |
| NAS mounting | Not supported | Supported |
| Disk mounting | Not supported | Supported |
| Collection of container logs | Logtail directly collects container logs from the host. | logtail sidecar |
| Pod Overhead | None | <ul style="list-style-type: none"> Sandboxed-Container V1: For example, if you set memory: 512 Mi for a pod overhead, it indicates that 512 MiB of memory is allocated to the pod sandbox. Pod overhead refers to the amount of resources consumed by the pod sandbox. In this case, if you set a memory limit of 512 MiB for containers in the pod, the pod will request a total memory of 1,024 MiB. Sandboxed-Container V2: The memory limit for a pod overhead is calculated based on the following formula: Memory for a pod overhead = 64 MiB + Requested memory of containers in a pod × 2%. If the result is greater than 512 MiB, the value is set to 512 Mi. If the result is smaller than 64 MiB, the value is set to 64 Mi. |

Differences in pod creation between runC and runV

You can connect to clusters of Container Service by using kubectl. For more information, see [Connect to a cluster through kubectl](#).

- Create a pod that uses runC
 - (Optional) Use `runtimeClassName: runc` to set the container runtime to runC.

Note The preceding command is optional. runC is the default container runtime.

- ii. Run the following commands to create a pod that uses runC:

```
cat <<EOF | kubectl create -f -
apiVersion: v1
kind: Pod
metadata:
  name: busybox-runc
labels:
  app: busybox-runc
spec:
  containers:
  - name: busybox
    image: registry.cn-hangzhou.aliyuncs.com/acs/busybox:v1.29.2
    command:
    - tail
    - -f
    - /dev/null
  resources:
    limits:
      cpu: 1000m
      memory: 512Mi
    requests:
      cpu: 1000m
      memory: 512Mi
EOF
```

- Create a pod that uses runV

- i. Use `runtimeClassName: runv` to set the container runtime to runV.
- ii. (Optional) Run the following command to verify that a RuntimeClass object named `runv` exists in the cluster.

```
kubectl get runtimeclass runv -o yaml
```

 **Note** A RuntimeClass object named `runv` is automatically created in a Kubernetes cluster that uses Sandboxed-Container.

- iii. Run the following command to create a pod that uses runV:

```
cat <<EOF | kubectl create -f -
apiVersion: v1
kind: Pod
metadata:
  name: busybox-runv
labels:
  app: busybox-runv
spec:
  runtimeClassName: runv
  nodeSelector:
    alibabacloud.com/container-runtime: Sandboxed-Container.runv
  containers:
  - name: busybox
    image: registry.cn-hangzhou.aliyuncs.com/acs/busybox:v1.29.2
    command:
    - tail
    - -f
    - /dev/null
  resources:
    limits:
      cpu: 1000m
      memory: 512Mi
    requests:
      cpu: 1000m
      memory: 512Mi
EOF
```

 **Notice** If the Kubernetes version is earlier than 1.16, add the following nodeSelector configuration:

```
nodeSelector:
  alibabacloud.com/container-runtime: Sandboxed-Container.runv
```

- iv. Run the following command to query the pod that you have created: If the output is `runv`, it indicates that the pod is running in a sandbox.

```
kubectl get pod busybox-runv -o jsonpath={.spec.runtimeClassName}
```

- v. Run the following command to log on to the pod and query its CPU and memory specifications:

```
kubectl exec -ti pod busybox-runv /bin/sh
/# cat /proc/meminfo | head -n1
MemTotal: 1130692 kB
/# cat /proc/cpuinfo | grep processor
processor :0
```

The output shows that the number of CPUs is not the same as that of the host. The total memory is the sum of pod memory and pod overhead. Be aware that the total memory is slightly smaller because the system also consumes some memory.

3.6.12.9. Compatibility notes

This topic describes the pod fields that are supported by Sandboxed-Container. This allows you to fully use the Sandboxed-Container runtime.

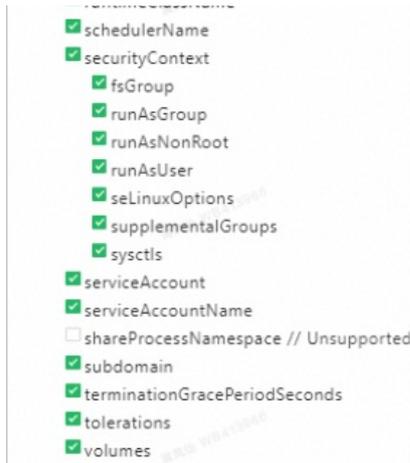
Context

Sandboxed-Container is a new runV container runtime that provides compatibility with runC in terms of pod networking, service networking (ClusterIP and NodePort), and image management. However, Sandboxed-Container does not support all pod fields. To use Sandboxed-Container, you do not need to change your development mode or image packaging method.

Supported pod fields

Sandboxed-Container supports the following pod fields that are marked by ticks:

```
• pod
  ◦ spec
    ✓ activeDeadlineSeconds
    ✓ affinity
    ✓ automountServiceAccountToken
    ✓ containers
      ✓ args
      ✓ command
      ✓ env
      ✓ envFrom
      ✓ image
      ✓ imagePullPolicy
      ✓ lifecycle
      ✓ livenessProbe
      ✓ name
      ✓ ports
      ✓ readinessProbe
      ✓ resources
      ✓ securityContext
        ✓ allowPrivilegeEscalation
        ✓ capabilities
         privileged // Unsupported
        ✓ procMount
        ✓ readOnlyRootFilesystem
        ✓ runAsGroup
        ✓ runAsNonRoot
        ✓ runAsUser
        ✓ seLinuxOptions
         windowsOptions // Unsupported
      ✓ startupProbe
      ✓ stdin
      ✓ stdinOnce
      ✓ terminationMessagePath
      ✓ terminationMessagePolicy
      ✓ tty
      ✓ volumeDevices
      ✓ volumeMounts
      ✓ workingDir
    ✓ dnsConfig
    ✓ dnsPolicy
    ✓ enableServiceLinks
    ✓ hostAliases
     hostIPC // Unsupported
     hostNetwork // Unsupported
     hostPID // Unsupported
    ✓ hostname
    ✓ imagePullSecrets
    ✓ initContainers
    ✓ nodeName
    ✓ nodeSelector
    ✓ priority
    ✓ priorityClassName
    ✓ readinessGates
    ✓ restartPolicy
    ✓ runtimeClassName
```



3.6.13. Use the Kubernetes event center

The event center feature allows you to log Kubernetes events, query events, and configure alerting. This topic describes how to enable and view the Kubernetes event center.

Enable the Kubernetes event center

Enable the Kubernetes event center when a Kubernetes cluster is created

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**. On the page that appears, click **Create Kubernetes Cluster** in the upper-right corner.
3. On the **Create Cluster** page, select **Inst all node-problem-detector and Create Event Center** in the **Log Service** field. For more information about other parameters, see [Create a Kubernetes cluster](#). Then, click **Create Cluster**.
4. On the **Confirm** dialog box, after all check items are verified, select the terms of service and disclaimer and click **OK** to create the cluster.

Enable the Kubernetes event center in App Catalog

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, choose **Market place > App Catalog**.
3. On the **App Catalog** page, find and click **ack-node-problem-detector**.
4. On the **App Catalog - ack-node-problem-detector** page, click the **Parameters** tab and set `eventer.sinks.sls.enabled` to `true`.

```
sinks:
  sls:
    enabled: true
    # If you want the monitoring results to be notified by sls, set enabled to true.
    topic: ""
    project: "k8s-log-cc7640381ac7f4a99971f55a2744a2748"
    # You can view the project information by logging in to the
    # SLS console. Please fill in the name of the project here.
    # eg: your project name is k8s-log-cc18a5f3443dhdss22654da,
    # then you can fill k8s-log-cc18a5f3443dhdss22654da to project label.
    logstore: "k8s-event"
    # You can view the project information by logging in to the
    # SLS console. Please fill the logstore address in here.
```

5. On the **App Catalog - ack-node-problem-detector** page, click the **Description** tab.
6. In the **Deploy** pane, select the cluster and click **Create**.

Access the Kubernetes event center

Access the Kubernetes event center in the Container Service console

1. [Log on to the Container Service console](#).
2. In the left-side navigation pane, click **Clusters**.
3. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
4. In the left-side navigation pane of the details page, choose **Operations > Event Center**.
5. Access the Kubernetes event center on the Event Center page

 **Note** If the Event Center page does not appear, check whether the Kubernetes event center is enabled.

- On the **Event Center** page, click the **Events Overview** tab to go to the overview page of the Kubernetes event center.
- On the **Event Center** page, click the **Cluster Events Query** tab to customize query conditions.
- On the **Event Center** page, click the **Pod Events** tab to query events of pods.

Access the Kubernetes event center by using Log Service

1. View the cluster ID.
 - i. [Log on to the Container Service console](#).
 - ii. In the left-side navigation pane, click **Clusters**.
 - iii. On the **Clusters** page, find the cluster that you want to manage and click **Details** in the **Actions** column.
 - iv. Click the **Basic Information** tab to view the cluster ID.
2. Log on to the Log Service console.
3. In the search box of the Projects section, enter and click `k8s-log-<CLUSTER_ID>`.

 **Note** Replace the `CLUSTER_ID` with the cluster ID that you obtained in Step .

4. In the Logstores section, choose **K8s-event > Visual Dashboards > Kubernetes Event Center V1.2**.
5. On the **Kubernetes Event Center V1.2** page, you can view the trends of **WARNING** events and **ERROR** events.

3.6.14. Use Log Service to collect log data from containers

Container Service is integrated with Log Service. When you create a cluster, you can enable Log Service to collect log data from containers, including standard outputs (stdout) and text files.

Activate Log Service

To activate Log Service, perform the following steps:

1. Log on to the Apsara Uni-manager Management Console. In the top navigation bar, choose **Products > Log Service** to go to the **Log Service** page.
2. Select the required organization and region.
3. Click **SLS** to go to the Log Service console.

Create a Kubernetes cluster and enable Log Service

To create a Kubernetes cluster, perform the following steps:

1. [Log on to the Container Service console](#).

Note The specified organization must be the same as the one that you selected when you activated Log Service. For more information, see [Activate Log Service](#).

- In the left-side navigation pane, click **Clusters**.
- On the Clusters page, click **Create Kubernetes Cluster**. For more information, see [Create a Kubernetes cluster](#).
- In the Component Configuration step, select **Enable Log Service** to install the Log Service component.
- When the Enable Log Service check box is selected, the system prompts you to create a Log Service project. You can only select **Create Project**.

After you select **Create Project**, the system creates a project. By default, the system names the project in the format of `k8s-log-{ClusterID}`. ClusterID indicates the unique ID of the cluster to be created.

- Click **Create Cluster** in the upper-right corner of the page.
- On the **Confirm Order** page, after all check items are verified, select the terms of service and disclaimer and click **OK** to start the deployment.
On the Clusters page, you can find the created cluster.

Install Log Service components in an existing Kubernetes cluster

If you created a Kubernetes cluster and activated Log Service, you can perform the following steps to use Log Service:

- Connect to the Kubernetes cluster by using CloudShell.
For more information, see [Connect to a Kubernetes cluster through kubectl](#).
- Run the `logtail-dedicated.sh` script to install Log Service components in the Kubernetes cluster.

```
#!/env/bin/bash
yaml=$(cat <<-END
---
apiVersion: v1
kind: ConfigMap
metadata:
  name: alibaba-log-config-file
  namespace: kube-system
data:
  ilogtail_config.json: |
    {
      "config_server_address": "http://logtail.$REGION.sls-pub.$INTERNET_DOMAIN",
      "data_server_address": "http://data.$REGION.sls-pub.$INTERNET_DOMAIN",
      "data_server_list":
      [
        {
          "cluster": "$REGION",
          "endpoint": "data.$REGION.sls-pub.$INTERNET_DOMAIN"
        }
      ],
      "shennong_unix_socket": false
    }
---
apiVersion: v1
```

```
apiVersion: extensions/v1beta1
kind: ConfigMap
metadata:
  name: alibaba-log-configuration
  namespace: kube-system
data:
  log-project: "k8s-log-$CLUSTER_ID"
  log-endpoint: "data.$REGION.sls-pub.$INTERNET_DOMAIN"
  log-machine-group: "k8s-group-$CLUSTER_ID"
  log-config-path: "/etc/ilogtail/conf/apsara/ilogtail_config.json"
  log-ali-uid: "$ALI_UID"
  log-access-id: "" # just use blank string
  log-access-key: "" # just use blank string
---
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: alibaba-log-controller
  namespace: kube-system
labels:
  k8s-app: alibaba-log-controller
annotations:
  component.version: "v0.1.3"
  component.revision: "v1"
spec:
  replicas: 1
  template:
    metadata:
      labels:
        k8s-app: alibaba-log-controller
      annotations:
        scheduler.alpha.kubernetes.io/critical-pod: ""
    spec:
      serviceAccountName: alibaba-log-controller
      tolerations:
        - operator: "Exists"
      containers:
        - name: alibaba-log-controller
          image: $IMAGE_REPO_URL/acs/log-controller-$ARCH:v0.1.3.0-527ff4d-aliyun
          resources:
            limits:
              memory: 100Mi
            requests:
              cpu: 50m
              memory: 100Mi
          env:
            - name: "ALICLOUD_LOG_PROJECT"
              valueFrom:
                configMapKeyRef:
                  name: alibaba-log-configuration
                  key: log-project
            - name: "ALICLOUD_LOG_ENDPOINT"
              valueFrom:
                configMapKeyRef:
                  name: alibaba-log-configuration
                  key: log-endpoint
            - name: "ALICLOUD_LOG_MACHINE_GROUP"
              valueFrom:
                configMapKeyRef:
                  name: alibaba-log-configuration
                  key: log-machine-group
```

```

    key: log-access-key
  - name: "ALICLOUD_ACS_K8S_FLAG"
    value: "ture"
  - name: "ALICLOUD_ACCESS_KEY_ID"
    valueFrom:
      configMapKeyRef:
        name: alibaba-log-configuration
        key: log-access-id
  - name: "ALICLOUD_ACCESS_KEY_SECRET"
    valueFrom:
      configMapKeyRef:
        name: alibaba-log-configuration
        key: log-access-key
nodeSelector:
  beta.kubernetes.io/os: linux
---
apiVersion: apiextensions.k8s.io/v1beta1
kind: CustomResourceDefinition
metadata:
  name: aliyunlogconfigs.log.alibabacloud.com
spec:
  group: log.alibabacloud.com
  version: v1alpha1
  names:
    kind: AliyunLogConfig
    plural: aliyunlogconfigs
    scope: Namespaced
---
apiVersion: rbac.authorization.k8s.io/v1beta1
kind: ClusterRoleBinding
metadata:
  name: alibaba-log-controller
subjects:
- kind: ServiceAccount
  name: alibaba-log-controller
  namespace: kube-system
roleRef:
  kind: ClusterRole
  name: alibaba-log-controller
  apiGroup: rbac.authorization.k8s.io
---
apiVersion: rbac.authorization.k8s.io/v1beta1
kind: ClusterRole
metadata:
  name: alibaba-log-controller
  labels:
    k8s-app: alibaba-log-controller
rules:
- apiGroups: ["log.alibabacloud.com"]
  resources:
  - aliyunlogconfigs
  verbs:
  - update
  - get
  - watch
  - list
- apiGroups: [""]
  resources:
  - configmaps
  verbs:
  - create

```

```
- create
- update
- get
- apiGroups: ["" ]
resources:
- events
verbs:
- create
- patch
- update
---
apiVersion: v1
kind: ServiceAccount
metadata:
name: alibaba-log-controller
namespace: kube-system
labels:
k8s-app: alibaba-log-controller
---
apiVersion: extensions/v1beta1
kind: DaemonSet
metadata:
name: logtail-ds
namespace: kube-system
labels:
k8s-app: logtail-ds
annotations:
component.version: "v0.16.16"
component.revision: "v0"
spec:
updateStrategy:
type: RollingUpdate
template:
metadata:
labels:
k8s-app: logtail-ds
annotations:
scheduler.alpha.kubernetes.io/critical-pod: "
spec:
tolerations:
- operator: "Exists"
containers:
- name: logtail
image: $IMAGE_REPO_URL/acs/logtail-$ARCH:v0.16.24.0-c46cd2fe-aliyun
resources:
limits:
memory: 512Mi
requests:
cpu: 100m
memory: 256Mi
livenessProbe:
exec:
command:
- /etc/init.d/ilogtaild
- status
initialDelaySeconds: 30
periodSeconds: 30
securityContext:
privileged: false
env:
name: "ALIBABA_LOGTAIL_CONFIG"
```

```

- name: ALIYUN_LOGTAIL_CONFIG
  valueFrom:
    configMapKeyRef:
      name: alibaba-log-configuration
      key: log-config-path
- name: "ALIYUN_LOGTAIL_USER_ID"
  valueFrom:
    configMapKeyRef:
      name: alibaba-log-configuration
      key: log-ali-uid
- name: "ALIYUN_LOGTAIL_USER_DEFINED_ID"
  valueFrom:
    configMapKeyRef:
      name: alibaba-log-configuration
      key: log-machine-group
- name: "ALICLOUD_LOG_DOCKER_ENV_CONFIG"
  value: "true"
- name: "ALICLOUD_LOG_ECS_FLAG"
  value: "ture"
- name: "ALICLOUD_LOG_DEFAULT_PROJECT"
  valueFrom:
    configMapKeyRef:
      name: alibaba-log-configuration
      key: log-project
- name: "ALICLOUD_LOG_ENDPOINT"
  valueFrom:
    configMapKeyRef:
      name: alibaba-log-configuration
      key: log-endpoint
- name: "ALICLOUD_LOG_DEFAULT_MACHINE_GROUP"
  valueFrom:
    configMapKeyRef:
      name: alibaba-log-configuration
      key: log-machine-group
- name: "ALICLOUD_LOG_ACCESS_KEY_ID"
  valueFrom:
    configMapKeyRef:
      name: alibaba-log-configuration
      key: log-access-id
- name: "ALICLOUD_LOG_ACCESS_KEY_SECRET"
  valueFrom:
    configMapKeyRef:
      name: alibaba-log-configuration
      key: log-access-key
- name: "ALIYUN_LOG_ENV_TAGS"
  value: "_node_name_|_node_ip_"
- name: "_node_name_"
  valueFrom:
    fieldRef:
      fieldPath: spec.nodeName
- name: "_node_ip_"
  valueFrom:
    fieldRef:
      fieldPath: status.hostIP
volumeMounts:
- name: sock
  mountPath: /var/run/docker.sock
- name: root
  mountPath: /logtail_host
  readOnly: true
  name: alibaba-log-config-file-volume

```

```

- name: alibaba-log-config-file-volume
  mountPath: /etc/ilogtail/conf/apsara
  readOnly: true
  terminationGracePeriodSeconds: 30
  nodeSelector:
    beta.kubernetes.io/os: linux
  volumes:
- name: sock
  hostPath:
    path: /var/run/docker.sock
    type: Socket
- name: root
  hostPath:
    path: /
    type: Directory
- name: alibaba-log-config-file-volume
  configMap:
    name: alibaba-log-config-file
END
)
echo "$yaml" > logtail.yml
kubectl create -f logtail.yml

```

3. Replace `<your server architecture>` , `<your k8s cluster region_id>` , `<your_k8s_cluster_id>` , `<k8s_cluster_domain_suffix>` , `<your_ali_uid>` , and `<your_image_repo_url>` with actual values, and run the following commands. This allows you to set the environment variables and deploy the components.

```

export ARCH=<your_server_architecture>
export REGION=<your_k8s_cluster_region_id>
export CLUSTER_ID=<your_k8s_cluster_id>
export INTERNET_DOMAIN=<k8s_cluster_domain_suffix>
export IMAGE_REPO_URL=<your_image_repo_url>
export ALI_UID=<your_ali_uid>
bash logtail-dedicated.sh // Run the script to install the components.

```

Note

- `<your_server_architecture>` : the server architecture, for example, amd64.
- `<your_k8s_cluster_region_id>` : the region where the Kubernetes cluster is deployed, for example, cn-qingdao-apsara-d01.
- `<your_k8s_cluster_id>` : the ID of the Kubernetes cluster.
- `<k8s_cluster_domain_suffix>` : the domain suffix of the Kubernetes cluster, for example, env28.internet.com.
- `<your_ali_uid>` : the ID of the Apsara Stack tenant account, for example, 1234074238634394.
- `<your_image_repo_url>` : the URL of the image repository, for example, registry.cn-hangzhou.aliyuncs.com.

Create an application and configure Log Service

When you create an application in Container Service, you can configure Log Service to collect log data from containers. You can only use YAML templates to configure Log Service.

1. Log on to the Container Service console. In the left-side navigation pane, choose **Applications > Deployments**. In the upper-right corner of the Deployments page, click **Create from Template**.
2. YAML templates follow the Kubernetes syntax. You can use environment variables to add **collection configurations** and **custom tags**. You must also configure volumeMounts and volumes. The following is an example of a pod for collecting log data:

```

apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  labels:
    app: logtail-test
    name: logtail-test
spec:
  replicas: 1
  template:
    metadata:
      labels:
        app: logtail-test
        name: logtail-test
    spec:
      containers:
        - name: logtail
          image: registry.acs.env28.intranet.com/acs/busybox:latest
          args:
            - ping
            - 127.0.0.1
          env:
            - name: aliyun_logs_log-stdout
              value: stdout
            - name: aliyun_logs_log-varlog
              value: /log/*.log
            - name: aliyun_logs_log_tags
              value: tag1=v1
          volumeMounts:
            - name: volumn-sls
              mountPath: /log
      volumes:
        - name: volumn-sls
          emptyDir: {}

```

- Specify the following configurations in order based on your business requirements:
- Use environment variables to add **collection configurations** and **custom tags**. All environment variables for collection configurations must use the `aliyun_logs_` prefix.
- Add log collection configurations in the following format:

```

- name: aliyun_logs_{Logstore name}
  value: {Log file path}

```

In the preceding YAML template, two environment variables are added to the log collection configuration. Environment variable `aliyun_logs_log-stdout` indicates that a Logstore named `log-stdout` is created to store the standard outputs collected from containers.

 **Note** The name of the Logstore cannot contain underscores (`_`). You can use hyphens (`-`) instead.

- **Custom tags** must be specified in the following format:

```

- name: aliyun_logs_{Tag name without underscores (_) }_tags
  value: {Tag name}={Tag value}

```

After a tag is added, the tag is automatically appended to the log files that are collected from the container.

- If you specify a log path to collect log files that are not standard outputs, you must configure `volumeMounts`.

In the preceding YAML template, the `mountPath` field in `volumeMounts` is set to `/var/log`. This allows the Log Service component to collect `/var/log/*.log` files.

3. After you modify the YAML template, click **Create** to submit the configurations.

Configure advanced parameters in the `env` field

You can specify multiple parameters in the `env` field to configure log collection. The following table describes the parameters.

| Parameter | Description | Example | Precaution |
|---|---|---|---|
| <code>aliyun_logs_{key}</code> | <ul style="list-style-type: none"> Required. <code>{key}</code> can contain only lowercase letters, digits, and hyphens (-), and cannot contain underscores (_). If the specified <code>aliyun_logs_{key}_logstore</code> does not exist, a Logstore named <code>{key}</code> is created. To collect standard outputs of container log files, set the value to <code>stdout</code>. You can also set the value to a path inside a container. | <pre>- name: aliyun_logs_catalina stdout - name: aliyun_logs_access- log /var/log/nginx/acces s.log</pre> | <ul style="list-style-type: none"> By default, the Log Service component collects log files in simple mode. In this case, the collected log files are not parsed. If you want to parse the log files, we recommend that you change the collection mode in the Log Service console. The value of <code>{key}</code> must be unique in the cluster. |
| <code>aliyun_logs_{key}_tags</code> | Optional. This parameter is used to add tags to log data. The value must be in the following format: <code>{tag-key}={tag-value}</code> . | <pre>- name: aliyun_logs_catalina_t ags app=catalina</pre> | - |
| <code>aliyun_logs_{key}_project</code> | Optional. This parameter specifies a project in Log Service. By default, the project that you specified when you create the cluster is used. | <pre>- name: aliyun_logs_catalina_p roject my-k8s-project</pre> | The region of the project must be the same as where your <code>logtail-ds</code> is deployed. |
| <code>aliyun_logs_{key}_logstore</code> | Optional. This variable specifies a Logstore in Log Service. By default, the Logstore is named after <code>{key}</code> . | <pre>- name: aliyun_logs_catalina_t ags my-logstore</pre> | - |
| <code>aliyun_logs_{key}_shard</code> | Optional. This variable specifies the number of shards in the Logstore. Valid values: 1 to 10. Default value: 2. | <pre>- name: aliyun_logs_catalina_s hard 4</pre> | - |

| Parameter | Description | Example | Precaution |
|--------------------------------|---|--|------------|
| aliyun_logs_{key}_ttl | Optional. This parameter specifies the number of days for which log data is retained. Valid values: 1 to 3650. <ul style="list-style-type: none"> To permanently retain log data, set the value to 3650. Default value: 90. | <pre>- name: aliyun_logs_catalina_t tl 3650</pre> | - |
| aliyun_logs_{key}_machinegroup | Optional. This parameter specifies the machine group of the application. By default, the machine group is the one where your logtail-ds is deployed. | <pre>- name: aliyun_logs_catalina_ machinegroup my-machine-group</pre> | - |

- Scenario 1: Collect log files from multiple applications and store them in the same Logstore

In this scenario, set the `aliyun_logs_{key}_logstore` parameter. The following example shows how to collect standard outputs from two applications and store the outputs in `stdout-logstore`.

Configure the following environment variables for Application 1:

```
##### Configure environment variables #####
- name: aliyun_logs_app1-stdout
  value: stdout
- name: aliyun_logs_app1-stdout_logstore
  value: stdout-logstore
```

Configure the following environment variables for Application 2:

```
##### Configure environment variables #####
- name: aliyun_logs_app2-stdout
  value: stdout
- name: aliyun_logs_app2-stdout_logstore
  value: stdout-logstore
```

- Scenario 2: Collect log files from different applications and store them in different projects

In this scenario, perform the following steps:

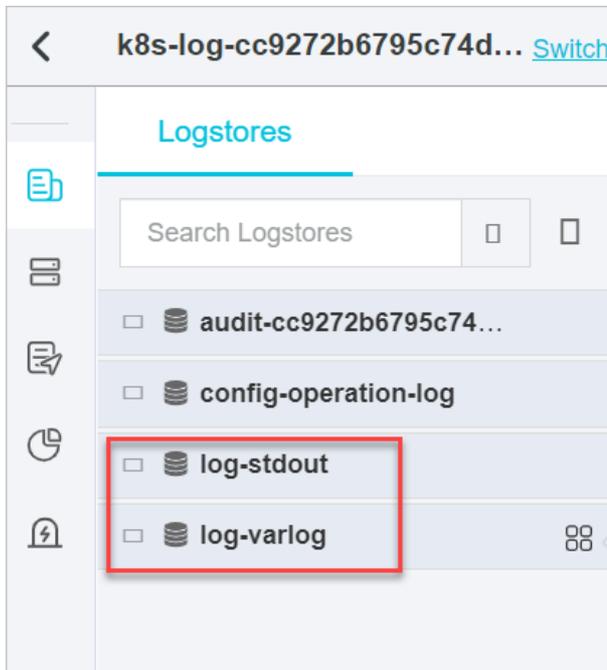
- Create a machine group in each project and set the machine group ID in the following format: `k8s-group-{cluster-id}`, where `{cluster-id}` is the ID of the cluster. You can customize the machine group name.
- Specify the project, Logstore, and machine group in the environment variables for each application.

```
##### Configure environment variables #####
- name: aliyun_logs_app1-stdout
  value: stdout
- name: aliyun_logs_app1-stdout_project
  value: app1-project
- name: aliyun_logs_app1-stdout_logstore
  value: app1-logstore
- name: aliyun_logs_app1-stdout_machinegroup
  value: app1-machine-group
```

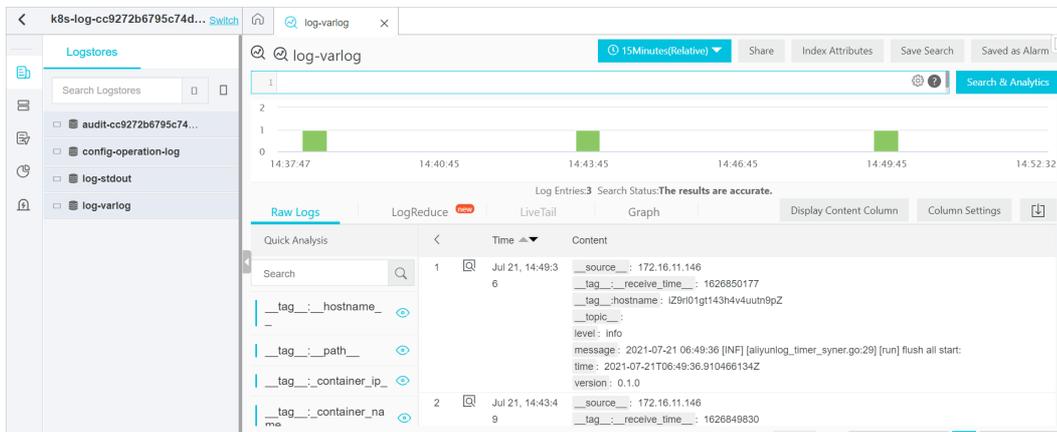
View log data

Perform the following steps to view log data:

1. Log on to the Log Service console. For more information, see [Activate Log Service](#).
2. Select the project that is associated with the Kubernetes cluster to go to the Logstores tab. By default, the project name is in the format of k8s-log-{Kubernetes cluster ID}.
3. In the list of Logstores, find the Logstores and click Search in the Log Search column for each Logstore. In this example, the Logstores are log-stdout and log-varlog.



4. On the Raw Logs tab, you can view raw log data in log-stdout and log-varlog.



4.Auto Scaling (ESS)

4.1. What is Auto Scaling?

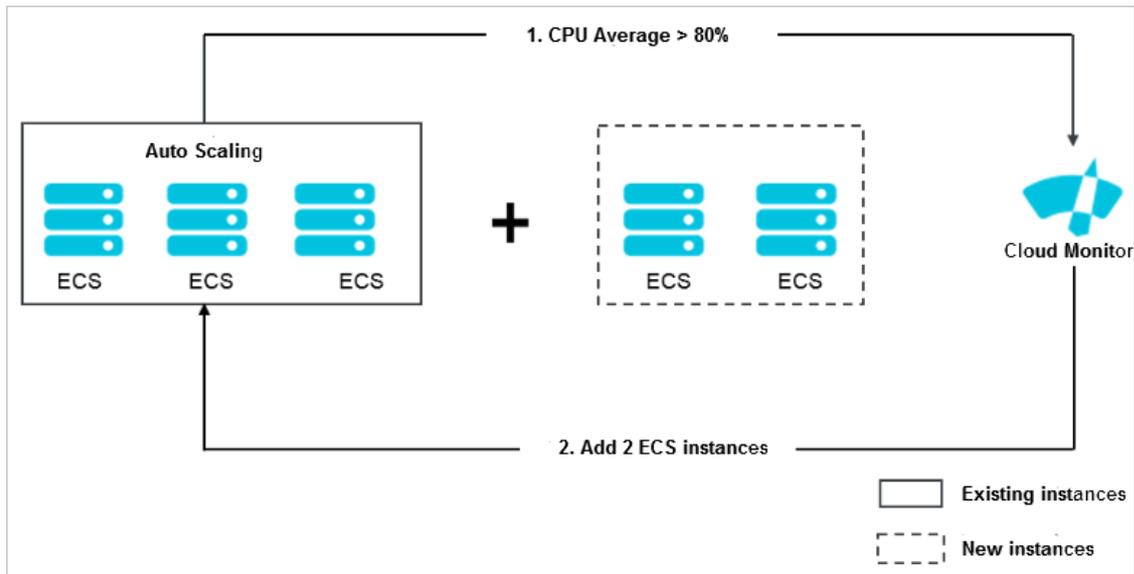
Auto Scaling is a management service that automatically adjusts your elastic computing resources based on your business needs and policies.

When business loads increase, Auto Scaling automatically adds ECS instances based on user-defined scaling rules to ensure sufficient computing capabilities. When business loads decrease, Auto Scaling automatically removes ECS instances to save costs.

Auto Scaling provides the following features:

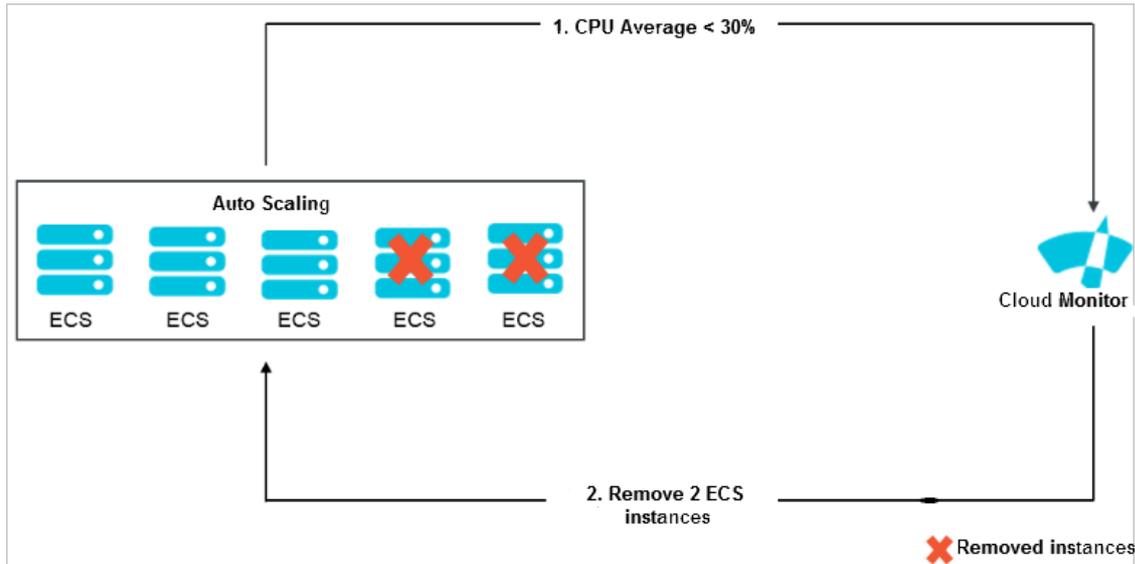
- Scale-out

When business loads surge above normal loads, Auto Scaling automatically increases underlying resources. This helps maintain access speed and ensures that resources are not overloaded. For example, if the CPU utilization of ECS instances exceeds 80%, Auto Scaling scales out ECS resources based on your configured rules. During the scale-out event, Auto Scaling automatically creates and adds ECS instances to a scaling group, and adds the new instances to the backend server groups of the associated SLB instances and the whitelists of the associated ApsaraDB RDS instances. The following figure shows the implementation of a scale-out event.



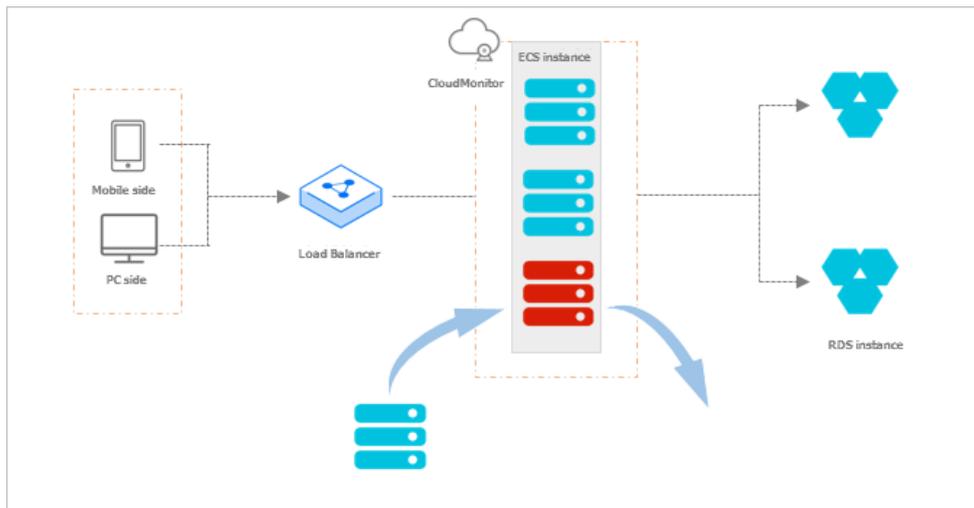
- Scale-in

When your business loads decrease, Auto Scaling automatically releases underlying resources to prevent resource wastage and reduce costs. For example, if the CPU utilization of ECS instances in a scaling group is less than 30%, Auto Scaling automatically scales in ECS instances based on your configured rules. During the scale-in event, Auto Scaling removes ECS instances from the scaling group and also from the backend server groups of the associated SLB instances as well as the whitelists of the associated ApsaraDB RDS instances. The following figure shows the implementation of a scale-in event.



- Elastic recovery

If ECS instances in a scaling group are not in the Running state, Auto Scaling considers the instances to be unhealthy. If an ECS instance is considered unhealthy, Auto Scaling automatically releases the instance and creates a new one. This process is called elastic recovery. It ensures that the number of healthy ECS instances in a scaling group does not fall below the minimum number of ECS instances that you specified for the scaling group. The following figure shows the implementation of elastic recovery.



4.2. Notes

4.2.1. Precautions

This topic describes the precautions when you use Auto Scaling (ESS).

Scaling rules

ESS uses scaling rules to scale ECS instances in a scaling group based on the minimum and maximum numbers of ECS instances specified for the scaling group. Assume that a scaling group can contain up to 45 ECS instances. If you configure a scaling rule to increase the number of ECS instances in the scaling group to 50, ESS only increases the number of ECS instances to 45 at most.

Scaling activities

- Only one scaling activity can be executed at a time in a scaling group.
- An ongoing scaling activity cannot be terminated. For example, if a scaling activity is being executed to create 20 ECS instances but only five have been created, you cannot forcibly terminate the scaling activity.
- If some ECS instances fail to be added to a scaling group during a scaling activity, ESS considers that the scaling activity is complete without trying to add the failed instances to the scaling group. ESS rolls back the ECS instances that fails to be added but not the scaling activity. For example, if ESS has created 20 ECS instances for a scaling group, and 19 of the instances are added to SLB instances, only the one ECS instance that failed to be added is automatically released.

Cooldown period

- During the cooldown period, if you manually execute a scaling task, such as a scaling rule or scheduled task, the task is immediately executed without waiting for the cooldown period to expire.
- The cooldown period starts after the last ECS instance is added to or removed from a scaling group during a scaling activity.

4.2.2. Manual intervention

If you manually intervene with Auto Scaling operations, Auto Scaling processes the intervention accordingly.

Auto Scaling does not prevent you from performing manual intervention, such as deleting automatically created ECS instances in the ECS console. The following table describes how Auto Scaling processes manual intervention.

| Resource | Manual intervention type | Processing method |
|----------|--|---|
| ECS | A user deletes an ECS instance from a scaling group by using the ECS console or calling API operations. | Auto Scaling performs health checks to determine whether the ECS instance is unhealthy. If the instance is unhealthy, Auto Scaling removes it from the scaling group. The internal IP address of the ECS instance is not automatically deleted from the whitelist of the associated ApsaraDB RDS (RDS) instance. If the total number of ECS instances in the scaling group falls below the lower limit after the ECS instance is removed, the scaling group automatically creates an ECS instance to ensure that the number of instances reaches the lower limit. |
| ECS | A user revokes the ECS API permissions granted to Auto Scaling. | Auto Scaling rejects all scaling activity requests. |
| SLB | A user manually removes an ECS instance from an SLB instance by using the SLB console or calling API operations. | Auto Scaling does not automatically detect this action or handle such exceptions. The ECS instance remains in the scaling group. When a scale-in event is triggered, Auto Scaling releases the ECS instance if the instance meets the removal policy. |
| SLB | A user manually deletes an SLB instance or disables the health check feature for an SLB instance by using the SLB console or calling API operations. | Auto Scaling does not add ECS instances to scaling groups that are associated with this SLB instance. Auto Scaling removes ECS instances from the scaling groups if a scaling task triggers a scale-in rule or the ECS instances are recognized as unhealthy after a health check is performed. |
| SLB | An SLB instance is unavailable because of system-related reasons. | All scaling activities fail except for instance removal tasks that are manually executed. |
| SLB | A user revokes the SLB API permissions granted to Auto Scaling. | Auto Scaling rejects all scaling activity requests for scaling groups with which SLB instances are associated. |

| Resource | Manual intervention type | Processing method |
|----------|---|---|
| RDS | A user manually removes the IP address of an ECS instance from the whitelist of the associated RDS instance by using the RDS console or calling API operations. | Auto Scaling does not automatically detect this action or handle such exceptions. The ECS instance remains in the scaling group. When a scale-in event is triggered, Auto Scaling releases the ECS instance if the instance meets the removal policy. |
| RDS | A user manually deletes an RDS instance by using the RDS console or calling API operations. | Auto Scaling does not add ECS instances that are associated with this RDS instance to scaling groups. Auto Scaling removes ECS instances from the scaling groups if a scaling task triggers a scale-in rule or the ECS instances are recognized as unhealthy after a health check is performed. |
| RDS | An RDS instance is unavailable because of system-related reasons. | All scaling activities fail except for instance removal tasks that are manually executed. |
| RDS | A user revokes the RDS API permissions granted to Auto Scaling. | Auto Scaling rejects all scaling activity requests for the scaling groups associated with RDS instances. |

4.2.3. Limits

This topic describes the limits of ESS.

- ESS does not support vertical scaling. It can only scale the number of ECS instances. The CPU, memory, and bandwidth configurations of ECS instances cannot be automatically adjusted.
- The following table describes the quantity limits that are applied to a scaling group.

| Item | Quota |
|-----------------------|--|
| Scaling configuration | You can create a maximum of 10 scaling configurations for a scaling group. |
| Scaling rule | You can create a maximum of 50 scaling rules for a scaling group. |
| ECS instance | A scaling group can contain a maximum of 1,000 ECS instances. |

4.2.4. Scaling group status

This topic describes the states of a scaling group in the console and in an API operation.

| State in the console | State in an API operation |
|----------------------|---------------------------|
| Creating | Inactive |
| Created | Inactive |
| Enabling | Inactive |
| Enabled | Active |
| Disabling | Inactive |
| Disabled | Inactive |
| Deleting | Deleting |

4.2.5. Scaling processes

Before you use Auto Scaling, you must understand the processes related to scaling activities.

Automatic scaling of a scaling group

- Automatic scale-out
 - i. Check the health status and boundary conditions of the scaling group.
 - ii. Assign the activity ID and execute the scaling activity.
 - iii. Create ECS instances.
 - iv. Modify Total Capacity.
 - v. Assign IP addresses to the created ECS instances.
 - vi. Add the ECS instances to the whitelist of the associated ApsaraDB RDS instance.
 - vii. Start the ECS instances.
 - viii. Associate the ECS instances with an SLB instance and set the weight to the SLB weight value that is specified when the scaling configuration is created.
 - ix. The cooldown period starts after the scaling activity is complete.
- Automatic scale-in
 - i. Check the health status and boundary conditions of the scaling group.
 - ii. Assign the activity ID and execute the scaling activity.
 - iii. Remove ECS instances from the associated SLB instance.
 - iv. Stop the ECS instances.
 - v. Remove the ECS instances from the whitelist of the associated ApsaraDB RDS instance.
 - vi. Release the ECS instances.
 - vii. Modify Total Capacity.
 - viii. The cooldown period starts after the scaling activity is complete.

Manually add or remove existing ECS instances

- Manually add instances
 - i. Check the health status and boundary conditions of the scaling group, and check the status and type of ECS instances.
 - ii. Assign the activity ID and execute the scaling activity.
 - iii. Add the ECS instances.
 - iv. Modify Total Capacity.
 - v. Add the ECS instances to the whitelist of the associated ApsaraDB RDS instance.
 - vi. Associate the ECS instances with an SLB instance and set the weight to the SLB weight value that is specified in the active scaling configuration.

 **Note** If you want to manually add an instance to a scaling group, the instance type of the instance must be the same as that specified in the active scaling configuration of the scaling group. Therefore, you must set the weight to the SLB weight value that is specified in the active scaling configuration.

- vii. The cooldown period starts after the scaling activity is complete.
- Manually remove instances
 - i. Check the health status and boundary conditions of the scaling group.

- ii. Assign the activity ID and execute the scaling activity.
- iii. SLB stops forwarding traffic to ECS instances.
- iv. Remove the ECS instances from SLB after 60 seconds.
- v. Remove the ECS instances from the whitelist of the associated ApsaraDB RDS instance.
- vi. Modify Total Capacity.
- vii. Remove the ECS instances from the scaling group.
- viii. After the scaling activity is complete, the cooldown period starts.

4.2.6. Remove unhealthy ECS instances

Before you use ESS, you must understand information about the removal of unhealthy ECS instances.

After an ECS instance is added to a scaling group, ESS checks the status of the instance on a regular basis. If the ECS instance is not in the Running state, ESS removes the ECS instance from the scaling group. The removal method depends on how the ECS instance is added:

- If an ECS instance is automatically created, ESS immediately removes and releases it.
- If an ECS instance is manually added, ESS immediately removes it, but does not stop or release it.

The removal of unhealthy ECS instances is not limited by the MinSize value. After the unhealthy ECS instances are removed, the number of ECS instances (Total Capacity) may fall below the MinSize value. In this case, ESS automatically creates ECS instances based on the difference between the actual instance number and MinSize value to ensure that the total number of ECS instances is equal to the MinSize value.

4.2.7. Instance rollback after a failed scaling activity

Before you use ESS, you must understand the mechanism of instance rollback after a failed scaling activity.

If some ECS instances fail to be added to a scaling group during a scaling activity, ESS considers that the scaling activity is complete without trying to add the failed instances to the scaling group. ESS rolls back ECS instances, not the scaling activity.

For example, if a scaling group has created 20 ECS instances, and 19 of the instances are added to SLB instances, only the one ECS instance that failed to be added is automatically released.

4.2.8. Instance lifecycle management

Before you use Auto Scaling (ESS), you must understand concepts related to the instance lifecycle.

Automatically created ECS instances

ECS instances are automatically created by ESS based on user-defined scaling configurations and rules.

ESS manages the entire lifecycle of automatically created ECS instances. ESS creates ECS instances during scale-out events, and stops and releases them during scale-in events.

Manually added ECS instances

ECS instances are manually added to a scaling group.

ESS does not manage the entire lifecycle of manually added ECS instances. These instances are not automatically created by ESS, but are manually added by a user to a scaling group. If the ECS instances are manually or automatically removed from the scaling group, ESS removes the instances but does not stop or release them.

Instance status

An ECS instance in a scaling group goes through the following states during its lifecycle:

- Pending: The ECS instance is being added to the scaling group. The instance is being created, added to an SLB instance, or added to the whitelist of the associated ApsaraDB RDS instance.

- **InService:** The ECS instance is added to the scaling group and is providing services normally.
- **Removing:** The ECS instance is being removed from the scaling group.

Instance health status

An ECS instance in a scaling group has the following health states:

- **Healthy**
- **Unhealthy**

If an ECS instance is not in the Running state, ESS considers the instance to be unhealthy and automatically removes it from the scaling group.

- ESS stops and releases automatically created ECS instances.
- ESS does not stop or release manually added ECS instances.

4.3. Quick start

4.3.1. Overview

This topic describes how to create a scaling group and how to add or remove ECS instances.

You can perform the following steps to create a scaling group, and add or remove ECS instances.

1. **Create a scaling group**

Set the parameters for the scaling group, such as the Maximum Capacity and Minimum Capacity of ECS instances.

2. **Create a scaling configuration**

Set the parameters for the scaling configuration, such as Instance Type and Image.

3. **Enable a scaling group**

Enable the scaling group after creating the scaling configuration.

4. **Create a scaling rule**

Specify how to add or remove ECS instances. For example, add an ECS instance to a scaling group.

5. **Create a scheduled task**

Create scheduled tasks to add or remove instances at a specified time point. Auto Scaling executes the scheduled tasks and scaling rules at the specified time. For example, Auto Scaling can trigger a task to execute a specific scaling rule at 08:00 everyday.

4.3.2. Log on to the Auto Scaling console

This topic describes how to log on to the Auto Scaling console.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

Note When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top navigation bar, choose **Products > Elastic Computing > Auto Scaling**.

4.3.3. Create a scaling group

This topic describes how to create a scaling group. A scaling group is a group of Elastic Compute Service (ECS) instances that is dynamically scaled based on the configured scenario. You can specify the minimum and maximum numbers of ECS instances in a scaling group.

Prerequisites

- A virtual private cloud (VPC) and a vSwitch are created. For more information, see [Create a VPC](#) and [Create a vSwitch](#) in *VPC User Guide*.
- To associate a scaling group with Server Load Balancer (SLB) instances, make sure that the following requirements are met:
 - You have one or more SLB instances in the **Running** state.
 - The SLB instances and the scaling group are in the same organization, resource set, and region.
- To associate a scaling group with ApsaraDB RDS instances, make sure that the following requirements are met:
 - You have one or more RDS instances in the **Running** state.
 - The RDS instances and the scaling group are in the same organization, resource set, and region.

Procedure

1. **Log on to the Auto Scaling console**.
2. In the left-side navigation pane, click **Scaling groups**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. In the upper-right corner of the Scaling Groups page, click **Create Scaling Group**.
5. Configure parameters for the scaling group.

| Parameter | Required | Description |
|-----------------------------|----------|---|
| Scaling Group | Yes | The name of the scaling group. It must be 2 to 64 characters in length, and can contain letters, digits, underscores (_), hyphens (-), and periods (.). It must start with a letter or a digit. |
| Organization/Resource Group | Yes | The organization and resource set in which to create the scaling group. |
| Maximum Capacity | Yes | The maximum number of instances that the scaling group can contain. Set the value based on your business requirements to control costs. Valid values: 0 to 1000. |

| Parameter | Required | Description |
|--------------------|----------|---|
| Minimum Capacity | Yes | <p>The minimum number of instances that a scaling group must contain. Set the value based on your business requirements to ensure service availability. When the scaling group is enabled, Auto Scaling automatically creates this number of ECS instances.</p> <p>Valid values: 0 to 1000.</p> |
| Cooldown (Seconds) | Yes | <p>The period of time after each scaling activity is complete. During the cooldown time, Auto Scaling rejects all scaling activity requests triggered by event-triggered tasks from CloudMonitor. However, scaling activities triggered by other types of tasks such as manually triggered tasks and scheduled tasks are not limited by the cooldown time and are immediately executed.</p> <p>The value must be an integer that is greater than or equal to zero. Unit: seconds.</p> |
| Scale-In Policy | No | <p>The policy for automatically removing ECS instances from the scaling group. This parameter contains the Select and From the list, select fields. You cannot specify the same values for the two fields. Valid values:</p> <ul style="list-style-type: none"> ◦ Earliest Instance Created Using Scaling Configuration: filters instances that were created based on the earliest scaling configuration. ◦ Earliest Created Instance: filters instances that were added to the scaling group at the earliest point in time. ◦ Newest Instances: filters instances that were most recently added to the scaling group. ◦ None: is available only for From the list, select. This value indicates that Auto Scaling does not filter instances based on the From the list, select field. <p>For example, if Auto Scaling filters instances based on the Earliest Created Instance value of Select, you can select one of the following values for From the list, select:</p> <ul style="list-style-type: none"> ◦ None: indicates that Auto Scaling does not filter instances based on the From the list, select field. ◦ Newest Instances: filters instances obtained based on the Select field and then filters instances that were most recently added to the scaling group. |
| Region/VPC | Yes | The region and VPC in which to create the scaling group. |
| Vswitch | Yes | The ID of the vSwitch with which to associate the scaling group. |

| Parameter | Required | Description |
|------------------------|----------|---|
| Associate SLB Instance | No | <p>After you associate SLB instances with the scaling group, ECS instances that are added to the scaling group are automatically added as SLB backend servers. You can specify a server group to which to add the ECS instances. ECS instances can be added to the following types of server groups:</p> <ul style="list-style-type: none"> Default server group: the group of ECS instances that are used to receive requests. If the listener is not configured with a vServer group or a primary/secondary server group, requests are forwarded to the ECS instances in the default server group. vServer group: If you want to distribute different requests to different backend servers or configure domain name- or URL-based routing methods, you can use vServer groups. |
| Associate RDS Instance | No | <p>After you associate RDS instances with the scaling group, the internal IP addresses of ECS instances that are added to the scaling group are automatically added to the whitelists of the RDS instances to allow communication over the internal network.</p> |

6. Click OK.

Result

The created scaling group is displayed in the scaling group list but is in the **Disabled** state. To enable the scaling group, you must create a scaling configuration. For more information, see [Create a scaling configuration](#).

4.3.4. Create a scaling configuration

This topic describes how to create a scaling configuration for a scaling group.

Prerequisites

At least one security group is available. If you do not have any security groups, create a security group. For more information, see *Create a security group in ECS User Guide*.

Context

You can create only a limited number of scaling configurations for a scaling group. For more information, see the Limits topic in *Auto Scaling Product Introduction*.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **Scaling Configuration**.
5. Choose **Create > Create Scaling Configuration**.
6. Configure parameters for the scaling configuration.

| Section | Parameter | Required | Description |
|---------|-----------|----------|-------------|
|---------|-----------|----------|-------------|

| Section | Parameter | Required | Description |
|----------------|-----------------|----------|---|
| Region | Region | Yes | The region where the ECS instance is located. |
| | Zone | Yes | The zone where the ECS instance is located. |
| Security Group | Security Group | Yes | The security group to which the ECS instance belongs. |
| Instance | Instance Family | Yes | The instance family to which the ECS instance belongs. |
| | Instance Type | Yes | The instance type of the ECS instance. |
| Image | Image Type | Yes | <ul style="list-style-type: none"> ◦ Public Image: Public images provided by Alibaba Cloud are fully licensed to offer a secure and stable operating environment for applications on ECS instances. ◦ Custom Image: You can create custom images to install software or deploy projects that have special requirements. |
| Storage | System Disk | Yes | Specify the category and size of the system disk. The operating system is installed on the system disk. You can select Ultra Disk or Standard SSD . |
| | Data Disk | No | Specify the category and size of the data disk. You can select Ultra Disk or Standard SSD . You can add a maximum of 16 data disks. The maximum capacity of each data disk is 32 TiB. You can set Release with Instance and Encrypt for each data disk. |
| | Set Password | Yes | Select when to set password. You can select Now or Later . If you select Later , you can use the Change Password feature in the console to set the password. For more information, see the Change Password topic in <i>ECS User Guide</i> . |

| Field | Parameter | Required | Description |
|----------------|--------------------|----------|---|
| Password | Logon Password | No | The password used to log on to the ECS instance. The password must be 8 to 30 characters in length and must contain at least three of the following character types: digits, uppercase letters, lowercase letters, and special characters.  Note The password is used to log on to the operating system and is not the VNC password. |
| | Confirm Password | No | Enter the password again. |
| Deployment Set | Deployment Set | No | The deployment set to which the instance belongs. |
| Instance Name | Configuration Name | No | The name of the scaling configuration. |
| | Instance Name | No | The name of the ECS instance. |
| User Data | User Data | No | Windows supports two formats: Bat and Powershell. Before you perform Base64 encoding, make sure to include <code>[bat]</code> or <code>[powershell]</code> as the first line. You can run shell scripts for Linux ECS instances. |
| Quantity | Quantity | No | The number of instances to purchase. |

7. Click **Submit**.

Result

After the scaling configuration is created, it is in the **Disabled** state and is displayed in your scaling configuration list. To automatically create ECS instances, you must apply a scaling configuration. For more information, see [Apply a scaling configuration](#).

4.3.5. Enable a scaling group

This topic describes how to enable a scaling group. You can enable a scaling group to trigger scaling activities.

Prerequisites

- The scaling group is in the **Disabled** state.
- The scaling group has a scaling configuration that is in the **Enable** state.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the top navigation bar, select an organization, a resource set, and a region.

3. Find the scaling group that you want to enable and click **Enable** in the **Actions** column.
4. In the message that appears, click **OK**.

Result

The status of the scaling group is changed from **Disabled** to **Enable** in the **Status** column.

4.3.6. Create a scaling rule

This topic describes how to create a scaling rule. You can create scaling rules to add or remove ECS instances. For example, you can add an ECS instance to a scaling group.

Context

- You can create only a limited number of scaling rules for a scaling group. For more information, see the Limits topic in *Auto Scaling Product Introduction*.
- After a scaling rule is executed, the resulting number of ECS instances in the scaling group may fall outside of the specified range. In this case, Auto Scaling automatically adjusts the number of ECS instances to ensure that the number of ECS instances in the scaling group is within the specified range.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **Scaling Rules**.
5. Click **Create Scaling Rule**.
6. Configure parameters for the scaling rule.

| Parameter | Required | Description |
|----------------------------|----------|--|
| Rule Name | Yes | The name of the scaling rule. It must be 2 to 64 characters in length and can contain letters, digits, underscores (_), hyphens (-), and periods (.). It must start with a letter or digit. |
| Scaling Activity | Yes | The operation that is performed when the scaling rule is triggered. The operations include: <ul style="list-style-type: none"> ◦ Change to N instances: After the scaling rule is executed, the number of instances in the scaling group is changed to N. ◦ Add N instances: After the scaling rule is executed, N instances are added to the scaling group. ◦ Remove N instances: After the scaling rule is executed, N instances are removed from the scaling group. |
| Default Cooldown (Seconds) | No | The cooldown period. If this parameter is not specified, the default value is used. |

7. Click **OK**.

4.3.7. Create a scheduled task

This topic describes how to create a scheduled task to scale computing resources in response to predictable business changes in the future. Scheduled tasks enable the system to automatically obtain sufficient computing resources before business peaks and release idle computing resources after the business peaks.

Context

A scheduled task is preconfigured to execute the specified scaling rule at the specified time. When the specified time arrives, the scheduled task automatically scales computing resources. This allows you to reduce costs and meet business requirements. You can also specify recurring schedules for scheduled tasks if business changes are regular.

If multiple scheduled tasks need to be executed in 1 minute, Auto Scaling executes the most recently created scheduled task.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the left-side navigation pane, choose **Event-driven Tasks > Timed tasks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. In the upper-right corner of the Scheduled Tasks page, click **Create Scheduled Task**.
5. In the dialog box that appears, configure parameters for the scheduled task.

| Parameter | Required | Description |
|--------------------------------|----------|--|
| Task Name | Yes | The name of the scheduled task. The name must be 2 to 64 characters in length and can contain letters, digits, underscores (_), hyphens (-), and periods (.). It must start with a letter or a digit. |
| Description | Yes | The description of the scheduled task. |
| Organization/Resource Group | Yes | The organization and resource set in which to create the scheduled task. |
| Start Time | Yes | The time to execute the scheduled task. |
| Scaling Rules | Yes | The scaling group to be monitored and the scaling rule to be executed. |
| Retry Interval (Seconds) | No | The period of time during which the system retries to execute the scheduled task. Unit: seconds. If a scaling activity fails to be executed at the specified time, Auto Scaling executes the scheduled task again within the period of time that is specified by the Retry Interval (Seconds) parameter. |
| Recurrence Settings (Advanced) | No | Specifies whether to execute the scheduled task on a recurring schedule. Select Recurrence Settings (Advanced) and set the Recurrence and Expire parameters. The valid values for Recurrence include Daily , Weekly , and Monthly . |

6. Click **OK**.

Result

The scheduled task that you created is displayed in the scheduled task list.

4.3.8. Create an event-triggered task

This topic describes how to create an event-triggered task associated with monitoring metrics in response to emergent or unpredictable business changes. After you create and enable an event-triggered task, Auto Scaling collects data for the specified metric in real time and triggers an alert when the specified condition is met. Then, Auto Scaling executes the corresponding scaling rule to scale Elastic Compute Service (ECS) instances in the scaling group.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the left-side navigation pane, choose **Event-driven Tasks > Alert Tasks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. In the upper-right corner of the Event-Triggered Tasks page, click **Create Alert**.
5. In the dialog box that appears, configure parameters for the event-triggered task.

| Parameter | Required | Description |
|---------------------------------|----------|--|
| Task Name | Yes | The name of the event-triggered task. It must be 2 to 64 characters in length and can contain letters, digits, underscores (_), hyphens (-), and periods (.). It must start with a letter or a digit. |
| Description | No | The description of the event-triggered task. |
| Organization/Resource Group | Yes | The organization and resource set in which to create the event-triggered task. |
| Monitoring Metrics/Scaling Rule | Yes | The scaling group to be monitored and the scaling rule to be executed. |
| Monitoring Type | Yes | System-Level Monitoring is selected by default. |
| Monitoring Metrics | Yes | The metrics that you want to monitor. Valid values: <ul style="list-style-type: none"> ◦ Average CPU Utilization ◦ Memory Usage ◦ Outbound Traffic ◦ Inbound Traffic ◦ Average System Load |
| Monitoring Period | Yes | The period during which data is aggregated and analyzed. The shorter the period, the higher the frequency that the alert is triggered. Unit: minutes. Valid values: <ul style="list-style-type: none"> ◦ 1 ◦ 2 ◦ 5 ◦ 15 |

| Parameter | Required | Description |
|---------------|----------|--|
| Statistic | Yes | <p>The rule that determines whether to trigger an alert. Select Average, Max Capacity, or Min Capacity, and specify a threshold value. For example, to trigger an alert when the CPU utilization exceeds 80%, you can use one of the following methods to specify the trigger condition:</p> <ul style="list-style-type: none"> ◦ Average: An alert is triggered when the average CPU utilization of all ECS instances in the scaling group exceeds 80%. ◦ Max Capacity: An alert is triggered when the highest CPU utilization among the ECS instances in the scaling group exceeds 80%. ◦ Min Capacity: An alert is triggered when the lowest CPU utilization among the ECS instances in the scaling group exceeds 80%. |
| Trigger After | Yes | <p>The number of consecutive times that the threshold must be exceeded before the alert is triggered. Valid values:</p> <ul style="list-style-type: none"> ◦ 1 ◦ 2 ◦ 3 ◦ 5 |

6. Click **OK**.

4.4. Scaling groups

4.4.1. Create a scaling group

This topic describes how to create a scaling group. A scaling group is a group of Elastic Compute Service (ECS) instances that is dynamically scaled based on the configured scenario. You can specify the minimum and maximum numbers of ECS instances in a scaling group.

Prerequisites

- A virtual private cloud (VPC) and a vSwitch are created. For more information, see [Create a VPC](#) and [Create a vSwitch](#) in *VPC User Guide*.
- To associate a scaling group with Server Load Balancer (SLB) instances, make sure that the following requirements are met:
 - You have one or more SLB instances in the **Running** state.
 - The SLB instances and the scaling group are in the same organization, resource set, and region.
- To associate a scaling group with ApsaraDB RDS instances, make sure that the following requirements are met:
 - You have one or more RDS instances in the **Running** state.
 - The RDS instances and the scaling group are in the same organization, resource set, and region.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the left-side navigation pane, click **Scaling groups**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. In the upper-right corner of the Scaling Groups page, click **Create Scaling Group**.
5. Configure parameters for the scaling group.

| Parameter | Required | Description |
|-----------------------------|----------|--|
| Scaling Group | Yes | The name of the scaling group. It must be 2 to 64 characters in length, and can contain letters, digits, underscores (_), hyphens (-), and periods (.). It must start with a letter or a digit. |
| Organization/Resource Group | Yes | The organization and resource set in which to create the scaling group. |
| Maximum Capacity | Yes | The maximum number of instances that the scaling group can contain. Set the value based on your business requirements to control costs. Valid values: 0 to 1000. |
| Minimum Capacity | Yes | The minimum number of instances that a scaling group must contain. Set the value based on your business requirements to ensure service availability. When the scaling group is enabled, Auto Scaling automatically creates this number of ECS instances. Valid values: 0 to 1000. |
| Cooldown (Seconds) | Yes | The period of time after each scaling activity is complete. During the cooldown time, Auto Scaling rejects all scaling activity requests triggered by event-triggered tasks from CloudMonitor. However, scaling activities triggered by other types of tasks such as manually triggered tasks and scheduled tasks are not limited by the cooldown time and are immediately executed. The value must be an integer that is greater than or equal to zero. Unit: seconds. |
| Scale-In Policy | No | The policy for automatically removing ECS instances from the scaling group. This parameter contains the Select and From the list, select fields. You cannot specify the same values for the two fields. Valid values: <ul style="list-style-type: none"> ◦ Earliest Instance Created Using Scaling Configuration: filters instances that were created based on the earliest scaling configuration. ◦ Earliest Created Instance: filters instances that were added to the scaling group at the earliest point in time. ◦ Newest Instances: filters instances that were most recently added to the scaling group. ◦ None: is available only for From the list, select. This value indicates that Auto Scaling does not filter instances based on the From the list, select field. For example, if Auto Scaling filters instances based on the Earliest Created Instance value of Select , you can select one of the following values for From the list, select : <ul style="list-style-type: none"> ◦ None: indicates that Auto Scaling does not filter instances based on the From the list, select field. ◦ Newest Instances: filters instances obtained based on the Select field and then filters instances that were most recently added to the scaling group. |

| Parameter | Required | Description |
|------------------------|----------|---|
| Region/VPC | Yes | The region and VPC in which to create the scaling group. |
| Vswitch | Yes | The ID of the vSwitch with which to associate the scaling group. |
| Associate SLB Instance | No | <p>After you associate SLB instances with the scaling group, ECS instances that are added to the scaling group are automatically added as SLB backend servers. You can specify a server group to which to add the ECS instances. ECS instances can be added to the following types of server groups:</p> <ul style="list-style-type: none"> ◦ Default server group: the group of ECS instances that are used to receive requests. If the listener is not configured with a vServer group or a primary/secondary server group, requests are forwarded to the ECS instances in the default server group. ◦ vServer group: If you want to distribute different requests to different backend servers or configure domain name- or URL-based routing methods, you can use vServer groups. |
| Associate RDS Instance | No | After you associate RDS instances with the scaling group, the internal IP addresses of ECS instances that are added to the scaling group are automatically added to the whitelists of the RDS instances to allow communication over the internal network. |

6. Click **OK**.

Result

The created scaling group is displayed in the scaling group list but is in the **Disabled** state. To enable the scaling group, you must create a scaling configuration. For more information, see [Create a scaling configuration](#).

4.4.2. Enable a scaling group

This topic describes how to enable a scaling group. You can enable a scaling group to trigger scaling activities.

Prerequisites

- The scaling group is in the **Disabled** state.
- The scaling group has a scaling configuration that is in the **Enable** state.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the scaling group that you want to enable and click **Enable** in the **Actions** column.
4. In the message that appears, click **OK**.

Result

The status of the scaling group is changed from **Disabled** to **Enable** in the **Status** column.

4.4.3. View scaling groups

This topic describes how to view the scaling group list and the details of a specific scaling group.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the top navigation bar, select an organization, a resource set, and a region.
The scaling groups that correspond to the specified organization, resource set, and region are displayed.
3. Select a filter option, enter the corresponding information, and then click **Search**.

You can select multiple filter options to narrow down the search results.

| Option | Description |
|------------------|---|
| Scaling Group | Enter a scaling group name to search for the scaling group. |
| Scaling Group ID | Enter a scaling group ID to search for the scaling group. |

4. Click the name of the scaling group in the **Scaling Group Name/ID** column.
5. View the details of the specified scaling group.

| Parameter | Description |
|-----------------------|--|
| Basic Information | The configurations of the scaling group, such as the scaling group ID, scaling group name, total instances, minimum number of instances, maximum number of instances, and scale-in policy. |
| ECS Instances | The details of ECS instances, such as the list of automatically created ECS instances, the list of manually added ECS instances, and the number of ECS instances that are in service. |
| Scaling Activities | All the scaling activities that have been executed in the scaling group. |
| Scaling Configuration | The information of scaling configurations in the scaling group. |
| Scaling Rules | The information of scaling rules. |

4.4.4. Modify a scaling group

This topic describes how to modify a scaling group. You can modify the parameters of a specific scaling group, such as the minimum and maximum numbers of ECS instances.

Context

After you modify the minimum or maximum number of ECS instances that a scaling group can have, if the number of instances in the scaling group is outside this range, Auto Scaling automatically creates or removes ECS instances until the number of instances are within the range.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click **Edit** in the **Actions** column.
4. Modify the parameters of the scaling group.

You can modify the scaling configuration and other parameters, but not the organization and resource set. For

more information about other parameters, see [Create a scaling group](#).

5. Click **OK**.

4.4.5. Disable a scaling group

This topic describes how to disable a scaling group.

Prerequisites

- The scaling group does not have scaling activities in progress.
- The scaling group is in the **Enable** state.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click **Disable** in the **Actions** column.
4. Click **OK**.

Result

The status of the scaling group is changed from **Enable** to **Disabled** in the **Status** column.

4.4.6. Delete a scaling group

This topic describes how to delete a scaling group. When you delete a scaling group, Auto Scaling removes and releases ECS instances that are automatically created, removes ECS instances that are manually added, and deletes the scaling configurations and rules in the scaling group. However, the scheduled tasks and event-triggered tasks that are associated with the scaling group are not deleted.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click **Delete** in the **Actions** column.
4. Click **OK**.

4.4.7. Query ECS instances

You can query all ECS instances in a scaling group and their states.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the left-side navigation pane, click **Scaling groups**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
5. In the left-side navigation pane, click **ECS instances**.
6. View the details of ECS instances.

| Category | Description |
|-------------------------------------|--|
| Automatically created ECS instances | The ECS instances that are automatically created based on the active scaling configuration when a scaling rule is triggered. |

| Category | Description |
|--|---|
| Manually added ECS instances | The ECS instances that are manually added to the specified scaling group. |
| The number of ECS instances in each state. | <p>The following section describes the states:</p> <ul style="list-style-type: none"> ◦ Total: all ECS instances in the scaling group ◦ In Service: the ECS instances that are in normal use ◦ On Standby: the ECS instances that are on standby ◦ Protected: the ECS instances that are protected ◦ Adding: the ECS instances that are being added to the scaling group ◦ Removing: the ECS instances that are being removed from the scaling group <div style="background-color: #e0f2f7; padding: 5px; margin-top: 10px;"> <p> Note The Disabled, Adding:wait, and Suspending states are unavailable.</p> </div> |

4.4.8. Put an ECS instance into the Standby state

This topic describes how to put an ECS instance into the Standby state. Auto Scaling does not perform health checks on or release ECS instances in the Standby state.

Context

After an ECS instance is put into the Standby state:

- The ECS instance stays in the Standby state until you change its status.
- Auto Scaling stops managing the lifecycle of the ECS instance. You must manually manage the lifecycle of the ECS instance.
- If a scale-in event is triggered, Auto Scaling will not remove the ECS instance.
- When the ECS instance is stopped or restarted, its health check status is not affected.
- To release the ECS instance, you must first remove it from the scaling group.
- If you delete the scaling group, the ECS instance is automatically put out of the Standby state and is released.
- You can also perform other operations on the ECS instance, such as stopping, restarting, changing the instance type of, or changing the operating system of the ECS instance.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **ECS Instances**.
5. Select the source of the ECS instance.
 - To select an automatically created ECS instance, click the **Auto Created** tab.
 - To select a manually added ECS instance, click the **Manually Added** tab.
6. Find the target ECS instance and choose **Actions > Switch to Standby** in the **Actions** column.
7. Click **OK**.

4.4.9. Remove an ECS instance from the Standby state

This topic describes how to remove an ECS instance from the Standby state. You can remove an instance from the Standby state to reuse it.

Context

After an ECS instance is removed from the Standby state:

- The ECS instance enters the In Service state.
- When the ECS instance is stopped or restarted, its health status is updated.
- Auto Scaling continues to manage the lifecycle of the ECS instance, and can remove the ECS instance from the scaling group during a scale-in event.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **ECS Instances**.
5. Select the source of the ECS instance.
 - To select an automatically created ECS instance, click the **Auto Created** tab.
 - To select a manually added ECS instance, click the **Manually Added** tab.
6. Find the target ECS instance and choose **Actions > Move Out Of Standby** in the **Actions** column.
7. Click **OK**.

4.4.10. Put an ECS instance into the Protected state

This topic describes how to put an ECS instance into the Protected state. Auto Scaling does not perform health checks on or release ECS instances that are in the Protected state.

Context

After an ECS instance is put into the Protected state:

- The ECS instance stays in the Protected state until you change its status.
- If a scale-in event is triggered, Auto Scaling will not remove the ECS instance. To release the ECS instance, you must remove the ECS instance from the Protected state and then remove it from the scaling group.
- When the ECS instance is stopped or restarted, its health check status is not affected.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **ECS Instances**.
5. Select the source of the ECS instance.
 - To select an automatically created ECS instance, click the **Auto Created** tab.
 - To select a manually added ECS instance, click the **Manually Added** tab.
6. Find the target ECS instance and choose **Actions > Switch to Protection** in the **Actions** column.

7. Click **OK**.

4.4.11. Remove an ECS instance from the Protected state

This topic describes how to remove an ECS instance from the Protected state. After an ECS instance is removed from the Protected state, Auto Scaling continues to manage the lifecycle of the ECS instance.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **ECS Instances**.
5. Select the source of the ECS instance.
 - To select an automatically created ECS instance, click the **Auto Created** tab.
 - To select a manually added ECS instance, click the **Manually Added** tab.
6. Find the target ECS instance and choose **Actions > Move Out Of Protection** in the **Actions** column.
7. Click **OK**.

4.5. Scaling configurations

4.5.1. Create a scaling configuration

This topic describes how to create a scaling configuration for a scaling group.

Prerequisites

At least one security group is available. If you do not have any security groups, create a security group. For more information, see [Create a security group](#) in *ECS User Guide*.

Context

You can create only a limited number of scaling configurations for a scaling group. For more information, see the [Limits](#) topic in *Auto Scaling Product Introduction*.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **Scaling Configuration**.
5. Choose **Create > Create Scaling Configuration**.
6. Configure parameters for the scaling configuration.

| Section | Parameter | Required | Description |
|---------|-----------|----------|---|
| Region | Region | Yes | The region where the ECS instance is located. |
| | Zone | Yes | The zone where the ECS instance is located. |

| Section | Parameter | Required | Description |
|----------------|-----------------|----------|---|
| Security Group | Security Group | Yes | The security group to which the ECS instance belongs. |
| Instance | Instance Family | Yes | The instance family to which the ECS instance belongs. |
| | Instance Type | Yes | The instance type of the ECS instance. |
| Image | Image Type | Yes | <ul style="list-style-type: none"> ◦ Public Image: Public images provided by Alibaba Cloud are fully licensed to offer a secure and stable operating environment for applications on ECS instances. ◦ Custom Image: You can create custom images to install software or deploy projects that have special requirements. |
| Storage | System Disk | Yes | Specify the category and size of the system disk. The operating system is installed on the system disk. You can select Ultra Disk or Standard SSD . |
| | Data Disk | No | Specify the category and size of the data disk. You can select Ultra Disk or Standard SSD . You can add a maximum of 16 data disks. The maximum capacity of each data disk is 32 TiB. You can set Release with Instance and Encrypt for each data disk. |
| Password | Set Password | Yes | <p>Select when to set password. You can select Now or Later.</p> <p>If you select Later, you can use the Change Password feature in the console to set the password. For more information, see the Change Password topic in <i>ECS User Guide</i>.</p> |
| | | | |

| Section | Parameter | Required | Description |
|----------------|--------------------|----------|---|
| | Logon Password | No | The password used to log on to the ECS instance. The password must be 8 to 30 characters in length and must contain at least three of the following character types: digits, uppercase letters, lowercase letters, and special characters.  Note The password is used to log on to the operating system and is not the VNC password. |
| | Confirm Password | No | Enter the password again. |
| Deployment Set | Deployment Set | No | The deployment set to which the instance belongs. |
| Instance Name | Configuration Name | No | The name of the scaling configuration. |
| | Instance Name | No | The name of the ECS instance. |
| User Data | User Data | No | Windows supports two formats: Bat and Powershell. Before you perform Base64 encoding, make sure to include <code>[bat]</code> or <code>[powershell]</code> as the first line. You can run shell scripts for Linux ECS instances. |
| Quantity | Quantity | No | The number of instances to purchase. |

7. Click **Submit**.

Result

After the scaling configuration is created, it is in the **Disabled** state and is displayed in your scaling configuration list. To automatically create ECS instances, you must apply a scaling configuration. For more information, see [Apply a scaling configuration](#).

4.5.2. View scaling configurations

This topic describes how to view scaling configurations.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the top navigation bar, select an organization, a resource set, and a region.
The scaling groups that correspond to the specified organization, resource set, and region are displayed.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **Scaling Configuration**.
5. View the list of scaling configurations.

4.5.3. Modify a scaling configuration

This topic describes how to modify a scaling configuration. You can modify the parameters of a scaling configuration based on your actual needs.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **Scaling Configuration**.
5. Find the target scaling configuration and click its name in the **Scaling Configuration Name/ID** column.
6. Modify the parameters of the scaling configuration.

For more information about parameters of the scaling configuration, see [Create a scaling configuration](#).

7. Click **OK**.

4.5.4. Apply a scaling configuration

This topic describes how to apply a scaling configuration. You can create multiple scaling configurations for a scaling group and apply one.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **Scaling Configuration**.
5. Find the target scaling configuration and click **Select** in the **Actions** column.

Only one scaling configuration can be in the **Enabled** state in a scaling group. After a scaling configuration is applied, other scaling configurations are put into the **Disabled** state.

6. Click **OK**.

Result

The status of the scaling configuration changes from **Disabled** to **Enable** in the **Status** column.

4.5.5. Delete a scaling configuration

This topic describes how to delete a scaling configuration that is no longer needed. After you delete a scaling configuration, existing ECS instances that are created from the scaling configuration are not removed.

Prerequisites

The scaling configuration is in the **Disabled** state.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **Scaling Configuration**.
5. Find the target scaling configuration and click **Delete** in the **Actions** column.

6. Click **OK**.

4.6. Scaling rules

4.6.1. Create a scaling rule

This topic describes how to create a scaling rule. You can create scaling rules to add or remove ECS instances. For example, you can add an ECS instance to a scaling group.

Context

- You can create only a limited number of scaling rules for a scaling group. For more information, see the Limits topic in *Auto Scaling Product Introduction*.
- After a scaling rule is executed, the resulting number of ECS instances in the scaling group may fall outside of the specified range. In this case, Auto Scaling automatically adjusts the number of ECS instances to ensure that the number of ECS instances in the scaling group is within the specified range.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **Scaling Rules**.
5. Click **Create Scaling Rule**.
6. Configure parameters for the scaling rule.

| Parameter | Required | Description |
|----------------------------|----------|--|
| Rule Name | Yes | The name of the scaling rule. It must be 2 to 64 characters in length and can contain letters, digits, underscores (_), hyphens (-), and periods (.). It must start with a letter or digit. |
| Scaling Activity | Yes | The operation that is performed when the scaling rule is triggered. The operations include: <ul style="list-style-type: none"> ◦ Change to N instances: After the scaling rule is executed, the number of instances in the scaling group is changed to N. ◦ Add N instances: After the scaling rule is executed, N instances are added to the scaling group. ◦ Remove N instances: After the scaling rule is executed, N instances are removed from the scaling group. |
| Default Cooldown (Seconds) | No | The cooldown period. If this parameter is not specified, the default value is used. |

7. Click **OK**.

4.6.2. View scaling rules

This topic describes how to view scaling rules.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the top navigation bar, select an organization, a resource set, and a region.
The scaling groups that correspond to the specified organization, resource set, and region are displayed.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **Scaling Rules**.
5. View the list of scaling rules.

4.6.3. Modify a scaling rule

This topic describes how to modify a scaling rule. You can modify the following parameters of a scaling rule: Rule Name, Scaling Activity, and Default Cooldown.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **Scaling Rules**.
5. Find the target scaling rule and click **Edit** in the **Actions** column.
6. Modify the Rule Name, Scaling Activity, and Default Cooldown parameters.
7. Click **OK**.

4.6.4. Delete a scaling rule

This topic describes how to delete a scaling rule that is no longer needed.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **Scaling Rules**.
5. Find the target scaling rule and click **Delete** in the **Actions** column.
6. In the message that appears, click **OK**.

4.7. Scaling tasks

4.7.1. Manually execute a scaling rule

This topic describes how to manually execute a scaling rule to add or remove ECS instances.

Prerequisites

- The scaling group to which the scaling rule belongs is in the **Enable** state.
- No scaling activity is in progress in the scaling group to which the scaling rule belongs.

Context

After the scaling rule is executed, if the number of ECS instances is greater than the maximum number or less than the minimum number, Auto Scaling automatically adjusts the number of ECS instances to be within the valid range.

Auto Scaling enables you to manually execute scaling rules. You can also associate an event-triggered task or scheduled task with the scaling rule to automatically adjust the number of ECS instances. For more information, see [Create a scheduled task](#) and [Create an event-triggered task](#).

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **Scaling Rules**.
5. Find the scaling rule that you want to execute and click **Run** in the **Actions** column.
6. In the message that appears, click **OK**.

Result

The **Scaling Activities** page appears. You can view the details of your scaling activity.

4.7.2. Manually add an ECS instance

This topic describes how to manually add an ECS instance to a scaling group. You can add existing ECS instances to a scaling group to take full advantage of the computing resources.

Prerequisites

The ECS instance to be added must meet the following conditions:

- The ECS instance and the scaling group to which to add the instance share the same region, organization, and resource set.
- The ECS instance is in the **Running** state.
- The ECS instance does not belong to any scaling groups.
- The ECS instance and the scaling group are in the same VPC.

The scaling group to which to add the ECS instance must meet the following conditions:

- The scaling group is in the **Enable** state.
- No scaling activity is in progress in the scaling group.

Context

- When no scaling activity is being executed in the scaling group, you can add an ECS instance to the scaling group without the need to wait for the cooldown time to expire.
- If the number of instances in the scaling group is greater than the maximum number of instances after an ECS instance is added to the scaling group, the ECS instance cannot be added.
- The ECS instances that are manually added to a scaling group are not limited by scaling configurations. The instance types of the manually added instances can be different from that of the scaling configuration in the **Enable** state.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **ECS Instances**.
5. Click **Add Instance**.
6. Select the ECS instance to be added and click **OK**.

Result

The manually added instance is displayed on the **Manually Added** tab.

4.7.3. Manually remove an ECS instance

This topic describes how to manually remove an ECS instance that is no longer needed from a scaling group.

Prerequisites

The scaling group must meet the following conditions:

- The scaling group is in the **Enable** state.
- No scaling activity is in progress in the scaling group.

Context

- When no scaling activity is being executed in the scaling group, you can immediately remove an ECS instance from the scaling group without the need to wait for the cooldown time to expire.
- After an ECS instances is removed from a scaling group, the number of instances in the scaling group must be greater than or equal to the minimum number of instances. Otherwise, the ECS instance cannot be removed.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the top navigation bar, select an organization, a resource set, and a region.
3. Find the target scaling group and click the name of the scaling group in the **Scaling Group Name/ID** column.
4. In the left-side navigation pane, click **ECS Instances**.
5. Select the source of the ECS instance.
 - To select an automatically created ECS instance, click the **Auto Created** tab.
 - To select a manually added ECS instance, click the **Manually Added** tab.
6. Use one of the following methods to remove one or more ECS instances from a scaling group:
Manually added ECS instances can only be removed, but cannot be released.
 - Find the ECS instance that you want to remove and choose **Actions > Remove from Scaling Group** in the **Actions** column.
 - Find the ECS instance that you want to remove and release, and choose **Actions > Remove from Scaling Group and Release** in the **Actions** column.
7. In the message that appears, click **OK**.

4.8. Scheduled tasks

4.8.1. Create a scheduled task

This topic describes how to create a scheduled task to scale computing resources in response to predictable business changes in the future. Scheduled tasks enable the system to automatically obtain sufficient computing resources before business peaks and release idle computing resources after the business peaks.

Context

A scheduled task is preconfigured to execute the specified scaling rule at the specified time. When the specified time arrives, the scheduled task automatically scales computing resources. This allows you to reduce costs and meet business requirements. You can also specify recurring schedules for scheduled tasks if business changes are regular.

If multiple scheduled tasks need to be executed in 1 minute, Auto Scaling executes the most recently created scheduled task.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the left-side navigation pane, choose **Event-driven Tasks > Timed tasks.**
3. In the top navigation bar, select an organization, a resource set, and a region.
4. In the upper-right corner of the Scheduled Tasks page, click **Create Scheduled Task.**
5. In the dialog box that appears, configure parameters for the scheduled task.

| Parameter | Required | Description |
|--------------------------------|----------|--|
| Task Name | Yes | The name of the scheduled task. The name must be 2 to 64 characters in length and can contain letters, digits, underscores (_), hyphens (-), and periods (.). It must start with a letter or a digit. |
| Description | Yes | The description of the scheduled task. |
| Organization/Resource Group | Yes | The organization and resource set in which to create the scheduled task. |
| Start Time | Yes | The time to execute the scheduled task. |
| Scaling Rules | Yes | The scaling group to be monitored and the scaling rule to be executed. |
| Retry Interval (Seconds) | No | The period of time during which the system retries to execute the scheduled task. Unit: seconds. If a scaling activity fails to be executed at the specified time, Auto Scaling executes the scheduled task again within the period of time that is specified by the Retry Interval (Seconds) parameter. |
| Recurrence Settings (Advanced) | No | Specifies whether to execute the scheduled task on a recurring schedule. Select Recurrence Settings (Advanced) and set the Recurrence and Expire parameters. The valid values for Recurrence include Daily , Weekly , and Monthly . |

6. Click **OK.**

Result

The scheduled task that you created is displayed in the scheduled task list.

4.8.2. View scheduled tasks

This topic describes how to view scheduled tasks.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the left-side navigation pane, choose **Event-driven Tasks > Timed tasks.**
3. In the top navigation bar, select an organization, a resource set, and a region.
The scheduled tasks that correspond to the specified organization, resource set, and region are displayed.
4. Select a filter option, enter the corresponding information, and then click **Search.**
You can select multiple filter options to narrow down the search results.

| Option | Description |
|-----------|---|
| Task Name | Enter a task name to search for the scheduled task. |
| Task ID | Enter a task ID to search for the scheduled task. |

5. View the scheduled task list.

4.8.3. Modify a scheduled task

This topic describes how to modify a scheduled task. You can modify parameters such as Start Time, Scaling Rules, and Retry Expiry Time for a scheduled task.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the left-side navigation pane, choose **Event-driven Tasks > Timed tasks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the scheduled task that you want to modify and click **Edit** in the **Actions** column.
5. Modify the parameters of the scheduled task.

You can modify the Recurrence and Expire parameters if you have enabled the Recurrence Settings (Advanced) feature when you create the scheduled task, but the Recurrence Settings (Advanced) feature cannot be disabled. For more information about other parameters of the scheduled task, see [Create a scheduled task](#).

6. Click **OK**.

4.8.4. Disable a scheduled task

This topic describes how to disable a scheduled task. You can disable a scheduled task that is no longer needed.

Prerequisites

The scheduled task is in the **Running** state.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the left-side navigation pane, choose **Event-driven Tasks > Timed tasks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the scheduled task that you want to disable and click **Disabled** in the **Actions** column.
5. In the message that appears, click **OK**.

Result

The status of the scheduled task is changed from **Running** to **Stop** in the **Status** column.

4.8.5. Enable a scheduled task

This topic describes how to enable a scheduled task. You can enable a scheduled task that has been disabled and use it to trigger scaling activities at the specified time point.

Prerequisites

The scheduled task is in the **Stop** state.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the left-side navigation pane, choose **Event-driven Tasks > Timed tasks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the scheduled task that you want to enable and click **Enable** in the **Actions** column.
5. In the message that appears, click **OK**.

Result

The status of the scheduled task is changed from **Stop** to **Running** in the **Status** column.

4.8.6. Delete a scheduled task

This topic describes how to delete a scheduled task that is no longer needed.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the left-side navigation pane, choose **Event-driven Tasks > Timed tasks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the scheduled task that you want to delete and click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

4.9. Event-triggered tasks

4.9.1. Create an event-triggered task

This topic describes how to create an event-triggered task associated with monitoring metrics in response to emergent or unpredictable business changes. After you create and enable an event-triggered task, Auto Scaling collects data for the specified metric in real time and triggers an alert when the specified condition is met. Then, Auto Scaling executes the corresponding scaling rule to scale Elastic Compute Service (ECS) instances in the scaling group.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the left-side navigation pane, choose **Event-driven Tasks > Alert Tasks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. In the upper-right corner of the Event-Triggered Tasks page, click **Create Alert**.
5. In the dialog box that appears, configure parameters for the event-triggered task.

| Parameter | Required | Description |
|---------------------------------|----------|---|
| Task Name | Yes | The name of the event-triggered task. It must be 2 to 64 characters in length and can contain letters, digits, underscores (_), hyphens (-), and periods (.). It must start with a letter or a digit. |
| Description | No | The description of the event-triggered task. |
| Organization/Resource Group | Yes | The organization and resource set in which to create the event-triggered task. |
| Monitoring Metrics/Scaling Rule | Yes | The scaling group to be monitored and the scaling rule to be executed. |
| Monitoring Type | Yes | System-Level Monitoring is selected by default. |

| Parameter | Required | Description |
|--------------------|----------|--|
| Monitoring Metrics | Yes | The metrics that you want to monitor. Valid values: <ul style="list-style-type: none"> ◦ Average CPU Utilization ◦ Memory Usage ◦ Outbound Traffic ◦ Inbound Traffic ◦ Average System Load |
| Monitoring Period | Yes | The period during which data is aggregated and analyzed. The shorter the period, the higher the frequency that the alert is triggered. Unit: minutes. Valid values: <ul style="list-style-type: none"> ◦ 1 ◦ 2 ◦ 5 ◦ 15 |
| Statistic | Yes | The rule that determines whether to trigger an alert. Select Average , Max Capacity , or Min Capacity , and specify a threshold value. For example, to trigger an alert when the CPU utilization exceeds 80%, you can use one of the following methods to specify the trigger condition: <ul style="list-style-type: none"> ◦ Average: An alert is triggered when the average CPU utilization of all ECS instances in the scaling group exceeds 80%. ◦ Max Capacity: An alert is triggered when the highest CPU utilization among the ECS instances in the scaling group exceeds 80%. ◦ Min Capacity: An alert is triggered when the lowest CPU utilization among the ECS instances in the scaling group exceeds 80%. |
| Trigger After | Yes | The number of consecutive times that the threshold must be exceeded before the alert is triggered. Valid values: <ul style="list-style-type: none"> ◦ 1 ◦ 2 ◦ 3 ◦ 5 |

6. Click **OK**.

4.9.2. View event-triggered tasks

This topic describes how to view event-triggered tasks.

Procedure

1. [Log on to the Auto Scaling console](#).
2. In the left-side navigation pane, choose **Event-driven Tasks > Alert Tasks**.
3. In the top navigation bar, select an organization, a resource set, and a region. The event-triggered tasks that correspond to the specific organization, resource set, and region are displayed.
4. Select a filter option, enter the corresponding information, and then click **Search**.

You can select multiple filter options to narrow down the search results.

| Option | Description |
|------------------|--|
| Alert Name | Enter an event-triggered task name to search for the event-triggered task. |
| Scaling Group ID | Enter a scaling group ID to search for the event-triggered task associated with the scaling group. |

4.9.3. Modify an event-triggered task

This topic describes how to modify an event-triggered task. You can modify parameters such as Scaling Rules, Monitoring Type, and Statistic for an event-triggered task.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the left-side navigation pane, choose **Event-driven Tasks > Alert Tasks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the event-triggered task that you want to modify and click **Edit** in the **Actions** column.
5. Modify the parameters of the event-triggered task.

For more information about other parameters of the scheduled task, see [Create an event-triggered task](#). The following parameters cannot be modified:

- Organization
 - Resource Group
 - Monitoring Metrics
 - Monitoring Period
6. Click **OK**.

4.9.4. Disable an event-triggered task

This topic describes how to disable an event-triggered task. You can disable an event-triggered task if you no longer want to use it to trigger scaling activities.

Prerequisites

The event-triggered task is in the **Normal**, **Alerts**, or **Insufficient Data** state.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the left-side navigation pane, choose **Event-driven Tasks > Alert tasks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the event-triggered task that you want to disable and click **Disable** in the **Actions** column.
5. In the message that appears, click **OK**.

Result

The status of the event-triggered task is changed to **Stopped** in the **Status** column.

4.9.5. Enable an event-triggered task

This topic describes how to enable an event-triggered task. You can enable an event-triggered task that has been disabled to continue to monitor metrics and trigger scaling activities for a scaling group.

Prerequisites

The event-triggered task is in the **Stopped** state.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the left-side navigation pane, choose **Event-driven Tasks > Alert tasks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the event-triggered task that you want to enable and click **Enable** in the **Actions** column.
5. In the message that appears, click **OK**.

Result

The status of the event-triggered task changes from **Stopped** to **Normal** in the **Status** column.

4.9.6. Delete an event-triggered task

This topic describes how to delete an event-triggered task. You can delete an event-triggered task that is no longer needed.

Procedure

1. [Log on to the Auto Scaling console.](#)
2. In the left-side navigation pane, choose **Event-driven Tasks > Alert Tasks**.
3. In the top navigation bar, select an organization, a resource set, and a region.
4. Find the event-triggered task that you want to delete and click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

5.Resource Orchestration Service (ROS)

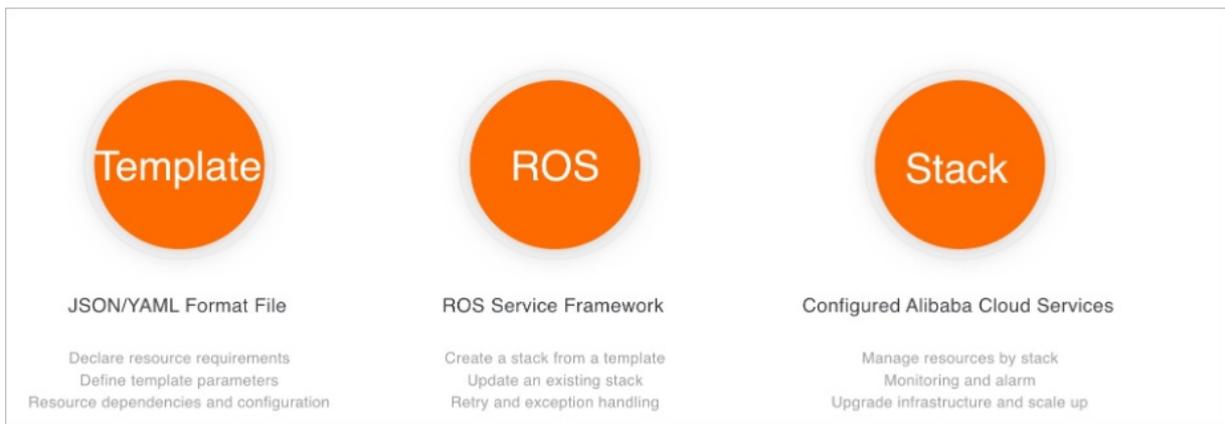
5.1. What is ROS?

Resource Orchestration Service (ROS) is an Apsara Stack service that can simplify the management of cloud computing resources. You can author stack templates based on the template specifications defined in ROS. Within a template, you can define required cloud computing resources such as Elastic Compute Service (ECS) and ApsaraDB RDS instances, and the dependencies between resources. The ROS engine automatically creates and configures all resources in a stack based on a template, which makes automatic deployment and O&M possible.

An ROS template is a readable, easy-to-author text file. You can directly edit a JSON-formatted template or use version control tools such as SVN and Git to control the template and infrastructure versions. You can use APIs and SDKs to integrate the orchestration capabilities of ROS with your own applications to implement Infrastructure as Code (IaC).

ROS templates are also a standardized way to deliver resources and applications. If you are an independent software vendor (ISV), you can use ROS templates to deliver a holistic system or solution that encompasses cloud resources and applications. ISVs can use this method to integrate Apsara Stack resources with their own software systems for centralized delivery.

ROS manages a group of cloud resources as a single unit called a stack. A stack is a group of Apsara Stack resources. You can create, delete, and clone cloud resources by stack.



5.2. Log on to the ROS console

This topic describes how to log on to the Resource Orchestration Service (ROS) console.

Prerequisites

- The URL, username, and password of the Apsara Uni-manager Management Console is obtained from the operations administrator before you log on.
- A browser is available. We recommend that you use Google Chrome.

Procedure

1. In the address bar, enter the URL used to log on to the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter the correct username and password.

The first time you log on to the Apsara Uni-manager Management Console, you must change the password of your username. For higher security, the password must meet the minimum complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase and lowercase letters

- Digits
 - Special characters including exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%)
3. Click **Login**.
 4. In the top navigation bar, choose **Products > Elastic Computing > Resource Orchestration Service**.

5.3. Create a stack

This topic describes how to create a stack in the Resource Orchestration Service (ROS) console.

Procedure

1. [Log on to the ROS console](#).
2. In the upper-right corner of the page, click **Create Stack**.
3. In the **Select Template** step, set **Organization**, **Resource Set**, and **regionId**.
4. In the **Prepare Template** section, enter template content in the JSON format. Click **Next**.
5. In the **Configure Template Parameters** step, specify the stack name and parameters, and click **Next**.

 **Note** The parameters that you need to specify vary based on the templates. Specify the parameters as instructed in the ROS console.

6. In the **Configure Stack** step, set **Rollback on Failure** and **Timeout Period**, and click **Next**.
7. In the **Confirm** step, check the template and stack configurations, and click **Create Stack**.

What to do next

On the **Stacks** page in the ROS console, you can perform the following operations:

- To delete a stack, click **Delete** in the Actions column corresponding to the stack.
- To update a stack, click **Update** in the Actions column corresponding to the stack.
- To recreate a stack, click **Recreate** in the Actions column corresponding to the stack.

 **Note**

- If you need only to modify the current template and configurations of a specified stack but do not need to change the region where the stack resides, update the stack.
- If you need to modify the current template and configurations of a specified stack and change the region where the stack resides, recreate the stack.

5.4. Template syntax

5.4.1. Template structure

A template is a UTF-8 encoded JSON file that is used to create stacks. Templates serve as the blueprint for underlying infrastructure and architecture. Templates define the configurations of Apsara Stack resources and dependencies between the resources.

ROS template structure

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Description": "The description of the template, which is used to provide information such as use scenarios and architecture of the template",
  "Metadata": {
    // The template metadata that provides information such as the layout for visualizations.
  },
  "Parameters": {
    // The parameters that you can specify when you create a stack.
  },
  "Mappings": {
    // The mapping tables. Mapping tables are nested tables.
  },
  "Conditions": {
    // The conditions defined by using internal condition functions. These conditions determine when to create associated resources.
  },
  "Resources": {
    // The detailed information of resources such as configurations and dependencies.
  },
  "Outputs": {
    // The outputs that are used to provide information such as resource properties. You can use the Resource Orchestration Service (ROS) console or API to obtain the information.
  }
}
```

ROSTemplateFormatVersion

Required. The template versions supported by ROS. Current version: 2015-09-01.

Description

Optional. The description of the template, which is used to provide information such as use scenarios and architecture of the template.

A detailed description can help users better understand the content of the template.

Metadata

Optional. The metadata of the template, in the JSON format.

Parameters

Optional. The parameters that you can specify when you create a stack. An Elastic Compute Service (ECS) instance type is often defined as a parameter. Parameters have default values. Parameters can improve the flexibility and reusability of the template. When you create a stack, select appropriate specifications.

Mappings

Optional. Mappings are defined as nested mapping tables. You can use `Fn::FindInMap` to retrieve values corresponding to keys. You can also use parameter values as keys. For example, you can search the region-image mapping table for desired images by region.

Conditions

Optional. The conditions defined by using `Fn::And`, `Fn::Or`, `Fn::Not`, and `Fn::Equals`. Multiple conditions are separated by commas (,). The system evaluates all conditions in the template before it creates or updates a stack. All resources associated with `true` conditions are created, and all resources associated with `false` conditions are ignored.

Resources

Optional. The detailed information of resources in the stack created based on the template. The information includes resource dependencies and configurations.

Outputs

Optional. The outputs that are used to provide information such as resource properties. You can use the ROS console or API to obtain the information.

5.4.2. Parameters

When you create a template, you can use the Parameters section to improve the flexibility and reusability of the template. When you create a stack, you can replace parameter values in the template.

For example, you have a web application requiring a stack that contains one Server Load Balancer (SLB) instance, two Elastic Compute Service (ECS) instances, and one ApsaraDB RDS instance. If the web application has a heavy workload, you can select ECS instances with advanced specifications when you create the stack. Otherwise, you can select ECS instances with basic specifications. The following example shows how to define the InstanceType parameter for an ECS instance:

```
"Parameters": {
  "InstanceType": {
    "Type": "String",
    "AllowedValues": ["ecs.t1.small", "ecs.s1.medium", "ecs.m1.medium", "ecs.c1.large"],
    "Default": "ecs.t1.small",
    "Label": "ECS instance type",
    "Description": "The type of the ECS instance that you want to create. Default value: ecs.t1.small. Valid values: ecs.t1.small, ecs.s1.medium, ecs.m1.medium, and ecs.c1.large."
  }
}
```

You can assign a value to the InstanceType parameter when you create stacks based on templates. If this parameter is not specified, the default value `ecs.t1.small` is used.

The following example shows how to reference the InstanceType parameter when you define a resource:

```
"Webserver": {
  "Type": "ALIYUN::ECS::Instance",
  "InstanceType": {
    "Ref": "InstanceType"
  }
}
```

Syntax

Each parameter consists of a name and properties. The parameter name can contain only letters and digits, and must be unique within the template. You can use the `Label` field to define the alias of a parameter.

The following table describes the parameter properties.

| Parameter property | Required | Description |
|--------------------|----------|-------------|
|--------------------|----------|-------------|

| Parameter property | Required | Description |
|-----------------------|----------|---|
| Type | Yes | <p>The data type of the parameter.</p> <ul style="list-style-type: none"> String: a string value. Example: <code>"ecs.s1.medium"</code>. Number: an integer or a floating-point number. Example: <code>3.14</code>. CommaDelimitedList: a set of strings separated by commas (<code>,</code>), which can be indexed by using the <code>Fn::Select</code> function. Example: <code>"80,foo,bar"</code>. Json: a JSON string. Example: <code>{"foo":"bar"}</code>. Boolean: a Boolean value. Example: <code>true</code> or <code>false</code>. |
| Default | No | The default value of the parameter. If you do not specify a value when you create a stack, Resource Orchestration Service (ROS) checks whether a default value is defined in the template. If a default value is found, ROS uses the default value. Otherwise, an error is returned. |
| AllowedValues | No | The list of one or more valid parameter values. |
| AllowedPattern | No | The regular expression that is used to check whether the specified parameter value is a string. If the input is not a string, an error is returned. |
| MaxLength | No | The integer value that determines the longest string allowed for a String-type parameter. |
| MinLength | No | The integer value that determines the shortest string allowed for a String-type parameter. |
| MaxValue | No | The numeric value that determines the maximum value allowed for a Number-type parameter. |
| MinValue | No | The numeric value that determines the minimum value allowed for a Number-type parameter. |
| NoEcho | No | Specifies whether to mask the parameter value when the <code>GetStack</code> operation is called. If you set this property to <code>true</code> , only asterisks (<code>*</code>) are returned. |
| Description | No | The string that describes the parameter. |
| ConstraintDescription | No | The description of the parameter constraints. |
| Label | No | The alias of the parameter, encoded in UTF-8. When you create a web form based on a template, the <code>Label</code> value can be mapped to the parameter name. |

| Parameter property | Required | Description |
|---------------------|----------|---|
| AssociationProperty | No | <p>The associated resource property. If you specify this parameter property, ROS verifies whether the specified parameter value is valid and provides a list of valid values based on the associated resource property.</p> <p>Valid values:</p> <ul style="list-style-type: none"> • ALIYUN::ECS::Instance::ImageId • ALIYUN::ECS::Instance::ZoneId • ALIYUN::ECS::VPC::VPCId • ALIYUN::ECS::VSwitch::VSwitchId <p>For example, if you set <code>AssociationProperty</code> to <code>ALIYUN::ECS::Instance::ImageId</code>, ROS verifies whether the specified image ID is valid and lists other valid values in a drop-down list.</p> |
| Confirm | No | <p>Specifies whether to enter the parameter value for a second time if the NoEcho property is set to <code>true</code>. Default value: <code>false</code>.</p> <div style="background-color: #e0f2f7; padding: 5px;"> <p> Notice The Confirm property can be set to <code>true</code> only when it is used with a String-type parameter and when the NoEcho parameter is set to <code>true</code>.</p> </div> |

Examples

In the following example, two parameters are defined in the Parameters section:

- `username`
 - Type: String
 - Valid values:
 - `anonymous`
 - `user-one`
 - `user-two`
 - Length: 6 to 12 characters

 **Notice** The default value `anonymous` must also meet the length and valid value requirements.

- `password`
 - Type: String
 - Length: 1 to 41 characters
 - The password can contain letters and digits.
 - If you set the NoEcho property to `true`, the GetStack operation does not return parameter values.

```
"Parameters": {
  "username": {
    "Label": "Username"
    "Description": "Enter the username"
    "Default": "anonymous",
    "Type": "String",
    "MinLength": "6",
    "MaxLength": "12",
    "AllowedValues": ["anonymous", "user-one", "user-two"]
  },
  "password": {
    "Label": "Password"
    "NoEcho": "True",
    "Description": "Enter the password"
    "Type": "String",
    "MinLength": "1",
    "MaxLength": "41",
    "AllowedPattern": "[a-zA-Z0-9]*"
  }
}
```

Pseudo parameters

Pseudo parameters are internal parameters provided by the ROS engine. They can be referenced in the same manner as user-defined parameters, and their values are determined when ROS is running. The following pseudo parameters are supported:

- `ALIYUN::StackName` : the name of the stack.
- `ALIYUN::StackId` : the ID of the stack.
- `ALIYUN::Region` : the region where the stack resides.
- `ALIYUN::AccountId` : the account ID of the stack.
- `ALIYUN::NoValue` : specifies whether the specific resource property is deleted when the resource is created or updated.

5.4.3. Resources

This topic describes the properties of each resource and dependencies between resources in a stack. A resource can be referenced by other resources and output items.

Syntax

Each resource consists of an ID and a description. All resource descriptions are enclosed in braces {}. Multiple resources are separated by commas (.). The following sample code shows the Resources syntax:

```
"Resources": {
  "Resource1 ID": {
    "Type": "The resource type",
    "Condition": "The condition that specifies whether to create the resource",
    "Properties": {
      The description of the resource properties
    }
  },
  "Resource2 ID": {
    "Type": "The resource type",
    "Condition": "The condition that specifies whether to create the resource",
    "Properties": {
      The description of the resource properties
    }
  }
}
```

Parameter description:

- The resource ID must be unique within the template. You can use the resource ID to reference the resource in other parts of the template.
- The Type parameter specifies the type of the resource that is being declared. For example, ALIYUN::ECS::Instance indicates that the resource is an Elastic Cloud Service (ECS) instance.
- The Properties section provides additional options that you can specify for a resource. For example, you must specify an image ID for each ECS instance. The image ID is one of the resource properties.

Examples

```
"Resources": {
  "ECSInstance": {
    "Type": "ALIYUN::ECS::Instance",
    "Properties": {
      "ImageId": "m-25l0r****"
    }
  }
}
```

If a resource does not need properties to be declared, omit the Properties section of that resource.

Property values can be text strings, string lists, Boolean values, referenced parameters, or return values of functions.

The following example shows how to declare different types of property values:

```
"Properties": {
  "String": "string",
  "LiteralList": [ "value1", "value2" ],
  "Boolean": "true"
  "ReferenceForOneValue": { "Ref": "ResourceID" },
  "FunctionResultWithFunctionParams": {
    "Fn::Join": [ "%", [ "Key=", { "Ref": "SomeParameter" } ] ]
  }
}
```

DeletionPolicy

The DeletionPolicy parameter specifies whether to retain a resource when its stack is deleted. The following sample code shows how to use the DeletionPolicy parameter to retain an ECS instance when its stack is deleted:

```
"Resources": {
  "ECSInstance": {
    "Type": "ALIYUN::ECS::Instance",
    "Properties": {
      "ImageId": "m-25l0r****"
    },
    "DeletionPolicy": "Retain"
  }
}
```

DependsOn

The DependsOn parameter allows you to create a specific resource after you create its dependent resource. If you specify the DependsOn parameter for a resource, the resource is created only after its dependent resource specified by the DependsOn parameter is created.

In the following example, WebServer is created only after DatabaseServer is created:

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "WebServer": {
      "Type": "ALIYUN::ECS::Instance",
      "DependsOn": "DatabaseServer"
    },
    "DatabaseServer": {
      "Type": "ALIYUN::ECS::Instance",
      "Properties": {
        "ImageId": "m-25l0r****",
        "InstanceType": "ecs.t1.small"
      }
    }
  }
}
```

Condition

The Condition parameter specifies whether to create the resource. The resource can be created only when the Condition parameter is set to true.

In the following example, WebServer is created only if the condition determined by the MaxAmount parameter is true:

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "MaxAmount": {
      "Type": "Number",
      "Default": 1
    }
  },
  "Conditions": {
    "CreateWebServer": {"Fn::Not": {"Fn::Equals": [0, {"Ref": "MaxAmount"}]}}
  }
  "Resources": {
    "WebServer": {
      "Type": "ALIYUN::ECS::InstanceGroup",
      "Condition": "CreateWebServer",
      "Properties": {
        "ImageId": "m-25l0r****",
        "InstanceType": "ecs.t1.small",
        "MaxAmount": {"Ref": "MaxAmount"}
      }
    },
    "DatabaseServer": {
      "Type": "ALIYUN::ECS::Instance",
      "Properties": {
        "ImageId": "m-25l0r****",
        "InstanceType": "ecs.t1.small"
      }
    }
  }
}
```

Resource declaration example

The following example shows how to declare a resource:

```

"Resources": {
  "WebServer": {
    "Type": "ALIYUN::ECS::Instance",
    "Properties": {
      "ImageId": "m-25l0r****",
      "InstanceType": "ecs.t1.small",
      "SecurityGroupId": "sg-25zwc****",
      "ZonId": "cn-beijing-b",
      "Tags": [
        {
          "Key": "Department1",
          "Value": "HumanResource"
        },
        {
          "Key": "Department2",
          "Value": "Finance"
        }
      ]
    }
  },
  "ScalingConfiguration": {
    "Type": "ALIYUN::ESS::ScalingConfiguration",
    "Properties": {
      "ImageId": "ubuntu1404_64_20G_aliaegis_2015****.vhd",
      "InstanceType": "ecs.t1.small",
      "InstanceId": "i-25xhh****",
      "InternetChargeType": "PayByTraffic",
      "InternetMaxBandwidthIn": 1,
      "InternetMaxBandwidthOut": 20,
      "SystemDisk_Category": "cloud",
      "ScalingGroupId": "bwhtvpcBcKYac9fe3vd0****",
      "SecurityGroupId": "sg-25zwc****",
      "DiskMappings": [
        {
          "Size": 10
        },
        {
          "Category": "cloud",
          "Size": 10
        }
      ]
    }
  }
}

```

5.4.4. Outputs

The Outputs section is used to define the values returned when the GetStack operation is called. For example, if you define an Elastic Compute Service (ECS) instance ID as an output item, the ECS instance ID is returned when the GetStack operation is called.

Syntax

Each output item consists of an ID and a description. All output descriptions are enclosed in braces {}. Multiple output items are separated by commas (.). Each output item can have multiple values in an array format. The following example shows the Outputs syntax:

```

"Outputs" : {
  "Output1 ID" : {
    "Description": "The description of the output item",
    "Condition": "The condition that specifies whether to provide resource properties",
    "Value": "The output value expression"
  },
  "Output2 ID" : {
    "Description": "The description of the output item",
    "Condition": "The condition that specifies whether to provide resource properties",
    "Value": [
      "Output value expression 1",
      "Output value expression 2",
      ...
    ]
  }
}

```

- Output ID: the ID of the output item. Duplicate IDs are not allowed within a template.
- Description: optional. The description of the output item.
- Value: required. The value returned when the GetStack operation is called.
- Condition: optional. The condition that specifies whether to create a resource and provide its information. The resource is created and its information is provided only when the specified condition is `true`.

In the following example, WebServer is created only if the condition determined by the MaxAmount parameter is true:

```

{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "MaxAmount": {
      "Type": "Number",
      "Default": 1
    }
  },
  "Conditions": {
    "CreateWebServer": {"Fn::Not": {"Fn::Equals": [0, {"Ref": "MaxAmount"}]}}
  }
  "Resources": {
    "WebServer": {
      "Type": "ALIYUN::ECS::InstanceGroup",
      "Condition": "CreateWebServer",
      "Properties": {
        "ImageId": "m-25l0r****",
        "InstanceType": "ecs.t1.small"
        "MaxAmount": {"Ref": "MaxAmount"}
      }
    }
  }
  "Outputs": {
    "WebServerIP": {
      "Condition": "CreateWebServer",
      "Value": {
        "Fn::GetAtt": ["WebServer", "PublicIps"]
      }
    }
  }
}

```

Examples

The following example contains two output items:

- The InstanceId value of WebServer
- The PublicIp and PrivateIp values of WebServer

```
"Outputs": {
  "InstanceId": {
    "Value": {"Fn::GetAtt": ["WebServer", "InstanceId"]}
  },
  "PublicIp & PrivateIp": {
    "Value": [
      {"Fn::GetAtt": ["WebServer", "PublicIp"]},
      {"Fn::GetAtt": ["WebServer", "PrivateIp"]}
    ]
  }
}
```

5.4.5. Functions

Resource Orchestration Service (ROS) provides several built-in functions to help you manage stacks. You can use built-in functions to define Resources and Outputs.

Fn::Base64Encode

The Fn::Base64Encode function is used to return the Base64 representation of the input string.

- Declaration

```
"Fn::Base64Encode": "stringToEncode"
```

- Parameters

`stringToEncode` : the string to be encoded in Base64.

- Return value

The Base64 representation of the input string.

- Examples

```
{"Fn::Base64Encode": "string to encode"}
```

`c3RyaW5nIHRvIGVudWY29kZQ==` is returned in this example.

Fn::Base64Decode

The Fn::Base64Decode function is used to return a string decoded from a Base64-encoded string.

- Declaration

```
{"Fn::Base64Decode": "stringToEncode"}
```

- Parameters

`stringToDecode` : the string decoded from the Base64-encoded string.

- Return value

The string decoded from the Base64-encoded string.

- Examples

```
{"Fn::Base64Decode": "c3RyaW5nIHRvIGVudWY29kZQ=="}
```

`string to encode` is returned in this example.

Fn::Base64

The Fn::Base64 function returns the Base64 representation of the input string.

- Declaration

```
"Fn::Base64": stringToEncode
```

- Parameters

`valueToEncode`: the string to be encoded in Base64.

- Return value

The Base64 representation of the input string.

- Examples

```
"Fn::Base64": "string to encode"
```

Fn::FindInMap

The Fn::FindInMap function is used to return the values based on keys in a two-level mapping that is declared in the Mappings section.

- Declaration

```
"Fn::FindInMap": ["MapName", "TopLevelKey", "SecondLevelKey"]
```

- Parameters

- `MapName` : the ID of a mapping declared in the Mappings section that contains keys and values.
- `TopLevelKey` : the top-level key name. The value is a list of key-value pairs.
- `SecondLevelKey` : the second-level key name. The value is a string or a number.

- Return value

The value that is assigned to the `SecondLevelKey` parameter.

- Examples

The `ImageId` property must be specified when you create a `WebServer` instance. The Mappings section describes the `ImageId` mappings by region. The Parameters section describes the regions that must be specified by template users. `Fn::FindInMap` finds the corresponding `ImageId` mapping in `RegionMap` based on the region specified by a user, and then finds the corresponding `ImageId` values in the mapping.

- `MapName` can be set to a custom value, which is `RegionMap` in this example.
- `TopLevelKey` is set to the region where the stack is created, which is `{"Ref": "regionParam"}` in this example.
- `SecondLevelKey` is set to the required architecture, which is `32` in this example.

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "regionParam": {
      "Description": "The region where the ECS instance is created",
      "Type": "String",
      "AllowedValues": [
        "hangzhou",
        "beijing"
      ]
    }
  },
  "Mappings": {
    "RegionMap": {
      "hangzhou": {
        "32": "m-25l0rcfjo",
        "64": "m-25l0rcfj1"
      },
      "beijing": {
        "32": "m-25l0rcfj2",
        "64": "m-25l0rcfj3"
      }
    }
  },
  "Resources": {
    "WebServer": {
      "Type": "ALIYUN::ECS::Instance",
      "Properties": {
        "ImageId": {
          "Fn::FindInMap": [
            "RegionMap",
            {"Ref": "regionParam"},
            "32"
          ]
        },
        "InstanceType": "ecs.t1.small",
        "SecurityGroupId": "sg-25zwc****",
        "ZoneId": "cn-beijing-b",
        "Tags": [
          {
            "Key": "key1",
            "Value": "value1"
          },
          {
            "Key": "key2",
            "Value": "value2"
          }
        ]
      }
    }
  }
}
```

- Supported functions
 - Fn::FindInMap
 - Ref

Fn::GetAtt

The Fn::GetAtt function is used to return the value of a property from a resource in a template.

- Declaration

```
"Fn::GetAtt": ["resourceID", "attributeName"]
```

- Parameters

- resourceID : the ID of the resource.
- attributeName : the name of the resource property.

- Return value

The value of the resource property.

- Examples

The ImageId property of MyEcsInstance is returned in this example.

```
{"Fn::GetAtt": ["MyEcsInstance", "ImageID"]}
```

Fn::Join

The Fn::Join function is used to combine a set of values into a single value that is separated by a specified delimiter.

- Declaration

```
{"Fn::Join": ["delimiter", ["string1", "string2", ... ]]}
```

- Parameters

- delimiter : the value used to divide the string. The delimiter value can be left blank so that all the values are directly combined.
- ["string1", "string2", ...] : the list of values that are combined into a string.

- Return value

The combined string.

- Examples

```
{"Fn::Join": [ "", ["a", "b", "c"]]}
```

"a,b,c" is returned in this example.

- Supported functions

- Fn::Base64Encode
- Fn::GetAtt
- Fn::Join
- Fn::Select
- Ref

Fn::Select

The Fn::Select function is used to return a single data element from a list of data elements by using an index.

- Declaration

- The following example assumes that the list of data elements is an array:

```
"Fn::Select": ["index", ["value1", "value2", ... ]]
```

- The following example assumes that the list of data elements is a mapping table:

```
"Fn::Select": ["index", {"key1": "value1", ... }]
```

- Parameters

`index` : the index of the object data element. If the list of data elements is an array, the index must be an integer ranging from 0 to N-1, where N indicates the number of elements in the array. If the list of data elements is a mapping table, the index must be a key in the mapping table.

If the corresponding value of the index cannot be found, the system returns an empty string.

- Return value

The object data element.

- Examples

- The following example assumes that the list of data elements is an array:

```
{"Fn::Select": ["1", ["apples", "grapes", "oranges", "mangoes"]]}
```

`"grapes"` is returned in this example.

- The following example assumes that the list of data elements is a mapping table:

```
{"Fn::Select": ["key1", {"key1": "grapes", "key2": "mangoes"}]}
```

`"grapes"` is returned in this example.

- The following example assumes that the list of data elements is a comma-delimited list:

```
"Parameters": {
  "userParam": {
    "Type": "CommaDelimitedList",
    "Default": "10.0.100.0/24, 10.0.101.0/24, 10.0.102.0/24"
  }
}
"Resources": {
  "resourceID": {
    "Properties": {
      "CidrBlock": {"Fn::Select": ["0", {"Ref": "userParam"}]}
    }
  }
}
```

- Supported functions

- For the `Fn::Select` index value, you can use the `Ref` function.
- For the `Fn::Select` list of data elements, you can use the following functions:
 - `Fn::Base64Encode`
 - `Fn::FindInMap`
 - `Fn::GetAtt`
 - `Fn::Join`
 - `Fn::Select`
 - `Ref`

Ref

The `Ref` function is used to return the value of a specified parameter or resource.

If the specified parameter is a resource ID, the value of the resource is returned. Otherwise, the system returns the value of the specified parameter.

- Declaration

```
"Ref": "logicalName"
```

- Parameters

`logicalName` : the logical name of the resource or parameter that you want to reference.

- Return value

The value of the resource or parameter.

- Examples

The following example demonstrates how to use the Ref function to specify regionParam as the region parameter for RegionMap of WebServer:

```

{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "regionParam": {
      "Description": "The region where the ECS instance is created",
      "Type": "String",
      "AllowedValues": [
        "hangzhou",
        "beijing"
      ]
    }
  },
  "Mappings": {
    "RegionMap": {
      "hangzhou": {
        "32": "m-25l0rcfjo",
        "64": "m-25l0rcfj1"
      },
      "beijing": {
        "32": "m-25l0rcfj2",
        "64": "m-25l0rcfj3"
      }
    }
  },
  "Resources": {
    "WebServer": {
      "Type": "ALIYUN::ECS::Instance",
      "Properties": {
        "ImageId": {
          "Fn::FindInMap": [
            "RegionMap",
            {"Ref": "regionParam"},
            "32"
          ]
        },
        "InstanceType": "ecs.t1.small",
        "SecurityGroupId": "sg-25zwc****",
        "ZoneId": "cn-beijing-b",
        "Tags": [
          {
            "Key": "tiantt",
            "Value": "ros"
          },
          {
            "Key": "tiantt1",
            "Value": "ros1"
          }
        ]
      }
    }
  }
}

```

- Supported function

When you use the Ref function, you cannot use other functions in it at the same time. You must specify a string value for the logical ID of a resource.

Fn::GetAZs

The Fn::GetAZs function is used to return a list of zones for a specified region.

- Declaration

```
"Fn::GetAZs": "region"
```

- Parameters

```
region : the ID of the region.
```

- Return value

The list of zones within the specified region.

- Examples

The following example demonstrates how to create an Elastic Compute Service (ECS) instance in the first zone of a specified region:

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "WebServer": {
      "Type": "ALIYUN::ECS::Instance",
      "Properties": {
        "ImageId": "centos7u2_64_40G_cloudinit_2016****.raw",
        "InstanceType": "ecs.n1.tiny",
        "SecurityGroupId": "sg-2zedcm7ep5quses0****",
        "Password": "Ros1****",
        "AllocatePublicIP": true,
        "InternetChargeType": "PayByTraffic",
        "InternetMaxBandwidthIn": 100,
        "InternetMaxBandwidthOut": 100,
        "SystemDiskCategory": "cloud_efficiency",
        "IoOptimized": "optimized",
        "ZoneId": {"Fn::Select": ["0", {"Fn::GetAZs": {"Ref": "ALIYUN::Region"}}]}
      }
    }
  },
  "Outputs": {
    "InstanceId": {
      "Value": {"Fn::GetAtt": ["WebServer", "InstanceId"]}
    },
    "PublicIp": {
      "Value": {"Fn::GetAtt": ["WebServer", "PublicIp"]}
    }
  }
}
```

- Supported functions

- Fn::Base64Encode
- Fn::FindInMap
- Fn::GetAtt
- Fn::Join
- Fn::Select
- Ref

Fn::Replace

The Fn::Replace function is used to replace a specified substring contained in a string with a new substring.

- Declaration

```
{"Fn::Replace": [{"object_key": "object_value"}, "object_string"]}
```

- Parameters
 - `object_key` : the substring to be replaced.
 - `object_value` : the new substring to replace the previous substring.
 - `object_string` : the string whose `object_key` is replaced.

- Return value

The string after replacement.

- Examples

The following example demonstrates how to replace "print" with "echo" in the specified script:

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "WebServer": {
      "Type": "ALIYUN::ECS::Instance",
      "Properties": {
        "ImageId": "centos_7_2_64_40G_base_2017****.vhd",
        "InstanceType": "ecs.n1.medium",
        "SecurityGroupId": "sg-94q49****",
        "Password": "MytestPassword****",
        "IoOptimized": "optimized",
        "VSwitchId": "vsw-94vdv8****",
        "VpcId": "vpc-949uz****",
        "SystemDiskCategory": "cloud_ssd",
        "UserData": {"Fn::Replace": [{"print": "echo"},
          {"Fn::Join": ["", [
            "#!/bin/sh\n",
            "mkdir ~/test_ros\n",
            "print hello > ~/1.txt\n"
          ]]}]}
      }
    }
  },
  "Outputs": {
    "InstanceId": {
      "Value": {"Fn::GetAtt": ["WebServer", "InstanceId"]}
    },
    "PublicIp": {
      "Value": {"Fn::GetAtt": ["WebServer", "PublicIp"]}
    }
  }
}
```

- Supported functions
 - Fn::Base64Encode
 - Fn::GetAtt
 - Fn::Join
 - Fn::Select
 - Ref

Fn::Split

The Fn::Split function is used to split a string into a list of values separated by a specified delimiter and return the list.

- Declaration

```
"Fn::Split": ["delim", "original_string"]
```

- Parameters

- `delim` : the specified delimiter, which can be a comma (,), semicolon (;), line break (\n), or indent (\t).
- `original_string` : the string to be split.

- Return value

A list of string values.

- Examples

- The following example assumes that the list of data elements is an array:

```
{"Fn::Split": [";", "foo; bar; achoo"]}
```

`["foo", " bar", " achoo "]` is returned in this example.

- The following example demonstrates how to use `Fn::Split` to split `InstanceIds`:

```
{
  "Parameters": {
    "InstanceIds": {
      "Type": "String",
      "Default": "instane1_id,instance2_id,instance2_id"
    }
  },
  "Resources": {
    "resourceID": {
      "Type": "ALIYUN::SLB::BackendServerAttachment",
      "Properties": {
        "BackendServerList": {
          "Fn::Split": [
            ";",
            {
              "Ref": "InstanceIds"
            }
          ]
        }
      }
    }
  }
}
```

- Supported functions

- `Fn::Base64Encode`
- `Fn::FindInMap`
- `Fn::GetAtt`
- `Fn::Join`
- `Fn::Select`
- `Fn::Replace`
- `Fn::GetAZs`
- `Fn::If`
- `Ref`

Fn::Equals

The Fn::Equals function is used to compare whether two values are equal. If the two values are equal, true is returned. If the two values are not equal, false is returned.

- Declaration

```
{"Fn::Equals":["value_1","value_2"]}
```

- Parameters

`value` : the values to be compared.

- Return value

true or false.

- Examples

The following example demonstrates how to use Fn::Equals to define a condition in the Conditions section:

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "EnvType": {
      "Default": "pre",
      "Type": "String"
    }
  },
  "Conditions": {
    "TestEqualsCond": {
      "Fn::Equals": [
        "prod",
        {"Ref": "EnvType"}
      ]
    }
  }
}
```

- Supported functions

- Fn::Or
- Fn::Not
- Fn::Equals
- Fn::FindInMap
- Fn::And
- Ref

Fn::And

The Fn::And function is used to represent the AND operator, and must contain at least two conditions. If all the specified conditions are evaluated as true, true is returned. If any condition is evaluated as false, false is returned.

- Declaration

```
{"Fn::And":["condition", {...]}
```

- Parameters

`condition` : the condition to be evaluated.

- Return value

true or false.

- Examples

The following example demonstrates how to use Fn::And to define a condition in the Conditions section:

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters":{
    "EnvType":{
      "Default":"pre",
      "Type":"String"
    }
  },
  "Conditions": {
    "TestEqualsCond": {"Fn::Equals": ["prod", {"Ref": "EnvType"}]},
    "TestAndCond": {"Fn::And": [{"TestEqualsCond"}, {"Fn::Equals": ["pre", {"Ref": "EnvType"}]}]}
  }
}
```

- Supported functions
 - Fn::Or
 - Fn::Not
 - Fn::Equals
 - Fn::FindInMap
 - Fn::And
 - Ref

Fn::Or

The Fn::Or function is used to represent the OR operator, and must contain at least two conditions. If any specified condition is evaluated as true, true is returned. If all the conditions are evaluated as false, false is returned.

- Declaration

```
{"Fn::Or": ["condition", {...]}
```

- Parameters

condition : the condition to be evaluated.

- Return value

true or false.

- Examples

The following example demonstrates how to use Fn::Or to define a condition in the Conditions section:

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters":{
    "EnvType":{
      "Default":"pre",
      "Type":"String"
    }
  },
  "Conditions": {
    "TestEqualsCond": {"Fn::Equals": ["prod", {"Ref": "EnvType"}]},
    "TestOrCond": {"Fn::And": [{"TestEqualsCond"}, {"Fn::Equals": ["pre", {"Ref": "EnvType"}]}]}
  }
}
```

- Supported functions
 - Fn::Or
 - Fn::Not

- Fn::Equals
- Fn::FindInMap
- Fn::And
- Ref

Fn::Not

The Fn::Not function is used to represent the NOT operator. If a condition is evaluated as false, true is returned. If a condition is evaluated as true, false is returned.

- Declaration

```
{"Fn::Not": "condition"}
```

- Parameters

`condition` : the condition to be evaluated.

- Return value

true or false.

- Examples

The following example demonstrates how to use Fn::Not to define a condition in the Conditions section:

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "EnvType": {
      "Default": "pre",
      "Type": "String"
    }
  },
  "Conditions": {
    "TestNotCond": {"Fn::Not": {"Fn::Equals": ["pre", {"Ref": "EnvType"}]}}
  }
}
```

- Supported functions

- Fn::Or
- Fn::Not
- Fn::Equals
- Fn::FindInMap
- Fn::And
- Ref

Fn::If

This function returns one of two possible values. If a specified condition is evaluated as true, one value is returned. If the specified condition is evaluated as false, the other value is returned. The property values of Resources and Outputs in templates support the Fn::If function. You can use the `ALIYUN::NoValue` pseudo parameter as the return value to delete the corresponding property.

- Declaration

```
{"Fn::If": ["condition_name", "value_if_true", "value_if_false"]}
```

- Parameters

- `condition_name` : the name of the condition in the Conditions section. A condition is referenced by using the condition name.

- `value_if_true` : If the specified condition is evaluated as true, this value is returned.
- `value_if_false` : If the specified condition is evaluated as false, this value is returned.

- Examples

The following example demonstrates how to determine whether to create a data disk based on input parameters:

```
{
  "ROSTemplateFormatVersion":"2015-09-01",
  "Parameters":{
    "EnvType":{
      "Default":"pre",
      "Type":"String"
    }
  },
  "Conditions":{
    "CreateDisk":{
      "Fn::Equals":[
        "prod",
        {
          "Ref":"EnvType"
        }
      ]
    }
  },
  "Resources":{
    "WebServer":{
      "Type":"ALIYUN::ECS::Instance",
      "Properties":{
        "DiskMappings":{
          "Fn::If":[
            "CreateDisk",
            [
              {
                "Category":"cloud_efficiency",
                "DiskName":"FirstDataDiskName",
                "Size":40
              },
              {
                "Category":"cloud_ssd",
                "DiskName":"SecondDataDiskName",
                "Size":40
              }
            ],
            {
              "Ref":"ALIYUN::NoValue"
            }
          ]
        },
        "VpcId":"vpc-2zew9pxh2yirtzqxd****",
        "SystemDiskCategory":"cloud_efficiency",
        "SecurityGroupId":"sg-2zece6wcqriejf1v****",
        "SystemDiskSize":40,
        "ImageId":"centos_6_8_64_40G_base_2017****.vhd",
        "IoOptimized":"optimized",
        "VSwitchId":"vsw-2zed9txvy7h2srqo6****",
        "InstanceType":"ecs.n1.medium"
      }
    }
  }
}
```

```

},
"Outputs":{
  "Instanceid":{
    "Value":{
      "Fn::GetAtt":[
        "WebServer",
        "Instanceid"
      ]
    }
  },
  "Zoneid":{
    "Value":{
      "Fn::GetAtt":[
        "WebServer",
        "Zoneid"
      ]
    }
  }
}
}
}

```

- Supported functions
 - Fn::Or
 - Fn::Not
 - Fn::Equals
 - Fn::FindInMap
 - Fn::And
 - Ref

Fn::ListMerge

The Fn::ListMerge function is used to merge multiple lists into one list.

- Declaration

```
{"Fn::ListMerge": [{"list_1_item_1", "list_1_item_2", ...}, {"list_2_item_1", "list_2_item_2", ...}]}
```

- Parameters

- ["list_1_item_1", "list_1_item_2", ...] : the first list to merge.
- ["list_2_item_1", "list_2_item_2", ...] : the second list to merge into the first list.

- Examples

The following example demonstrates how to attach two ECS instance groups to a Server Load Balancer (SLB) instance:

```

{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "LoadBalancer": {
      "Type": "ALIYUN::SLB::LoadBalancer",
      "Properties": {
        "LoadBalancerName": "ros",
        "AddressType": "internet",
        "InternetChargeType": "paybybandwidth",
      }
    },
    "BackendServer1": {
      "Type": "ALIYUN::ECS::InstanceGroup",
      "Properties": {
        "ImageId": "m-2ze9uqi7wo61hwep****",
        "InstanceType": "ecs.t1.small",
        "SecurityGroupId": "sg-2ze8yxgempcdsq3****",
        "MaxAmount": 1,
        "MinAmount": 1
      }
    },
    "BackendServer2": {
      "Type": "ALIYUN::ECS::InstanceGroup",
      "Properties": {
        "ImageId": "m-2ze9uqi7wo61hwep****",
        "InstanceType": "ecs.t1.small",
        "SecurityGroupId": "sg-2ze8yxgempcdsq3iu****",
        "MaxAmount": 1,
        "MinAmount": 1
      }
    },
    "Attachment": {
      "Type": "ALIYUN::SLB::BackendServerAttachment",
      "Properties": {
        "LoadBalancerId": {"Ref": "LoadBalancer"},
        "BackendServerList": { "Fn::ListMerge": [
          {"Fn::GetAtt": ["BackendServer1", "InstanceIds"]},
          {"Fn::GetAtt": ["BackendServer2", "InstanceIds"]}
        ]
      }
    }
  }
}

```

- Supported functions
 - Fn::Base64Encode
 - Fn::GetAtt
 - Fn::Join
 - Fn::Select
 - Ref
 - Fn::If

Fn::GetJsonValue

The Fn::GetJsonValue function is used to resolve a JSON string and obtain its key value from the first layer.

- Declaration

```
{"Fn::GetJsonValue":["key","json_string"]}
```

- Parameters
 - `key` : the key value.
 - `json_string` : the specified JSON string to be resolved.
- Examples

In the following example, the WebServer instance executes UserData and returns a JSON string, and the WebServer2 instance then obtains the corresponding key value from the string.

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "WebServer": {
      "Type": "ALIYUN::ECS::Instance",
      "Properties": {
        "ImageId": "m-2ze45uwova5fedlu****",
        "InstanceType": "ecs.n1.medium",
        "SecurityGroupId": "sg-2ze7pxymaix640qr****",
        "Password": "Wenqiao****",
        "IoOptimized": "optimized",
        "VSwitchId": "vsw-2zei67xd9nhcqxe****",
        "VpcId": "vpc-2zevx9ios1rszqv0a****",
        "SystemDiskCategory": "cloud_ssd",
        "UserData": {"Fn::Join": ["", [
          "#!/bin/sh\n",
          "mkdir ~/test_ros\n",
          "print hello > ~/1.txt\n",
          "Fn::GetAtt": ["WaitConHandle", "CurlCli"],
          "\n",
          "Fn::GetAtt": ["WaitConHandle", "CurlCli"],
          "-d '{\"id\": \"1\", \"data\": [\"111\", \"222\"]}'\n"
        ]]},
        "PrivateIpAddress": "192.168.XX.XX",
        "HostName": "userdata-1
      }
    },
    "WaitConHandle": {
      "Type": "ALIYUN::ROS::WaitConditionHandle"
    },
    "WaitCondition": {
      "Type": "ALIYUN::ROS::WaitCondition",
      "Properties": {
        "Handle": {"Ref": "WaitConHandle"},
        "Timeout": 900
      }
    },
    "WebServer2": {
      "Type": "ALIYUN::ECS::Instance",
      "Properties": {
        "ImageId": "m-2ze45uwova5fedlu****",
        "InstanceType": "ecs.n1.medium",
        "SecurityGroupId": "sg-2ze7pxymaix640qr****",
        "Password": "Wenqiao****",
        "IoOptimized": "optimized",
        "VSwitchId": "vsw-2zei67xd9nhcqzec****",
        "VpcId": "vpc-2zevx9ios1rszqv0a****",
        "SystemDiskCategory": "cloud_ssd",
        "UserData":
```

```

    {"Fn::Join": ["", [
      "#!/bin/sh\n",
      "mkdir ~/test_ros\n",
      "echo hello > ~/1.txt\n",
      "server_1_token=",
      {"Fn::GetJsonValue": ["1", {"Fn::GetAtt": ["WaitCondition", "Data"]}],
      "\n"
    ]}],
    "PrivateIpAddress": "192.168.XX.XX",
    "HostName": "userdata-2"
  }
},
},
"Outputs": {
  "Instanceid": {
    "Value": {"Fn::GetAtt": ["WebServer", "Instanceid"]}
  },
  "PublicIp": {
    "Value": {"Fn::GetAtt": ["WebServer", "PublicIp"]}
  }
}
}
}

```

- Supported functions
 - Fn::Base64Encode
 - Fn::GetAtt
 - Fn::Join
 - Fn::Select
 - Ref
 - Fn::If

Fn::MergeMapToList

The Fn::MergeMapToList function is used to merge multiple mappings into a list of mapping elements.

- Declaration

```

{"Fn::MergeMapToList": [{"key_1": ["key_1_item_1", "key_1_item_2", ...]}, {"key_2": ["key_2_item_1", "key_2_item_2", ...]}, ...]}

```

- Parameters

- {"key_1": ["key_1_item_1", "key_1_item_2", ...]} : the first mapping to merge. The "key_1" value must be a list. "key_1" is the key for each mapping in the list of merged mappings. The "key_1" value is "key_1_item_1" for the first merged mapping and "key_1_item_2" for the second merged mapping. All values follow the same format. The length of the final list of merged mappings is the length of the longest list "key_x" from all mappings being merged. If a "key_y" list is shorter, the last element of the list is repeated until the list is the longest.
- {"key_2": ["key_2_item_1", "key_2_item_2", ...]} : the second mapping to merge into the first mapping. The "key_2" value must be a list. "key_2" is the key for each mapping in the merged list. The "key_2" value is "key_2_item_1" for the first merged mapping and "key_2_item_2" for the second merged mapping. All values follow the same format.

- Examples

- The following example demonstrates how to merge three mappings. The length of the list based on the key values for each mapping is the same.

```
{
  "Fn::MergeMapToList": [
    {"key_1": ["kye_1_item_1", "kye_1_item_2"]},
    {"key_2": ["kye_2_item_1", "kye_2_item_2"]},
    {"key_3": ["kye_3_item_1", "kye_3_item_2"]}
  ]
}
```

The following code shows the merged result:

```
[
  {
    "key_1": "kye_1_item_1",
    "key_2": "kye_2_item_1",
    "key_3": "kye_3_item_1"
  },
  {
    "key_1": "kye_1_item_2",
    "key_2": "kye_2_item_2",
    "key_3": "kye_3_item_2"
  }
]
```

- The length of the list based on the key values for each mapping varies in the following example:

```
{
  "Fn::MergeMapToList": [
    {"key_1": ["kye_1_item_1", "kye_1_item_2"]},
    {"key_2": ["kye_2_item_1", "kye_2_item_2", "key_2_item_3"]},
    {"key_3": ["kye_3_item_1", "kye_3_item_2"]}
  ]
}
```

The following code shows the merged result:

```
[
  {
    "key_1": "kye_1_item_1",
    "key_2": "kye_2_item_1",
    "key_3": "kye_3_item_1"
  },
  {
    "key_1": "kye_1_item_2",
    "key_2": "kye_2_item_2",
    "key_3": "kye_3_item_2"
  },
  {
    "key_1": "kye_1_item_2",
    "key_2": "kye_2_item_3",
    "key_3": "kye_3_item_2"
  }
]
```

- o In the following template example, all instances created in WebServer are added to the vServer group of an SLB instance:

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "WebServer": {
      "Type": "ALIYUN::ECS::InstanceGroupClone",
      "Properties": {
        "SourceInstanceId": "i-xxxxx",
        "Password": "Hello****",
        "MinAmount": 1,
        "MaxAmount": 1
      }
    },
    "CreateVServerGroup": {
      "Type": "ALIYUN::SLB::VServerGroup",
      "Properties": {
        "LoadBalancerId": "lb-****",
        "VServerGroupName": "VServerGroup-****",
        "BackendServers": {
          "Fn::MergeMapToList": [
            {"Port": [6666, 9090, 8080]},
            {"ServerId": {"Fn::GetAtt": ["WebServer", "InstanceIds"]},
             {"Weight": [20, 100]}
          ]
        }
      }
    }
  }
}
```

- Supported functions
 - o Fn::Base64Encode
 - o Fn::GetAtt
 - o Fn::Join
 - o Fn::Select
 - o Ref
 - o Fn::If
 - o Fn::ListMerge
 - o Fn::GetJsonValue

Fn::Avg

The Fn::Avg function is used to return the average value of a set of numbers.

- Declaration

```
{"Fn::Avg": [ndigits, [number1, number2, ... ]]}
```

- Parameters
 - o **ndigits** : the number of decimal places to display. This parameter value must be an integer.
 - o **[number1, number2, ...]** : the set of numbers for which the average value will be calculated. Each element in the group must be a number or a string that can be converted into a number.
- Return value

The average value of the set of numbers.

- Examples

```
{ "Fn::Avg": [ 1, [1, 2, 6.0]] }
{ "Fn::Avg": [ 1, ['1', '2', '6.0']] }
```

3.0 is returned in this example.

- Supported functions

- Fn::GetAtt
- Ref

Fn::SelectMapList

The Fn::SelectMapList function is used to return a list of map elements.

- Declaration

```
{ "Fn::SelectMapList": ["key2", [{"key1": "value1-1", "key3": "value1-3"}, {"key1": "value2-1", "key2": "value2-2"}, {"key1": "value3-1", "key2": "value3-2"}, ...] ] }
```

- Parameters

- key2 : the key to be queried in the map.
- [{"key1": "value1-1", "key3": "value1-3"}, ...] : the list of maps.

- Return value

A list of key values for all maps in the map list.

- Examples

```
{
  "Fn::SelectMapList": [
    "key2",
    [
      {"key1": "value1-1", "key3": "value1-3"},
      {"key1": "value2-1", "key2": "value2-2"},
      {"key1": "value3-1", "key2": "value3-2"}
    ]
  ]
}
```

["value2-2", "value3-2"] is returned in this example.

Fn::Add

The Fn::Add function is used to sum the values of parameters.

- Declaration

```
{ "Fn::Add": [{"Product": "ROS"}, {"Fn": "Add"}] }
```

- Parameters

- The parameters must be arranged as a list.
- The parameters in the list can be of the Number, List, or Dictionary type. All the parameters must be of the same type. The list must contain at least two parameters.

- Return value

If the parameter values are numbers, sum the parameter values. If the parameter values are lists, concatenate the values. If the parameter values are dictionaries, merge the values. If the two parameters have the same key, overwrite the former parameter value with the latter.

- Examples

```
{
  "Fn::Add": [
    {"Product": "ROS"},
    {"Fn": "Add"}
  ]
}
```

`{"Fn":"Add","Product":"ROS"}` is returned in this example.

5.4.6. Mappings

The Mappings section is a key-value mapping table. When mappings are used in resource and output definitions, use `Fn::FindInMap` to find their values by specifying corresponding keys.

Syntax

A mapping consists of key-value pairs, where both the keys and values can be strings or numbers. Multiple mappings are separated by commas (,). Each mapping name must be unique. Mappings must be pure data and cannot parse functions.

Examples

The following example shows a correct mapping definition:

```
"Mappings": {
  "ValidMapping": {
    "TestKey1": {"TestValu1": "value1"},
    "TestKey2": {"TestValu2": "value2"},
    1234567890: {"TestValu3": "value3"},
    "TestKey4": {"TestValu4": 1234}
  }
}
```

The following example shows an incorrect mapping definition:

```
"Mappings": {
  "InvalidMapping1": {
    "ValueList": ["foo", "bar"],
    "ValueString": "baz"
  },
  "InvalidMapping2": ["foo", {"bar": "baz"}],
  "InvalidMapping3": "foobar"
}
```

The following example shows how to use `Fn::FindInMap` to find the return value:

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "regionParam": {
      "Description": "The region where the ECS instance is created",
      "Type": "String",
      "AllowedValues": [
        "hangzhou",
        "beijing"
      ]
    }
  },
  "Mappings": {
    "RegionMap": {
      "hangzhou": {
        "32": "m-25l0rcfj0",
        "64": "m-25l0rcfj1"
      },
      "beijing": {
        "32": "m-25l0rcfj2",
        "64": "m-25l0rcfj3"
      }
    }
  },
  "Resources": {
    "WebServer": {
      "Type": "ALIYUN::ECS::Instance",
      "Properties": {
        "ImageId": {
          "Fn::FindInMap": [
            "RegionMap",
            {
              "Ref": "regionParam"
            }
          ],
          "32"
        }
      },
      "InstanceType": "ecs.t1.small",
      "SecurityGroupId": "sg-25zwc****",
      "ZoneId": "cn-beijing-b",
      "Tags": [
        {
          "Key": "Department1",
          "Value": "HumanResource"
        },
        {
          "Key": "Department2",
          "Value": "Finance"
        }
      ]
    }
  }
}
```

5.4.7. Conditions

Condition bodies are defined by Fn::And, Fn::Or, Fn::Not, and Fn::Equals operators. These operators, along with the parameters that you specify when you create or update a stack, are used to evaluate each condition. You can reference other conditions, parameters, and mappings in your condition. Conditions are used in resource and output definitions to establish dependencies. Use Fn::If or Condition in resource and output definitions to implement conditions.

Syntax

Each condition consists of a condition name and a condition body. The condition name is a string. The condition body starts with Fn::And, Fn::Or, Fn::Not, or Fn::Equals. You can reference other conditions in your condition. Separate multiple conditions with comma (.). Each condition name must be unique.

The following functions can be used, but not as the outermost functions:

"Fn::Select", "Fn::Join", "Fn::Split", "Fn::Replace", "Fn::Base64Encode", "Fn::Base64Decode", "Fn::MemberListToMap", "Fn::If", "Fn::ListMerge", "Fn::GetJsonValue", "Fn::MergeMapToList", "Fn::SelectMapList", "Fn::Add", "Fn::Avg", "Fn::Str", "Fn::Calculate", "Ref" (parameter references only), and "Fn::FindInMap".

Examples

- The following example shows how to define conditions:

```
"Conditions": {
  "DevEnv": {"Fn::Equals": ["Dev", {"Ref": "EnvType"}]},
  "UTEnv": {"Fn::Equals": ["UT", {"Ref": "EnvType"}]},
  "PREEnv": {"Fn::Not": {"Fn::Or": ["DevEnv", "UTEnv"]}},
  "ProdEnv": {"Fn::And": [{"Fn::Equals": ["Prod", {"Ref": "EnvType"}]}, "PREEnv"]}
}
```

- The following example shows how to use conditions in a resource definition.

In this example, a condition is used to determine whether to create a data disk and an Object Storage Service (OSS) bucket for an Elastic Compute Service (ECS) instance based on the EnvType value.

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "EnvType": {
      "Default": "pre",
      "Type": "String"
    }
  },
  "Conditions": {
    "CreateProdRes": {
      "Fn::Equals": [
        "prod",
        {
          "Ref": "EnvType"
        }
      ]
    }
  },
  "Resources": {
    "WebServer": {
      "Type": "ALIYUN::ECS::Instance",
      "Properties": {
        "DiskMappings": {
          "Fn::If": [
            "CreateProdRes",
            [
              {
                "Category": "cloud_efficiency",
```


5.5.1.1. ALIYUN::ECS::AutoSnapshotPolicy

ALIYUN::ECS::AutoSnapshotPolicy is used to create an automatic snapshot policy.

Statement

```
{
  "Type" : "ALIYUN::ECS::AutoSnapshotPolicy",
  "Properties" : {
    "TimePoints" : String,
    "RepeatWeekdays" : String,
    "RetentionDays" : Integer,
    "DiskIds" : String,
    "AutoSnapshotPolicyName" : String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|------------|------|----------|----------|---|--|
| TimePoints | List | Retained | Yes | The points in time at which automatic snapshots are created. Unit: hours. | <p>Value range: [0, 23], represents 24 time points from 00:00 to 23:00. For example:</p> <ul style="list-style-type: none"> 1 indicating 01:00. To schedule multiple automatic snapshot creation tasks in a day, you can set the TimePoints parameter as an array. The maximum number of time points allowed is 24. Use one format for multiple time points like [0, 1,... 23]. Separate time points with commas (,). |

| Parameter | Type | Required | Editable | Description | Constraint |
|----------------|---------|----------|----------|--|---|
| RepeatWeekdays | List | Retained | Yes | The days of a week on which automatic snapshots are created. | <p>Value range:[1, 7], 1 indicates Monday. To schedule multiple automatic snapshot creation tasks in a week, you can set the RepeatWeekdays parameter as an array.</p> <ul style="list-style-type: none"> You can specify up to 7 days over a one week period. Use one format for multiple time points like [1, 2,... 7]. Separate the time points with commas (,). |
| RetentionDays | Integer | Retained | Yes | The number of days for which you want to retain automatic snapshots. | <p>Default value: -1. Valid values:</p> <ul style="list-style-type: none"> -1: The automatic snapshots are retained indefinitely. [1, 65536]: The automatic snapshots are retained for the specified number of days. <p>Default value: -1.</p> |

| Parameter | Type | Required | Editable | Description | Constraint |
|------------------------|--------|----------|----------|--|---|
| DiskIds | List | Retained | Yes | The ID of the destination disk. When you want to apply the automatic snapshot policy to multiple disks, you can set the diskids "d-zzzzzzzz". Separate multiple disk IDs with commas (,). | None |
| AutoSnapshotPolicyName | String | Yes | True | The name of the automatic snapshot policy. | <ul style="list-style-type: none"> The name must be 2 to 128 characters in length It can contain letters, digits, colons (:), underscores (_), and hyphens (-). It cannot start with http:// or https://. <p>This parameter is empty by default.</p> |

Response parameters

Fn::GetAtt

AutoSnapshotPolicyId: the ID of the automatic snapshot policy.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "AutoSnapshotPolicy": {
      "Type": "ALIYUN::ECS::AutoSnapshotPolicy",
      "Properties": {
        "TimePoints": ["0"],
        "RepeatWeekdays": ["1"],
        "RetentionDays": 10,
        "DiskIds": ["<DiskId>"],
        "AutoSnapshotPolicyName": "MyAutoSnapshotPolicy"
      }
    }
  }
}
```

5.5.1.2. ALIYUN::ECS::BandwidthPackage

ALIYUN::ECS::BandwidthPackage is used to create a service plan for a NAT gateway.

Syntax

```
{
  "Type": "ALIYUN::ECS::BandwidthPackage",
  "Properties": {
    "Description": String,
    "NatGatewayId": String,
    "ZoneId": String,
    "BandwidthPackageName": String,
    "Bandwidth": Integer,
    "IpCount": Integer
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|--------------|---------|----------|----------|---|--|
| NatGatewayId | String | Yes | No | The ID of the NAT gateway to which you want to bind the service plan. | None |
| Bandwidth | Integer | Yes | No | The bandwidth. | Valid values: 5 to 5000. Unit: Mbit/s. Default value: 5. |
| IpCount | Integer | Yes | No | The number of public IP addresses assigned to the NAT gateway. | Valid values: 1 to 5. |
| Description | String | No | No | The description of the service plan. | The description must be 2 to 256 characters in length. |
| ZoneId | String | No | No | The ID of the zone where the NAT gateway resides. | None |

| Property | Type | Required | Editable | Description | Constraint |
|----------------------|--------|----------|----------|-------------------------------|---|
| BandwidthPackageName | String | No | No | The name of the service plan. | The name must be 2 to 128 characters in length and can contain letters, digits, underscores (_), hyphens (-), and periods (.). It must start with a letter. |

Response parameters

Fn::GetAtt

- BandwidthPackageId: the ID of the service plan.
- BandwidthPackageIps: all IP addresses included in the service plan.

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "BandwidthPackage": {
      "Type": "ALIYUN::ECS::BandwidthPackage",
      "Properties": {
        "BandwidthPackageName": "pkg_2",
        "Description": "my_bandwidth",
        "NatGatewayId": "ngw-h1xox****",
        "IpCount": 2,
        "Bandwidth": 5,
        "ZoneId": "cn-beijing-c"
      }
    }
  },
  "Outputs": {
    "BandwidthPackageId": {
      "Value": {"Fn::GetAttr": ["BandwidthPackage", "BandwidthPackageId"]}
    },
    "BandwidthPackageIps": {
      "Value": {"Fn::GetAttr": ["BandwidthPackage", "BandwidthPackageIps"]}
    }
  }
}
```

5.5.1.3. ALIYUN::ECS::Command

ALIYUN::ECS::Command is used to create a Cloud Assistant command.

Statement

```
{
  "Type": "ALIYUN::ECS::Command",
  "Properties": {
    "Name": String,
    "WorkingDir": String,
    "CommandContent": String,
    "Timeout": Integer,
    "Type": String,
    "Description": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|----------------|--------|----------|----------|---|------------|
| Name | String | Yes | True | The name of the command, which supports all character sets. The name can be up to 30 characters in length. | None |
| WorkingDir | String | Yes | True | The working directory on the ECS instance where the command will be run. | None |
| CommandContent | String | Yes | Released | The Base64-encoded content of the command. When you specify request parameters Type you must also specify this parameter. The parameter value must be Base64-encoded and cannot exceed 16 KB in size after encoding. | None |
| Timeout | String | No. | True | The timeout period that is specified for the command to run on ECS instances. Unit: seconds. If the command fails to run within the specified period, the command execution will time out and the process will be forcibly terminated. Default value: 3600. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------|--------|----------|----------|--|------------|
| Type | String | No | No | The command type. Valid values: <ul style="list-style-type: none"> RunBatScript: Creates a Bat script for a Windows instance. RunPowerShellScript: Create a PowerShell script to run on a Windows instance. RunShellScript: Creates a Shell script for Linux-based instances. | None |
| Description | String | Yes | True | The description of the command, which supports all character sets. The description can be up to 100 characters in length. | None |

Response parameters

Fn::GetAtt

CommandId: the ID of the command.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "WorkingDir": {
      "Type": "String",
      "Description": "The path where command will be executed in the instance."
    },
    "CommandContent": {
      "Type": "String",
      "Description": "The content of command. Content requires base64 encoding. Maximum size support 16KB."
    },
    "Type": {
      "Type": "String",
      "Description": "The type of command."
    },
    "Description": {
      "Type": "String",
      "Description": "The description of command."
    },
    "Timeout": {
      "Type": "Number",
      "Description": "Total timeout when the command is executed in the instance. Input the time unit as second. Default is 3600s."
    },
    "Name": {
      "Type": "String",
      "Description": "The name of command."
    }
  },
  "Resources": {
```

```
resources : {
  "Command": {
    "Type": "ALIYUN::ECS::Command",
    "Properties": {
      "WorkingDir": {
        "Ref": "WorkingDir"
      },
      "CommandContent": {
        "Ref": "CommandContent"
      },
      "Type": {
        "Ref": "Type"
      },
      "Description": {
        "Ref": "Description"
      },
      "Timeout": {
        "Ref": "Timeout"
      },
      "Name": {
        "Ref": "Name"
      }
    }
  },
  "Outputs": {
    "CommandId": {
      "Description": "The id of command created.",
      "Value": {
        "Fn::GetAtt": [
          "Command",
          "CommandId"
        ]
      }
    }
  }
}
```

5.5.1.4. ALIYUN::ECS::CustomImage

ALIYUN::ECS::CustomImage is used to create a custom image.

Statement

```
{
  "Type": "ALIYUN::ECS::CustomImage",
  "Properties": {
    "Description": String,
    "Instanceld": String,
    "ImageName": String,
    "ImageVersion": String,
    "SnapshotId": String,
    "Tag": List,
    "ResourceGroupId": String,
    "Platform": String,
    "DiskDeviceMapping": List,
    "Architecture": String
  }
}
```

Properties

| Parameter | Type | Required or Not | Editable | Description | Constraint |
|--------------|--------|-----------------|----------|-------------------------------|---|
| Description | String | Yes | Released | The description of the image. | The description can be up to 256 characters in length. This parameter is empty by default. It cannot start with http:// or https://. |
| Instanceld | String | Yes | Released | The ID of the ECS instance. | If this parameter is specified, an ECS instance will be used to create the custom image. |
| ImageName | String | Yes | Released | The name of the image. | The name must be 2 to 128 characters in length and can contain letters, digits, underscores (_), and hyphens(-). It must start with a letter but cannot start with http:// or https://. |
| ImageVersion | String | Yes | Released | The image version. | The image version must be 1 to 40 characters in length. |
| SnapshotId | String | Yes | Released | The ID of the snapshot. | <ul style="list-style-type: none"> If this parameter is specified, a snapshot will be used to create the custom image. If both this parameter and the Instanceld parameter are specified, this parameter will be ignored and an instance will be used to create the custom image. |

| Parameter | Type | Required or Not | Editable | Description | Constraint |
|-------------------|--------|-----------------|----------|---|--|
| Tags | List | Erased | Released | The tags of the image. | None |
| ResourceGroupId | String | Yes | Released | The ID of the resource group to which the custom image belongs. | None |
| Platform | String | Yes | Released | If you specify a data disk snapshot to be used to create the system disk of the custom image, you must use the Platform parameter to determine the release version of the operating system for the system disk. | None |
| DiskDeviceMapping | List | Erased | Released | The mappings between images and snapshots. | None |
| Architecture | String | Yes | Released | If you specify a data disk snapshot to be used to create the system disk of the custom image, you must use the Architecture parameter to determine the architecture of the system disk. Default value: x86_64. | Valid values: <ul style="list-style-type: none"> i386 x86_64 |

Tag syntax

```
"Tag": [
  {
    "Key": String,
    "Value": String
  }
]
```

Tag properties

| Parameter | Type | Required or Not | Editable | Description | Constraint |
|-----------|------|-----------------|----------|-------------|------------|
|-----------|------|-----------------|----------|-------------|------------|

| Parameter | Type | Required or Not | Editable | Description | Constraint |
|-----------|--------|-----------------|----------|-----------------------------|--|
| Key | String | Yes | Released | The tag key of the image. | The tag key cannot be a null string. The key can be up to 64 characters in length. It cannot start with aliyun or acs: and cannot contain http:// or https://. |
| Value | String | Yes | Released | The tag value of the image. | The tag value can be an empty string. The value can be up to 128 characters in length. It cannot start with aliyun or acs: and cannot contain http:// or https://. |

DiskDeviceMapping

```
"DiskDeviceMapping": [
  {
    "Device": String,
    "SnapshotId": String,
    "Size": Integer,
    "DiskType": String
  }
]
```

DiskDeviceMapping properties

| Parameter | Type | Required or Not | Editable | Description | Constraint |
|------------|--------|-----------------|----------|---|---|
| Device | String | Yes | Released | The device name of disk N in the custom image. | The system allocates a device name in alphabetical order from /dev/xvda to /dev/xvdz. |
| SnapshotId | String | Yes | Released | The ID of the snapshot that is used to create the custom image. | None |

| Parameter | Type | Required or Not | Editable | Description | Constraint |
|-----------|--------|-----------------|----------|--|---|
| Size | String | Optional | Released | The size of disk N. Unit: GiB. | Valid values: 5 to 2000. <ul style="list-style-type: none"> The default value is the size of the snapshot specified by the DiskDeviceMapping.N.Sn aphotoId parameter. If the DiskDeviceMapping.N.Sn aphotoId parameter is not specified, the default disk size is 5 GiB. The disk size must be greater than or equal to the size of the snapshot specified by the DiskDeviceMapping.N.Sn aphotoId parameter. |
| DiskType | String | Yes | Released | The type of disk N in the custom image. You can specify this parameter to create the system disk of the custom image from a data disk snapshot. If you do not specify this parameter, the disk type is determined by the corresponding snapshot. | Valid values: <ul style="list-style-type: none"> system: indicates a system disk. data: indicates a data disk. |

Response parameters

Fn::GetAtt

ImageId: the ID of the custom image.

Sample request

```

{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "WebServer": {
      "Type": "ALIYUN::ECS::InstanceGroup",
      "Properties": {
        "VpcId": "vpc-2zevx9ios1rszqv0a****",
        "MinAmount": 1,
        "SecurityGroupId": "sg-2ze7pxymaix640qr****",
        "ImageId": {
          "Ref": "CustomImage"
        },
        "IoOptimized": "optimized",
        "SystemDisk_Description": "SystemDisk.Description",
        "SystemDisk_DiskName": "SystemDisk.DiskName",
        "SystemDisk_Category": "cloud_ssd",
        "VSwitchId": "vsw-2zei67xd9nhcqxzec****",
        "Password": "Wenqiao****",
        "InstanceType": "ecs.n1.medium",
        "MaxAmount": 1
      }
    },
    "CustomImage": {
      "Type": "ALIYUN::ECS::CustomImage",
      "Properties": {
        "InstancelId": "i-2zefq1f3ynrrr89q****",
        "SnapshotId": "s-2ze0ibk1pvak4mw6****",
        "ImageName": "image-test-****",
        "ImageVersion": "verison-6-1"
      }
    }
  },
  "Outputs": {
    "CustomImage": {
      "Value": {
        "Fn::GetAtt": [
          "CustomImage",
          "ImageId"
        ]
      }
    },
    "InstancelIds": {
      "Value": {
        "Fn::GetAtt": [
          "WebServer",
          "InstancelIds"
        ]
      }
    }
  }
}

```

5.5.1.5. ALIYUN::ECS::DedicatedHost

ALIYUN::ECS::DedicatedHost is used to create a dedicated host.

Statement

```
{
  "Type": "ALIYUN::ECS::DedicatedHost",
  "Properties": {
    "DedicatedHostType": String,
    "DedicatedHostName": String,
    "AutoReleaseTime": String,
    "Description": String,
    "AutoPlacement": String,
    "Tags": List,
    "ActionOnMaintenance": String,
    "NetworkAttributesSlbUdpTimeout": Integer,
    "ChargeType": String,
    "ResourceGroupId": String,
    "ZoneId": String,
    "NetworkAttributesUdpTimeout": Integer,
    "Quantity": Integer
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------------|--------|----------|----------|---------------------------------|---|
| DedicatedHostType | String | No | No | The dedicated host type. | None |
| DedicatedHostName | String | Yes | Released | The name of the dedicated host. | <ul style="list-style-type: none"> The name must be 2 to 128 characters in length and can contain letters, digits, colons (:), underscores (_), and hyphens (-). Must start with an uppercase or lowercase letter, and cannot start with <code>http://</code> or <code>https://</code> the beginning. It can contain digits, colons (:), underscores (_), and hyphens (-). |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|---|---|
| AutoReleaseTime | String | Yes | Released | <p>The time scheduled for the dedicated host to be automatically released. If you do not specify the AutoReleaseTime parameter, the dedicated host will not be automatically released.</p> <ul style="list-style-type: none"> The minimum release time must be at least 30 minutes after the current time. The maximum release time must be at most three years from the current time. If the value of <code>ss</code> is not <code>00</code>, the start time is automatically rounded down to the nearest minute based on the value of <code>mm</code>. | None |
| Description | String | Yes | Released | The description of the dedicated host. | None |
| ZoneId | String | Yes | Released | <p>The ID of the zone where the dedicated host resides.</p> <p>This parameter is empty by default. If this parameter is not specified, the system will automatically select a zone.</p> | None |
| ChargeType | String | Yes | Released | The billing method of the dedicated host. | Valid values: PostPaid and pay-as-you-go. |

| Parameter | Type | Required | Editable | Description | Constraint |
|---------------|--------|----------|----------|---|--|
| AutoPlacement | String | Yes | Released | Specifies whether to add the dedicated host to the resource pool for automatic deployment. If you do not specify a DedicatedHostId when you create an instance on a DDH, Alibaba Cloud automatically selects a DDH from the resource pool to host the instance. | <p>Valid values:</p> <ul style="list-style-type: none"> on off <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note</p> <p>If you do not specify this parameter, the dedicated host is added to the automatic deployment resource pool.</p> <p>If you do not want to add the dedicated host to the resource pool for automatic deployment, set the value to off.</p> </div> |
| Tags | List | Erased | Released | The custom tags of the instance. | <p>A maximum of 20 tags are supported. The format is as follows:</p> <pre>[{"Key": "tagKey", "Value": "tagValue"}, {"Key": "tagKey2", "Value": "tagValue2"}]</pre> |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------------------|--------|----------|----------|---|--|
| ActionOnMaintenance | String | Yes | Released | The method used to migrate the instances on the DDH when the DDH fails or needs to be repaired online. | Valid values: <ul style="list-style-type: none"> Migrate: specifies that the instances are migrated to another physical server and restarted. Stop: specifies that all the instances on the DDH are stopped. If the DDH cannot be repaired, the instances are migrated to another physical server and restarted. The default value is "Migrate" for a dedicated host and "Stop" for a local disk. |
| NetworkAttributesSlbUdpTimeout | String | Optional | Released | The timeout period for a UDP session. | Valid values: 15 to 310. Unit: seconds. |
| ResourceGroupId | String | Yes | Released | The ID of the resource group to which the dedicated host belongs. | None |
| NetworkAttributesUdpTimeout | String | Optional | Released | The timeout period for UDP sessions that users can access for cloud services running on the dedicated host. | Valid values: 15 to 310. Unit: seconds. |
| Quantity | String | Optional | Released | The number of DDHs that you want to create this time. | Valid values: 1 to 100. Default value: 1 |

Response parameters

Fn::GetAtt

- OrderId: the ID of the order.
- DedicatedHostIds: the list of host IDs.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "AutoRenewPeriod": {
      "Type": "Number",
      "Description": "The time period of auto renew. When the parameter InstanceChargeType is PrePaid, it will take effect. It could be 1, 2, 3, 6, 12. Default value is 1.",
      "Default": "1"
    }
  }
}
```

```

    "AllowedValues": [
      1,
      2,
      3,
      6,
      12
    ],
    "Default": 1
  },
  "Description": {
    "Type": "String",
    "Description": "The description of host."
  },
  "ZoneId": {
    "Type": "String",
    "Description": "The zone to create the host."
  },
  "DedicatedHostName": {
    "Type": "String",
    "Description": "The name of the dedicated host, [2, 128] English or Chinese characters. It must begin with an uppercase/lowercase letter or a Chinese character, and may contain numbers, '_' or '-'. It cannot begin with http:// or https://."
  },
  "ChargeType": {
    "Type": "String",
    "Description": "Instance Charge type, allowed value: Prepaid and Postpaid. If specified Prepaid, please ensure you have sufficient balance in your account. Or instance creation will be failure. Default value is Postpaid.",
    "AllowedValues": [
      "PrePaid",
      "PostPaid"
    ],
    "Default": "PostPaid"
  },
  "AutoRenew": {
    "Type": "String",
    "Description": "Whether renew the fee automatically? When the parameter InstanceChargeType is PrePaid, it will take effect. Range of value: True: automatic renewal. False: no automatic renewal. Default value is False.",
    "AllowedValues": [
      "True",
      "False"
    ],
    "Default": "False"
  },
  "Period": {
    "Type": "Number",
    "Description": "Prepaid time period. Unit is month, it could be from 1 to 9 or 12, 24, 36, 48, 60. Default value is 1.",
    "AllowedValues": [
      1,
      2,
      3,
      4,
      5,
      6,
      7,
      8,
      9,
      12,
      24,
      36,
      48,
      60
    ]
  }
}

```

```

    },
    "Default": 1
  },
  "DedicatedHostType": {
    "Type": "String",
    "Description": "The instance type of host."
  },
  "PeriodUnit": {
    "Type": "String",
    "Description": "Unit of prepaid time period, it could be Week/Month. Default value is Month.",
    "AllowedValues": [
      "Week",
      "Month"
    ],
    "Default": "Month"
  },
  "AutoReleaseTime": {
    "Type": "String",
    "Description": "Auto release time for created host, Follow ISO8601 standard using UTC time. format is 'yyyy-MM-ddTHH:mm:ssZ'. Not bigger than 3 years from this day onwards"
  },
  "Resources": {
    "Host": {
      "Type": "ALIYUN::ECS::DedicatedHost",
      "Properties": {
        "Description": {
          "Ref": "Description"
        },
        "ZoneId": {
          "Ref": "ZoneId"
        },
        "DedicatedHostName": {
          "Ref": "DedicatedHostName"
        },
        "ChargeType": {
          "Ref": "ChargeType"
        },
        "DedicatedHostType": {
          "Ref": "DedicatedHostType"
        },
        "PeriodUnit": {
          "Ref": "PeriodUnit"
        },
        "AutoReleaseTime": {
          "Ref": "AutoReleaseTime"
        }
      }
    },
    "Outputs": {
      "OrderId": {
        "Description": "The order id list of created instance.",
        "Value": {
          "Fn::GetAtt": [
            "Host",
            "OrderId"
          ]
        }
      }
    }
  }
}

```

```
"DedicatedHostIds": {
  "Description": "The host id list of created hosts",
  "Value": {
    "Fn::GetAtt": [
      "Host",
      "DedicatedHostIds"
    ]
  }
}
}
```

5.5.1.6. ALIYUN::ECS::Disk

ALIYUN::ECS::Disk is used to create an ECS Disk.

Statement

```
{
  "Type": "ALIYUN::ECS::Disk",
  "Properties": {
    "DiskName": String,
    "Description": String,
    "Tags": List,
    "AutoSnapshotPolicyId": String,
    "Encrypted": Boolean,
    "ZoneId": String,
    "ResourceGroupId": String,
    "SnapshotId": String,
    "DiskCategory": String,
    "PerformanceLevel": String,
    "DeleteAutoSnapshot": Boolean,
    "Size": Integer
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|---|------------|
| ResourceGroupId | String | Yes | Released | The ID of the resource group to which the instance belongs. | None |
| ZoneId | String | No | No | The ID of the zone where the instance resides. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------|--------|----------|----------|----------------------------------|---|
| DiskName | String | Yes | Released | The name of the disk. | <ul style="list-style-type: none"> The name must be 2 to 128 characters in length And can contain letters, digits, periods, underscores (_), and hyphens (-). It cannot start with http:// or https://. The disk name will be displayed in the ECS console. |
| Description | String | Yes | Released | The description of the disk. | <ul style="list-style-type: none"> The description must be 2 to 256 characters in length. Cannot start with http:// or https:// the beginning. The disk description will be displayed in the ECS console. |
| Tags | List | Erased | Released | The custom tags of the instance. | <p>Up to four tags are supported. Example values: [{"Key": "tagKey", "Value": "tagValue"}, {"Key": "tagKey2", "Value": "tagValue2"}].</p> |
| DiskCategory | String | Yes | Released | The type of the data disk. | <p>Value range</p> <ul style="list-style-type: none"> cloud: indicates a basic disk. cloud_efficiency: indicates an ultra disk. cloud_ssd: indicates a standard SSD. cloud_essd: enhanced SSD (ESSD) <p>Default value: cloud.</p> |

| Parameter | Type | Required | Editable | Description | Constraint |
|----------------------|---------|----------|----------|---|--|
| SnapshotId | String | Yes | Released | The ID of the snapshot used to create the data disk. | <ul style="list-style-type: none"> If both this parameter and 'Size' are specified, the value of this parameter prevails. The actual size of the created disk is the size of the specified snapshot. Snapshots created on or before July 15, 2013 cannot be used to create disks. |
| PerformanceLevel | String | Yes | Released | Specifies the performance level of an ESSD when you create the ESSD. | Default value: PL1. Valid values: <ul style="list-style-type: none"> PL1: A single enhanced SSD delivers up to 50,000 random read/write IOPS. PL2: A single ESSD delivers up to 100,000 random read/write IOPS. PL3: maximum random read/write IOPS of 100,000 per disk. |
| Size | String | Optional | Released | The size of the disk. Unit: GiB. The value of this parameter must be equal to or greater than the capacity of the specified snapshot. | Valid values: <ul style="list-style-type: none"> cloud: 5 to 2000 cloud_efficiency: 20 to 32768 cloud_ssd: 20 to 32768 cloud_essd: 20 to 32768 |
| AutoSnapshotPolicyId | String | Yes | Released | The ID of each automatic snapshot policy. | None |
| Encrypted | Boolean | Erased | Released | Specifies whether to encrypt the disk. | Valid values: <ul style="list-style-type: none"> true false Default value: false. |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------|---------|----------|----------|--|---|
| DeleteAutoSnapshot | Boolean | Erased | Released | Specifies whether to delete the automatic snapshots of the disk when the disk is released. | Valid values: <ul style="list-style-type: none"> • true • false Default value: true. |

Tags syntax

```
"Tags": [
  {
    "Value": String,
    "Key": String
  }
]
```

Tags properties

| Parameter | Type | Required | Editable |
|-----------|--------|----------|----------|
| Key | String | No | No |
| Value | String | Yes | Released |

Response parameters

Fn::GetAtt

- DiskId: the ID of the disk.
- Status: The Status of the disk.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "DataDisk": {
      "Type": "ALIYUN::ECS::Disk",
      "Properties": {
        "Size": 10,
        "ZoneId": "cn-beijing-a",
        "DiskName": "DataDisk",
        "Description": "ECSDataDisk"
      }
    }
  },
  "Outputs": {
    "DiskId": {
      "Value": {"Fn::GetAtt": ["DataDisk", "DiskId"]}
    },
    "Status": {
      "Value": {"Fn::GetAtt": ["DataDisk", "Status"]}
    }
  }
}
```

5.5.1.7. ALIYUN::ECS::DiskAttachment

ALIYUN::ECS::DiskAttachment is used to attach an ECS disk.

Statement

```
{
  "Type": "ALIYUN::ECS::DiskAttachment",
  "Properties": {
    "DiskId": String,
    "InstanceId": String,
    "Device": String,
    "DeleteWithInstance": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|------------|--------|----------|----------|-------------------------|---|
| InstanceId | String | No | No | The ID of the instance. | None |
| DiskId | String | No | No | The ID of the disk. | The disk and the ECS instance must belong to the same zone. |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------|---------|----------|----------|--|---|
| Device | String | Yes | Released | The name of the disk. | If you do not set this parameter, the system will automatically allocate a device name in alphabetical order from /dev/xvdb to /dev/xvdz. |
| DeleteWithInstance | Boolean | Erased | Released | Specifies whether the disk is to be released together with the instance. | Valid values: <ul style="list-style-type: none"> • true: The disk will be released when the instance is released. • false: The disk will be retained when the instance is released. |

Response parameters

Fn::GetAtt

- DiskId: the ID of the disk.
- Status: The Status of the disk.
- The name of the Device: disk.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "DiskAttachment": {
      "Type": "ALIYUN::ECS::DiskAttachment",
      "Properties": {
        "InstancedId": {
          "Ref": "InstancedId"
        },
        "Device": {
          "Ref": "Device"
        },
        "DeleteWithInstance": {
          "Ref": "DeleteWithInstance"
        },
        "DiskId": {
          "Ref": "DiskId"
        }
      }
    }
  },
  "Parameters": {
    "InstancedId": {
```

```

    "Type": "String",
    "Description": "The ID of the instance to attach the disk."
  },
  "Device": {
    "Type": "String",
    "Description": "The device where the volume is exposed on the instance. The device name could be /dev/xvd[a-z]. If this parameter is not specified, the default value will be used."
  },
  "DeleteWithInstance": {
    "Type": "Boolean",
    "Description": "If this parameter is set to true, the disk will be deleted while the instance is deleted. If this parameter is set to false, the disk will be retained after the instance is deleted.",
    "AllowedValues": [
      "True",
      "true",
      "False",
      "false"
    ]
  },
  "DiskId": {
    "Type": "String",
    "Description": "The ID of the disk to be attached."
  }
},
"Outputs": {
  "Status": {
    "Description": "The disk status now.",
    "Value": {
      "Fn::GetAtt": [
        "DiskAttachment",
        "Status"
      ]
    }
  },
  "Device": {
    "Description": "The device where the volume is exposed on the ECS instance.",
    "Value": {
      "Fn::GetAtt": [
        "DiskAttachment",
        "Device"
      ]
    }
  },
  "DiskId": {
    "Description": "The ID of the created disk.",
    "Value": {
      "Fn::GetAtt": [
        "DiskAttachment",
        "DiskId"
      ]
    }
  }
}
}
}
}
}

```

5.5.1.8. ALIYUN::ECS::ForwardEntry

ALIYUN::ECS::ForwardEntry is used to configure the DNAT table of a NAT Gateway.

Statement

```
{
  "Type": "ALIYUN::ECS::ForwardEntry",
  "Properties": {
    "ExternalIp": String,
    "ExternalPort": String,
    "ForwardTableId": String,
    "InternalIp": String,
    "IpProtocol": String,
    "InternalPort": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|----------------|--------|----------|----------|---|--|
| ExternalIp | String | No | No | The public IP address of the NAT gateway. | It must be an IP address that is included in the shared NAT Gateway of the bandwidth plan to which the DNAT table belongs. |
| ExternalPort | String | No | No | The public port number. | Valid values: 1 to 65535. |
| ForwardTableId | String | No | No | The ID of the DNAT table. | None |
| InternalIp | String | No | No | The destination IP address to which the request is forwarded. | This IP address is a private IP address. |
| IpProtocol | String | No | No | The type of the protocol. | Valid values: TCP, UDP, and Any. |
| InternalPort | String | No | No | The destination private port. | Valid values: 1 to 65535. |

Response parameters

Fn::GetAtt

ForwardEntryId: the ID of each entry in the DNAT table.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "ForwardEntry": {
      "Type": "ALIYUN::ECS::ForwardEntry",
      "Properties": {
        "ForwardTableId": "my_forwardtable",
        "ExternalIp": "101.201.XX.XX",
        "ExternalPort": "8080",
        "IpProtocol": "TCP",
        "InternalIp": "10.2.XX.XX",
        "InternalPort": "80"
      }
    }
  },
  "Outputs": {
    "ForwardEntryId": {
      "Value": {"Fn::GetAttr": ["ForwardEntry", "ForwardEntryId"]}
    }
  }
}
```

5.5.1.9. ALIYUN::ECS::Instance

ALIYUN::ECS::Instance is used to create an ECS instance.

Statement

```
{
  "Type": "ALIYUN::ECS::Instance",
  "Properties": {
    "RamRoleName": String,
    "IoOptimized": String,
    "PrivateIpAddress": String,
    "KeyPairName": String,
    "SystemDiskDiskName": String,
    "Description": String,
    "Tags": List,
    "HostName": String,
    "ImageId": String,
    "ResourceGroupId": String,
    "VSwitchId": String,
    "Password": String,
    "InstanceType": String,
    "SystemDiskCategory": String,
    "UserData": String,
    "SystemDiskSize": Number,
    "ZoneId": String,
    "VpcId": String,
    "InstanceName": String,
    "InternetMaxBandwidthIn": Integer,
    "DeletionProtection": Boolean,
    "DeploymentSetId": String,
    "SecurityGroupId": String,
    "HpcClusterId": String,
    "SystemDiskDescription": String,
    "DiskMappings": List
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|---|------------|
| ResourceGroupId | String | Yes | Released | The ID of the resource group to which the instance belongs. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|--|--|
| ImageId | String | No | Yes | The ID of the image used to start the ECS instance. You can use a public image, a custom image, or an Alibaba Cloud Marketplace image. | <p>When editing a template, you can specify the image type and version or only the image type. ROS automatically selects an appropriate public image ID.</p> <p>You can use the wildcard character (*) to represent part of an image ID.</p> <p>Take all Ubuntu public images provided by Alibaba Cloud as an example. You can use one of the following methods to specify the public image ID for the ECS instance:</p> <ul style="list-style-type: none"> If you enter ubuntu, the system matches it with the following ID: ubuntu16_0402_64_20G_alibase_20170818.vhd If you enter ubuntu_14, the system matches it with the following ID: ubuntu_14_0405_64_20G_alibase_20170824.vhd If you enter ubuntu*14*32, the system matches it with the following ID: ubuntu_14_0405_32_40G_alibase_20170711.vhd If you enter ubuntu_16_0402_32, the system matches it with the following ID: ubuntu_16_0402_32_40G_alibase_20170711.vhd |
| InstanceType | String | No | No | The type of the ECS instance. | None |
| SecurityGroupId | String | No | No | The ID of the security group to which the created instance will belong. | None |
| Description | String | Yes | Released | The description of created instances. | The description can be up to 256 characters in length. |

| Parameter | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|--|---|
| InstanceName | String | Yes | Released | The name of a created instance. | The name can be up to 128 characters in length and can contain letters, digits, underscores (_), periods (.), and hyphens (-). |
| Password | String | Yes | Released | The password used to log on to the ECS instance. | <p>The characters in length is 8 to 30.</p> <p>And must contain at least one of the following character types: uppercase letters, lowercase letters, digits, and special character.</p> <p>Special characters include () ` ~ ! @ # \$ % ^ & * - + = { } [] ; ' < > , . ? / - / - If you specify the Password parameter in the API request, use HTTPS to secure the API and protect your password.</p> |
| HostName | String | Yes | Released | The hostname of the instance. | <p>The password must be at least 2 characters in length.</p> <p>It cannot And hyphens (-) cannot start or end the hostname and cannot be used consecutively.</p> <p>On Windows, the hostname can be up to 15 characters in length and can contain letters, digits, and hyphens (-). It cannot contain periods (.) and cannot be composed of only digits.</p> <p>On other OSs such as Linux, the hostname can contain a maximum of 30 characters, including periods (.), each segment can contain uppercase or lowercase letters, digits, and hyphens (-).</p> |
| PrivateIpAddress | String | Yes | Released | The private IP address of an ECS instance in a VPC. The specified IP address must not be used by other instances in the VPC. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|------------------------|--------|----------|----------|---|---|
| InternetMaxBandwidthIn | String | Optional | Released | The maximum inbound bandwidth from the Internet. | Valid values: 1 to 100. Default value: 100. Unit: Mbit/s. |
| IoOptimized | String | Yes | Released | Specifies whether an I/O optimized instance is created. | Valid values: <ul style="list-style-type: none"> • none (non-I/O optimized) • optimized Default value: none. |
| DiskMappings | List | Erased | Released | The data disks to be attached to the instance. | A maximum of 16 disks can be attached to each instance. |
| SystemDiskCategory | String | Yes | Released | The type of the system disk. | Valid values: <ul style="list-style-type: none"> • cloud • cloud_efficiency • cloud_ssd • ephemeral_ssd |
| SystemDiskDescription | String | Yes | Released | The description of the ECS instance system disk. | None |
| SystemDiskDiskName | String | Yes | Released | The name of the ECS instance system disk. | None |
| SystemDiskSize | Number | No. | True | The size of the system disk. Unit: GB. | Valid values: 40 to 500. If a custom image is used to create a system disk, make sure that the size of the system disk is greater than that of the custom image. |
| Tags | List | Erased | Released | The custom tags of the instance. | A maximum of 20 tags are supported. The format is as follows: <pre>[{"Key": "tagKey", "Value": "tagValue"}, {"Key": "tagKey2", "Value": "tagValue2"}]</pre> |
| UserData | String | Yes | Released | The user data that you provide when you create ECS instances. | The user data can be up to 16KB in size. You do not need to use Base64. you must use special characters. \ Escape. |
| ZoneId | String | Yes | Released | The ID of the zone where the instance resides. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------|---------|----------|----------|--|--|
| HpcClusterId | String | Yes | Released | The ID of the HPC cluster to which the ECS instance belongs. | None |
| VpcId | String | Yes | Released | The ID of the VPC to which the ECS instance belongs. | None |
| VSwitchId | String | Yes | Released | The ID of the VSwitch for the ECS instance. | None |
| KeyPairName | String | Yes | Released | The name of the key pair that is used to connect to created ECS instances. | <ul style="list-style-type: none"> For Windows-based instances, this parameter is empty by default. In the Linux, if this parameter is specified, the Password content is still set to the instance, but the Password logon method is disabled by default. The key pair is used to verify the logon. |
| RamRoleName | String | Yes | Released | The RAM role name of the instance. You can call the ListRoles operation to query the role name. | None |
| DeletionProtection | Boolean | Erased | Released | The release protection property of created instances. It specifies whether the instances can be released from the ECS console or through the DeleteInstance operation. | Valid values: <ul style="list-style-type: none"> True False |
| DeploymentSetId | String | Yes | True | Deployment Set ID. | None |

DiskMappings syntax

```
"DiskMappings": [
  {
    "Category": String,
    "DiskName": String,
    "Description": String,
    "Device": String,
    "SnapshotId": String,
    "Size": String
  }
]
```

DiskMappings properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------|--------|----------|----------|--|---|
| Size | String | No | No | The size of data disk N. Unit: GB. | None |
| Category | String | Yes | Released | The type of the data disk. | Valid values: <ul style="list-style-type: none"> cloud cloud_efficiency cloud_ssd ephemeral_ssd Default value: cloud. |
| DiskName | String | Yes | Released | The name of data disk N. | The name can be up to 128 characters in length and can contain letters, digits, underscores (_), periods (.), and hyphens (-). |
| Description | String | Yes | Released | The description of data disk N. | Valid values: 2 to 256. Default value: Null. |
| Device | String | Yes | Released | The device name of the data disk. | If you do not specify this parameter, the system automatically allocates a device name in alphabetical order from /dev/xvdb to /dev/xvdz. |
| SnapshotId | String | Yes | Released | The ID of the snapshot used to create the data disk. | None |

Tags syntax

```
"Tags": [
  {
    "Value": String,
    "Key": String
  }
]
```

Tags properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|-------------|---|
| Key | String | No | No | None | None |
| Value | String | Yes | Released | None | Valid values: 1025 to 10000. You cannot use the following commonly used port numbers: 2222, 4500, 4510, 4560, 7505, 9000, 9001, and 9002. |

Response parameters

Fn::GetAtt

- **Instanceld**: the ID of the instance. An ID is a globally unique identifier (GUID) generated by the system for an instance.
- **PrivateIp**: The private IP address of the instance in a VPC. This parameter takes effect only when the **NetworkType** parameter is set to VPC.
- **InnerIp**: The private IP address of the instance in a Classic network. This parameter takes effect only when the **NetworkType** parameter is set to Classic.
- **PublicIp**: the public IP address of the instance in a Classic network. This parameter takes effect only when the **NetworkType** parameter is set to Classic.
- **ZoneId**: the zone ID.
- **HostName**: the hostname of the instance.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "WebServer": {
      "Type": "ALIYUN::ECS::Instance",
      "Properties": {
        "ImageId": "m-25l0rc****",
        "InstanceType": "ecs.t1.small",
        "SecurityGroupId": "sg-25zwc****",
        "ZoneId": "cn-beijing-b",
        "Tags": [{
          "Key": "tiantt",
          "Value": "ros"
        }],
        "Key": "tiantt1",
        "Value": "ros1"
      }
    }
  },
  "Outputs": {
    "Instanceid": {
      "Value": {"get_attr": ["WebServer", "Instanceid"]}
    },
    "PublicIp": {
      "Value": {"get_attr": ["WebServer", "PublicIp"]}
    }
  }
}
```

5.5.1.10. ALIYUN::ECS::InstanceClone

ALIYUN::ECS::InstanceClone is used to clone an ECS instance.

Statement

```
{
  "Type": "ALIYUN::ECS::InstanceClone",
  "Properties": {
    "DeletionProtection": Boolean,
    "DiskMappings": List,
    "LoadBalancerIdToAttach": String,
    "Description": String,
    "BackendServerWeight": Integer,
    "Tags": List,
    "SecurityGroupId": String,
    "RamRoleName": String,
    "ImageId": String,
    "ResourceGroupId": String,
    "SpotPriceLimit": String,
    "InstanceChargeType": String,
    "SourceInstanceId": String,
    "Period": Number,
    "SpotStrategy": String,
    "Password": String,
    "InstanceName": String,
    "InternetMaxBandwidthIn": Integer,
    "ZoneId": String,
    "KeyPairName": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|------------------------|--------|----------|----------|---|--|
| ResourceGroupId | String | Yes | Released | The ID of the resource group to which the instance belongs. | None |
| SourceInstanceId | String | No | No | The ID of the ECS instance to be cloned. | The clone operation clones the specified instance, including its instance type, image, bandwidth billing method, bandwidth limit, and network type. If the source ECS instance belongs to multiple security groups, the cloned instance is added only to the first of these security groups. |
| BackendServerWeight | String | Optional | Released | The weight of the ECS instance in the Server Load Balancer instance. | Value range:[0, 100]. Default value: 100. |
| LoadBalancerIdToAttach | String | Yes | Released | The ID of the SLB instance to which the ECS instance is to be attached. | None |
| Description | String | Yes | Released | The description of created instances. | The description can be up to 256 characters in length. |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|---|---|
| ImageId | String | Yes | True | The ID of the image used to start created instances. You can use a public image, a custom image, or an Alibaba Cloud Marketplace image. | <p>When editing a template, you can specify the image type and version or only the image type. ROS automatically selects an appropriate public image ID.</p> <p>You can use the wildcard character (*) to represent part of an image ID.</p> <p>Take all Ubuntu public images provided by Alibaba Cloud as an example. You can use one of the following methods to specify the public image ID for the ECS instance:</p> <p>If you enter ubuntu,</p> <p>the system matches it with the following ID: ubuntu16_0402_64_20G_alibase_20170818.vhd</p> <p>If you enter ubuntu_14,</p> <p>the system matches it with the following ID: ubuntu_14_0405_64_20G_alibase_20170824.vhd</p> <p>If you enter ubuntu*14*32,</p> <p>the system matches it with the following ID: ubuntu_14_0405_32_40G_alibase_20170711.vhd</p> <p>If you enter ubuntu_16_0402_32,</p> <p>the system matches it with the following ID: ubuntu_16_0402_32_40G_alibase_20170711.vhd</p> |
| SecurityGroupId | String | Yes | Released | The ID of the security group to which the created instance will belong. | None |
| InstanceName | String | Yes | Released | The name of a created instance. | The name can be up to 128 characters in length and can contain letters, digits, underscores (_), periods (.), and hyphens (-). |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------|--------|----------|----------|--|--|
| Password | String | Yes | Released | The password used to log on to the ECS instance. | <p>The password must be 8 to 30 characters in length.</p> <p>The password must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters.</p> <p>Special characters include () ` ~ ! @ # \$ % ^ & * - + = { } [] : ; < > , . ? /</p> <p>If you specify the password parameter in the API request, use HTTPS to secure the API and protect your password.</p> |
| DiskMappings | List | Erased | Released | The disks to be attached to created instances. | A maximum of 16 disks can be attached to each instance. |
| Tags | List | Erased | Released | The custom tags of the instance. | <p>A maximum of 20 tags can be specified in the</p> <pre>[{"Key": "tagKey", "Value": "tagValue"}, {"Key": "tagKey2", "Value": "tagValue2"}]</pre> |
| ZoneId | String | Yes | Released | The ID of the zone where the instance resides. | None |
| InstanceChargeType | String | Yes | Released | The billing method of the new ECS instance. | <p>Valid values: PrePaid and PostPaid.</p> <p>Default value: Postpaid. If you set this parameter to Prepaid, make sure that you have sufficient balance in your account. Otherwise, the instance fails to be created.</p> |
| Period | Number | Erased | Released | The subscription period of the new ECS instance. This parameter is required when the InstanceChargeType parameter is set to PrePaid. This parameter is ignored when the InstanceChargeType parameter is set to PostPaid. | Valid values: 1, 2, 3, 4, 5, 6, 7, 8, 9, 12, 24, and 36. Unit: month. |

| Parameter | Type | Required | Editable | Description | Constraint |
|------------------------|---------|----------|----------|---|---|
| KeyPairName | String | Yes | Released | The name of the key pair that is used to connect to created ECS instances. | For Windows-based instances, this parameter is empty by default. In the Linux, if this parameter is specified, the Password content is still set to the instance, but the Password logon method is disabled by default. The key pair is used to verify the logon. |
| RamRoleName | String | Yes | Released | The RAM role name of the instance. You can call the ListRoles operation to query the role name. | None |
| SpotPriceLimit | String | Yes | Released | The maximum hourly price of the instance. | Parameter SpotStrategy this parameter takes effect only when the value is SpotWithPriceLimit. |
| SpotStrategy | String | Yes | Released | The bidding policy for the pay-as-you-go instance. | This parameter is valid only when the InstanceChargeType parameter is set to PostPaid. Valid values: NoSpot: applies to regular pay-as-you-go instances. SpotWithPriceLimit: applies to preemptible instances with a maximum hourly price. SpotAsPriceGo: applies to pay-as-you-go instances priced at the market price at the time of purchase. Default value: NoSpot. |
| InternetMaxBandwidthIn | String | Optional | Released | The maximum inbound bandwidth from the Internet. Unit: Mbit/s. | Valid values: 1 to 200. |
| DeletionProtection | Boolean | Erased | Released | Specifies whether to enable instance release protection in the console or by calling an API operation. (DeleteInstance) Release instances. | Valid values: true and false. |

DiskMappings syntax

```
"DiskMappings": [
{
  "Category": String,
  "DiskName": String,
  "Description": String,
  "Device": String,
  "SnapshotId": String,
  "Size": String
}
]
```

DiskMappings properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------|--------|----------|----------|--|---|
| Size | String | No | No | The size of data disk N. Unit: GB. | None |
| Category | String | Yes | Released | The type of the data disk. | Valid values: cloud, cloud_efficiency, cloud_ssd, and ephemeral_ssd Default value: cloud. |
| DiskName | String | Yes | Released | Disk name. | The name can be up to 128 characters in length and can contain letters, digits, underscores (_), periods (.), and hyphens (-). |
| Description | String | Yes | Released | The description of data disk N. | The description must be 2 to 256 characters in length. This parameter is empty by default. |
| Device | String | Yes | Released | The device name of the data disk. | If you do not specify this parameter, the system automatically allocates a device name in alphabetical order from /dev/xvdb to /dev/xvdz. |
| SnapshotId | String | Yes | Released | The ID of the snapshot used to create the data disk. | None |

Tags syntax

```
"Tags": [
{
  "Value": String,
  "Key": String
}
]
```

Tags properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|--|---|
| Key | String | No | No | The tag key, which cannot be an empty string. It can be up to 64 characters in length, cannot start with acs: or aliyun, and cannot contain http:// or https://. | None |
| Value | String | Yes | Released | The tag value, which can be an empty string. It can be up to 128 characters in length and cannot start with acs: or aliyun. It cannot contain http:// or https://. | Valid values: 1025 to 10000. You cannot use the following commonly used port numbers: 2222, 4500, 4510, 4560, 7505, 9000, 9001, and 9002. |

Response parameters

Fn::GetAtt

- **Instanceld:** the ID of the instance. An ID is a globally unique identifier (GUID) generated by the system for an instance.
- **PrivateIp:** The private IP address of the instance in a VPC. This parameter takes effect only when the **NetworkType** parameter is set to VPC.
- **InnerIp:** The private IP address of the instance in a Classic network. This parameter takes effect only when the **NetworkType** parameter is set to Classic.
- **PublicIp:** the public IP address of the instance in a Classic network. This parameter takes effect only when the **NetworkType** parameter is set to Classic.
- **Zoneld:** the zone ID.
- **HostName:** the hostname of the instance.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "WebServer": {
      "Type": "ALIYUN::ECS::InstanceClone",
      "Properties": {
        "SourceInstanceId": "i-25zsk****",
        "SecurityGroupId": "sg-25zwc****",
        "ZoneId": "cn-beijing-b",
        "DiskMappings": [
          {"Size": 10, "Category": "cloud"},
          {"Size": 10, "Category": "cloud", "SnapshotId": "s-25wsw****"}
        ]
      }
    }
  },
  "Outputs": {
    "InstanceId": {
      "Value": {"get_attr": ["WebServer", "InstanceId"]}
    },
    "PublicIp": {
      "Value": {"get_attr": ["WebServer", "PublicIp"]}
    }
  }
}
```

5.5.1.11. ALIYUN::ECS::InstanceGroup

ALIYUN::ECS::InstanceGroup is used to create an ECS instance group.

Syntax

```
{
  "Type": "ALIYUN::ECS::InstanceGroup",
  "Properties": {
    "SystemDiskAutoSnapshotPolicyId": String,
    "DedicatedHostId": String,
    "LaunchTemplateName": String,
    "RamRoleName": String,
    "IoOptimized": String,
    "PrivateIpAddress": String,
    "KeyPairName": String,
    "SystemDiskDiskName": String,
    "Description": String,
    "Tags": List,
    "HostName": String,
    "ImageId": String,
    "ResourceGroupId": String,
    "VSwitchId": String,
    "EniMappings": List,
    "Password": String,
    "InstanceType": String,
    "MaxAmount": Integer,
    "AutoReleaseTime": String,
    "SystemDiskCategory": String,
    "UserData": String,
    "LaunchTemplateId": String,
    "LaunchTemplateVersion": String,
    "SystemDiskSize": Number,
    "ZoneId": String,
    "VpcId": String,
    "InternetMaxBandwidthIn": Integer,
    "DeletionProtection": Boolean,
    "DeploymentSetId": String,
    "Ipv6AddressCount": Integer,
    "SecurityGroupId": String,
    "HpcClusterId": String,
    "SystemDiskDescription": String,
    "Ipv6Addresses": List,
    "NetworkType": String,
    "DiskMappings": List,
    "SystemDiskPerformanceLevel": String
  }
}
```

Properties

| Attribute | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|---|------------|
| ResourceGroupId | String | No | Yes | The ID of the resource group to which a created instance belongs. | None. |

| Attribute | Type | Required | Editable | Description | Constraint |
|-----------------|---------|-----------|----------|---|--|
| HpcClusterId | String | No | Yes | The ID of the HPC cluster to which a created instance belongs. | None. |
| MaxAmount | Integer | Supported | Yes | The maximum number of ECS instances that can be created at a time. | Valid values: 1 to 100. The MaxAmount parameter must be set to a value greater than or equal to the value of MinAmount. |
| MinAmount | String | Yes | Yes | The minimum number of ECS instances that can be created at a time. | Valid values: 1 to 100. The MinAmount parameter must be set to a value less than or equal to the value of MaxAmount. |
| Description | String | No | Yes | The description of created instances. | The description can be up to 256 characters in length. |
| InstanceType | String | Yes | Yes | The type of the ECS instance. | None. |
| ImageId | String | Yes | Yes | The ID of the image used to start an ECS instance. You can use a public image, a custom image, or an Alibaba Cloud Marketplace image. | You can specify a partial public image ID instead of providing the complete ID. The following example shows how to use a CNAME record: <ul style="list-style-type: none"> • Ubuntu is specified and ubuntu_16_0402_64_20G_alibase_20170818.vhd are matched. • If ubuntu1432 is specified, ubuntu_14_0405_32_40G_alibase_20170711.vhd is matched. |
| SecurityGroupId | String | No | No | The ID of the security group to which created instances belong. | None. |

| Attribute | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|--|--|
| InstanceName | String | No | No | The name of an instance. | The names can be up to 128 characters in length. It can contain English letters, Chinese characters, digits, underscores (_), periods (.), and hyphens (-). By <code>name_prefix[begin number,bits]name_suffix</code> format to specify different instance name for each ECS instance. |
| Password | String | No | Yes | The password used to log on to created ECS instances. | <ul style="list-style-type: none"> The password must be 8 to 30 characters in length and must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. The following special characters are supported: <code>:() ~ ! @ # \$ % ^ & * - + = { } [] ; ' < > , . ? /</code> <p>If you specify the Password parameter in the API request, use HTTPS to secure the API and protect your password.</p> |
| HostName | String | No | No | The hostname of created ECS instances. | The hostname must contain at least two characters in length. The periods and hyphens (-) cannot start or end the hostname and cannot be used together. |
| AutoReleaseTime | String | No | No | The time scheduled for created ECS instances to be automatically released. | The time format must comply with ISO8601 specifications. for example. <code>"vvvv-MM-ddTHH:mm:ssZ"</code> . The maximum release time must be within three years from the current time. |
| PrivateIpAddress | String | No | No | The private IP address of an ECS instance in a VPC. | The specified IP address must not be used by other instances in the VPC. |

| Attribute | Type | Required | Editable | Description | Constraint |
|------------------------|---------|----------|----------|--|---|
| DiskMappings | List | No | Yes | The data disks to be attached to created instances. | None. |
| InternetMaxBandwidthIn | Integer | No | No | The maximum inbound bandwidth from the Internet. | Unit: Mbit/s. Valid values: 1 to 100 Default value: 100. |
| IoOptimized | String | No | No | Specifies whether the created instances are I/O optimized. | Valid values: <ul style="list-style-type: none"> • none (non-I/O optimized) • optimized Default value: none. |
| SystemDiskCategory | String | No | Yes | The category of the system disk. | Valid values: cloud: basic disk cloud_efficiency: the ultra disk cloud_essd: standard SSDs cloud_essd: enhanced SSD ephemeral_essd: local SSDs |
| SystemDiskDescription | String | No | Yes | The description of the ECS instance system disk. | None. |
| SystemDiskDiskName | String | No | Yes | The name of the ECS instance system disk. | None. |
| SystemDiskSize | Number | No | Yes | The size of the system disk. | Valid values: 40 to 500. If a custom image is used to create a system disk, make sure that the size of the system disk is greater than that of the custom image. |
| Tags | List | No | Yes | The custom tags of a created instance. | A maximum of 20 tags are supported. The format is as follows: [{"Key": "tagKey", "Value": "tagValue"}, {"Key": "tagKey2", "Value": "tagValue2"}] |

| Attribute | Type | Required | Editable | Description | Constraint |
|--------------------|--------|----------|----------|--|---|
| UserData | String | No | Yes | The user data that you provide when you create ECS instances. | The user data can be up to 16 KB in size. You do not need to use Base64 for transcoding. Special characters need to be escaped. |
| ZoneId | String | No | No | The zone ID of the disk. | None. |
| VpcId | String | No | No | The ID of the VPC. | None. |
| VSwitchId | String | No | No | The ID of the VSwitch for the ECS instance. | None. |
| KeyPairName | String | No | Yes | The name of the key pair that is used to connect to created ECS instances. | For Windows-based ECS instances, this parameter is ignored, and it is empty by default. For Linux-based ECS instances, the Password parameter still takes effect if this parameter is specified. However, logon by password is disabled, and the KeyPairName value is used. |
| RamRoleName | String | No | Yes | The name of the instance RAM role. | You can call the ListRoles operation to query the role name. |
| DedicatedHostId | String | No | No | The ID of the dedicated host. | None. |
| LaunchTemplateName | String | No | Yes | The name of the launch template for the instance. | None. |
| EniMappings | List | No | Yes | The elastic network interfaces (ENIs) to be attached to created instances. | Only one ENI can be attached to each instance. |
| LaunchTemplateId | String | No | Yes | The ID of the launch template. | None. |

| Attribute | Type | Required | Editable | Description | Constraint |
|-----------------------|---------|----------|----------|--|--|
| LaunchTemplateVersion | String | No | Yes | The version of the launch template. | If you do not specify a version, the default version is used. |
| NetworkType | String | No | No | The network type of created ECS instances. | Valid values: <ul style="list-style-type: none"> vpc classic Default value: classic. |
| DeletionProtection | Boolean | No | No | The release protection attribute of the instance. It specifies whether you can use the ECS console or call the DeleteInstance operation to release the instance. | Valid values: <ul style="list-style-type: none"> true false |
| DeploymentSetId | String | No | Yes | Deployment Set ID. | None. |

DiskMappings syntax

```
"DiskMappings": [
{
  "Category": String,
  "DiskName": String,
  "Description": String,
  "Device": String,
  "SnapshotId": String,
  "Size": String,
  "Encrypted": String,
  "KMSKeyId": String,
  "PerformanceLevel": String,
  "AutoSnapshotPolicyId": String
}
]
```

DiskMappings properties

| Attribute | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|--------------------------|------------|
| Size | String | Yes | No | The size of a data disk. | Unit: GB. |

| Attribute | Type | Required | Editable | Description | Constraint |
|----------------------|---------|----------|----------|---|--|
| Category | String | No | No | The type of the data disk. | Valid values: <ul style="list-style-type: none"> cloud cloud_efficiency cloud_ssd cloud_essd ephemeral_ssd For I/O optimized instances, the default value is cloud_efficiency. For non-I/O optimized instances, the default value is cloud. |
| DiskName | String | No | No | The name of data disk N. | The name can be up to 128 characters in length. It can contain English letters, Chinese characters, digits, underscores (_), periods (.), and hyphens (-). |
| Description | String | No | No | The description of data disk N. | The description must be 2 to 256 characters in length. The description cannot start with <code>http://</code> or <code>https://</code> . |
| Device | String | No | No | The device name of the data disk. | The system allocates a device name in alphabetical order from <code>/dev/xvda</code> to <code>/dev/xvdz</code> . |
| SnapshotId | String | No | No | The ID of the snapshot. | None. |
| Encrypted | Boolean | No | No | Specifies whether to encrypt the data disk. | Valid values: <ul style="list-style-type: none"> true false Default value: false |
| KMSKeyId | String | No | No | The ID of the KMS key corresponding to the data disk. | None. |
| AutoSnapshotPolicyId | String | No | Yes | The ID of the automatic snapshot policy. | None. |

| Attribute | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|--|--|
| PerformanceLevel | String | No | No | The performance level of the enhanced SSD used as the data disk. | <ul style="list-style-type: none"> (Default): Maximum random read/write IOPS of 50,000 per disk PL2: A single enhanced SSD delivers up to 100,000 random read/write IOPS. PL3: A single enhanced SSD delivers up to 1,000,000 random read/write IOPS. |

Tags syntax

```
"Tags": [
  {
    "Value": String,
    "Key": String
  }
]
```

Tags properties

| Attribute | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|---------------------|---|
| Key | String | Yes | No | The key of tag N. | It must be 1 to 128 characters in length, and cannot start with <code>aliyun</code> and <code>acs:</code> beginning. cannot contain <code>http://</code> or <code>https://</code> . |
| Value | String | No | No | The value of tag N. | It must be 0 to 128 characters in length and cannot start with <code>aliyun</code> and <code>acs:</code> beginning. cannot contain <code>http://</code> or <code>https://</code> . |

EniMappings syntax

```
"EniMappings": [
  {
    "SecurityGroup": String,
    "VSwitch": String,
    "Description": String,
    "NetworkInterfaceName": String,
    "PrimaryIpAddress": String
  }
]
```

EniMappings properties

| Attribute | Type | Required | Editable | Description | Constraint |
|----------------------|--------|----------|----------|--|--|
| SecurityGroupId | String | Yes | Yes | The ID of the security group to which an instance belongs. | The security group and the instance must be in the same VPC. |
| VSwitchId | String | Yes | No | The ID of the VSwitch to which the instance is connected. | None. |
| Description | String | No | Yes | The description of the ENI. | It can contain 2 to 256 English letters or Chinese character. It cannot start with <code>http://</code> or <code>https://</code> the beginning. |
| NetworkInterfaceName | String | No | Yes | The ENI name. | <ul style="list-style-type: none"> The name must be 2 to 128 characters in length Must start with an uppercase or lowercase letter, and cannot start with <code>http://</code> or <code>https://</code> the beginning. It must start with a letter and cannot start with <code>http://</code> or <code>https://</code>. |
| PrimaryIpAddress | String | No | No | The primary private IP address of ENI. | The specified IP address must be available within the CIDR block of the VSwitch. If this parameter is not specified, an available IP address in the VSwitch CIDR block will be selected at random. |

Return value

Fn::GetAtt

- `InstanceIds`: the IDs of created instances in the ECS instance group. An ID is a system-generated globally unique identifier (GUID) for an instance.
- `PrivateIps`: the list of private IP addresses of instances in a VPC. This parameter takes effect only when the `NetworkType` parameter is set to VPC. For example, a json-formatted Array: `["172.16.XX.XX", "172.16.XX.XX", … "172.16.XX.XX"]` the maximum number of IP addresses that can be specified. Separate multiple IP addresses with commas (,).
- `InnerIps`: the list of private IP addresses of instances in a classic network. This parameter takes effect only when the `NetworkType` parameter is set to Classic. For example, a json-formatted Array: `["10.1.XX.XX", "10.1.XX.XX", … "10.1.XX.XX"]`

ip; "10.1.XX.XX"] the maximum number of IP addresses that can be specified. Separate multiple IP addresses with commas (,).

- **PublicIps**: the list of public IP addresses of instances in a classic network. This parameter takes effect only when the **NetworkType** parameter is set to **Classic**. For example, a json-formatted Array: ["42.1.XX.XX", "42.1.XX.XX",… "42.1.XX.XX"] the maximum number of IP addresses that can be specified. Separate multiple IP addresses with commas (,).
- **HostNames**: the list of hostnames of all instances.
- **OrderId**: the list of order IDs of all instances.
- **ZoneIds**: the IDs of the zones where created instances reside.
- **RelatedOrderIds**: the list of related order IDs of created ECS instances.

Examples

```
{
  "ROSTemplateFormatVersion":"2015-09-01",
  "Resources":{
    "WebServer":{
      "Type": "ALIYUN::ECS::InstanceGroup",
      "Properties": {
        "ImageId": "m-25l0r****",
        "InstanceType": "ecs.t1.small",
        "SecurityGroupId": "sg-25zwc****",
        "ZoneId": "cn-beijing-b",
        "MaxAmount":1,
        "MinAmount":1,
        "Tags": [{
          "Key": "tiantt",
          "Value": "ros"
        }],
        "Key": "tiantt1",
        "Value": "ros1"
      }
    }
  },
  "Outputs":{
    "InstanceIds": {
      "Value":{"get_attr": ["WebServer", "InstanceIds"]}
    },
    "PublicIps": {
      "Value":{"get_attr": ["WebServer", "PublicIps"]}
    }
  }
}
```

5.5.1.12. ALIYUN::ECS::InstanceGroupClone

ALIYUN::ECS::InstanceGroupClone is used to clone an ECS instance group.

Statement

```
{
  "Type": "ALIYUN::ECS::InstanceGroupClone",
  "Properties": {
    "BackendServerWeight": Integer,
    "DiskMappings": List,
    "LaunchTemplateName": String,
    "SpotPriceLimit": String,
    "ResourceGroupId": String,
    "KeyPairName": String,
    "SystemDiskDiskName": String,
    "PeriodUnit": String,
    "Description": String,
    "Tags": List,
    "ImageId": String,
    "SpotStrategy": String,
    "SourceInstanceid": String,
    "EniMappings": List,
    "Password": String,
    "MaxAmount": Integer,
    "AutoReleaseTime": String,
    "SystemDiskCategory": String,
    "LoadBalancerIdToAttach": String,
    "LaunchTemplateId": String,
    "LaunchTemplateVersion": String,
    "ZoneId": String,
    "InstanceName": String,
    "InternetMaxBandwidthIn": Integer,
    "DeletionProtection": Boolean,
    "DeploymentSetId": String,
    "SecurityGroupId": String,
    "RamRoleName": String,
    "HpcClusterId": String,
    "SystemDiskDescription": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|--|--|
| ResourceGroupId | String | Yes | Released | The ID of the enterprise resource group to which a created instance belongs. | None |
| HpcClusterId | String | Yes | True | The ID of the E-HPC cluster to which a created instance belongs. | None |
| SourceInstanceid | String | No | No | The ID of the ECS instance to be cloned. | The clone operation clones the specified instance, including its instance type, image, bandwidth limit, and network type. If the source ECS instance belongs to multiple security groups, the cloned instance is added only to the first of these security groups. |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|------|----------|----------|-------------|------------|
|-----------|------|----------|----------|-------------|------------|

| | | | | | |
|------------------------|---------|----------|----------|---|---|
| MaxAmount | Integer | Retained | Yes | The maximum number of ECS instances to be created. | Valid values: 1 to 100. The MaxAmount parameter must be set to a value greater than or equal to the value of the MinAmount parameter. |
| MinAmount | String | No | Yes | The minimum number of ECS instances to be created. | Valid values: 1 to 100. The MinAmount parameter must be set to a value less than or equal to the value of the MaxAmount parameter. |
| BackendServerWeight | String | Optional | Released | The weight assigned to the ECS instance in the Server Load Balancer instance. | Valid values: 0 to 100. Default value: 100. |
| LoadBalancerIdToAttach | String | Yes | Released | The ID of the SLB instance to which the ECS instance is to be attached. | None |
| Description | String | Yes | Released | The description of created instances. | The description can be up to 256 characters in length. |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|---|--|
| ImageId | String | Yes | True | The ID of the image used to start created instances. You can use a public image, a custom image, or an Alibaba Cloud Marketplace image. | <p>You can specify a partial public image ID instead of providing the complete ID. When editing a template used to deploy an ECS instance, you can specify the image type and version or only the image type. ROS automatically selects an appropriate public image ID. You can use the wildcard character (*) to represent part of an image ID.</p> <p>You can use one of the following methods to specify the public image ID for the ECS instance:</p> <ul style="list-style-type: none"> If you set the parameter to ubuntu, ubuntu_16_0402_64_20G_alibase_20170818.vhd. If this parameter is set to ubuntu_14, ubuntu_14_0405_64_20G_alibase_20170824.vhd is returned. Specify: ubuntu1432, which will eventually match: ubuntu_14_0405_32_40G_alibase_20170711.vhd Specify: ubuntu_16_0402_32, which will eventually match: ubuntu_16_0402_32_40G_alibase_20170711.vhd |
| SecurityGroupId | String | Yes | Released | The ID of the security group to which created instances belong. | None |
| InstanceName | String | Yes | Released | The name of a created instance. | The name can be up to 128 characters in length and can contain letters, digits, underscores (_), periods (.), and hyphens (-). |
| Password | String | Yes | Released | The password used to log on to the ECS instance. | The password must be 8 to 30 characters in length and must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special character. Special characters include () ` ~ ! @ # \$ % ^ & * - + = { } [] ; ' < > , . ? / If the Password parameter is specified, you must use HTTPS to call the operation to ensure that the Password remains confidential. |
| DiskMappings | List | Erased | Released | The data disks to be attached to the instance. | A maximum of 16 disks can be attached to each instance. |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------|--------|----------|----------|---|---|
| Tags | List | Erased | Released | The custom tags of the instance. | You can specify a maximum of 20 tags. The format is as follows: <pre>[{"Key": "tagKey", "Value": "tagValue"}, {"Key": "tagKey2", "Value": "tagValue2"}]</pre> |
| ZoneId | String | Yes | Released | The ID of the zone where the instance resides. | None |
| KeyPairName | String | Yes | Released | The name of the key pair that is used to connect to the ECS instance. For Windows-based ECS instances, this parameter is ignored. Default value: empty. For Linux-based instances, the Password parameter still takes effect if this parameter is specified. However, logon by Password is disabled, and the KeyPairName value is used. | None |
| RamRoleName | String | Yes | Released | The RAM role name of the instance. | You can use RAM API ListRoles you can call this operation to query the RAM role name of an instance. |
| SpotPriceLimit | String | Yes | Released | The maximum hourly price of the instance. | This parameter supports up to three decimal places. Parameter SpotStrategy this parameter takes effect only when the value is SpotWithPriceLimit. |
| SystemDiskDiskName | String | Yes | True | The name of the system disk of created instances. | The name must be 2 to 128 characters in length Must start with an uppercase or lowercase letter. and cannot start with <code>http://</code> or <code>https://</code> and can contain digits, colons (:), underscores (_), and hyphens (-). |
| PeriodUnit | String | Yes | True | The billing cycle for created ECS instances. | Valid values: <ul style="list-style-type: none"> Week. 1, 2, 3, and 4} when the value of the PeriodUnit parameter is Week. AutoRenewPeriod values are 1, 2, "3". Month 1, 2, 3, 4, 5, 6, and 7 when the PeriodUnit parameter is set to Month "8", "9", "12", "24", "36", "48", "60"}, AutoRenewPeriod can be {"1", "2", "3", "6", "12"}. Default value: Month. |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|---|---|
| EniMappings | List | No. | True | The elastic network interfaces (ENIs) to be attached to a created instance. | Only a single ENI can be attached to each instance. |
| AutoReleaseTime | String | Yes | Released | <p>The time scheduled for a created ECS instance to be automatically released. Specify the time in the ISO 8601 standard in the YYYY-MM-DDThh:mmZ format. in the yyyy-MM-ddTHH:mm:ssZ format. The time must be in UTC.</p> <ul style="list-style-type: none"> • If the value of seconds (ss) is a value other than 00, the start time is automatically rounded down to the nearest minute based on the value of mm. • The minimum release time must be at least 30 minutes later than the current time. • The maximum release time must be at most three years from the current time. <p>If you do not specify the AutoReleaseTime it indicates that the auto release feature is disabled and the ECS instance will not be automatically released.</p> | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|------------------------|---------|----------|----------|--|---|
| SystemDiskCategory | String | Yes | True | The type of the system disk. | <p>Valid values:</p> <ul style="list-style-type: none"> cloud: basic disk cloud_efficiency: indicates an ultra disk. cloud_ssd: indicates a standard SSD. ephemeral_ssd: local SSD. cloud_essd: indicates an enhanced SSD (ESSD). ESSDs are still in public preview and only available in some regions. <p>For phased-out instance types that are not I/O optimized, the default value is cloud. For other instances, the default value is cloud_efficiency.</p> |
| LaunchTemplateName | String | Yes | True | The name of the launch template. | None |
| LaunchTemplateVersion | String | Yes | True | The version of the launch template. If you do not specify this parameter, the default version is used. | None |
| InternetMaxBandwidthIn | String | Optional | Released | The maximum inbound bandwidth from the Internet. Unit: Mbit/s. | <p>Valid values: 1 to 200.</p> <p>Default value: 200.</p> |
| LaunchTemplateId | String | Yes | True | The ID of the launch template. | None |
| SystemDiskDescription | String | Yes | True | The description of the system disk. | The description must be 2 to 256 characters in length and cannot start with http:// or https://. This parameter is empty by default. |
| DeletionProtection | Boolean | Erased | Released | The release protection attribute of the instance. Specifies whether the ECS console or API (DeleteInstance) to release the instance. | <p>Valid values:</p> <ul style="list-style-type: none"> true false |
| DeploymentSetId | String | Yes | True | Deployment Set ID. | None |

DiskMappings syntax

```
"DiskMappings": [
  {
    "Category": String,
    "DiskName": String,
    "Description": String,
    "Encrypted": String,
    "KMSKeyId": String,
    "Device": String,
    "SnapshotId": String,
    "Size": String
  }
]
```

DiskMappings properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------|--------|----------|----------|---|--|
| Encrypted | String | Yes | Released | Specifies whether to encrypt the data disk. | Valid values: <ul style="list-style-type: none"> true false Default value: false |
| KMSKeyId | String | Yes | Released | The KMS key ID for data disk N. | None |
| Size | String | No | No | The size of data disk N. Unit: GB. | None |
| Category | String | Yes | Released | The type of the data disk. | Valid values: <ul style="list-style-type: none"> cloud cloud_efficiency cloud_essd ephemeral_essdDefault |
| DiskName | String | Yes | Released | The name of data disk N. | The name can be up to 128 characters in length and can contain letters, digits, underscores (_), periods (.), and hyphens (-). |
| Description | String | Yes | Released | The description of data disk N. | Valid values: 2 to 256. Default value: Null. |
| Device | String | Yes | Released | The device name of the data disk attached to an ECS instance. | The system allocates a device name in alphabetical order from /dev/xvda to /dev/xvdz. |
| SnapshotId | String | Yes | Released | Create a data disk by using a snapshot. | None |

EniMappings syntax

```
"EniMappings": [
  {
    "SecurityGroupId": String,
    "VSwitchId": String,
    "Description": String,
    "NetworkInterfaceName": String,
    "PrimaryIpAddress": String
  }
]
```

EniMappings properties

| Parameter | Type | Required | Editable | Description | Constraint |
|----------------------|--------|----------|----------|--|--|
| SecurityGroupId | String | No | Yes | The ID of the security group to which the ENI belongs. | None |
| VSwitchId | String | No | No | The ID of the VSwitch to which the ENI belongs. | None |
| Description | String | Yes | True | The description of the ENI. | It can contain 2 to 256 English letters or Chinese character. It cannot start with <code>http://</code> and <code>https://</code> the beginning. |
| NetworkInterfaceName | String | Yes | True | The name of the ENI. | None |
| PrimaryIpAddress | String | Yes | Released | The primary IP address of the ENI. | None |

Tags syntax

```
"Tags": [
  {
    "Value": String,
    "Key": String
  }
]
```

Tags properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|-------------|------------|
| Key | String | No | No | None | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|-------------|---|
| Value | String | Yes | Released | None | Valid values: 1025 to 10000. You cannot use the following commonly used port numbers: 2222, 4500, 4510, 4560, 7505, 9000, 9001, and 9002. |

Response parameters

Fn::GetAtt

- **InstanceIds:** the IDs of instances in the ECS instance group. An ID is a globally unique identifier (GUID) generated by the system for an instance.
- **PrivateIps:** the private IP addresses of VPC-type instances. This parameter is valid only when the NetworkType parameter is set to VPC. The parameter value is a JSON-formatted array, containing up to 100 IP addresses separated by commas (.). Example: ["172.16.XX.XX", "172.16.XX.XX", ... "172.16.XX.XX"].
- **InnerIps:** the private IP addresses of instances in the classic network. This parameter is valid only when the NetworkType parameter is set to Classic. The parameter value is a JSON-formatted array, containing up to 100 IP addresses separated by commas (.). Example: ["10.1.XX.XX", "10.1.XX.XX", ... "10.1.XX.XX"].
- **PublicIps:** the public IP addresses of instances in the classic network. This parameter is applicable only when the NetworkType parameter is set is Classic. The parameter value is a JSON-formatted array, containing up to 100 IP addresses separated by commas (.). Example: ["42.1.XX.XX", "42.1.XX.XX", ... "42.1.XX.XX"].
- **HostNames:** the host names of all instances. The parameter value is a JSON-formatted array. Example: ["host1", "host2", ... "host3"].
- **OrderId:** the order IDs of all instances.
- **ZoneIds:** the IDs of the zones where created instances reside.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "WebServer": {
      "Type": "ALIYUN::ECS::InstanceGroupClone",
      "Properties": {
        "SourceInstanceId": "i-25zsk****",
        "ImageId": "m-25l0r****",
        "SecurityGroupId": "sg-25zwc****",
        "ZoneId": "cn-beijing-b",
        "MaxAmount": 1,
        "MinAmount": 1
      }
    }
  },
  "Outputs": {
    "InstanceIds": {
      "Value": {"get_attr": ["WebServer", "InstanceIds"]}
    },
    "PublicIps": {
      "Value": {"get_attr": ["WebServer", "PublicIps"]}
    }
  }
}
```

5.5.1.13. ALIYUN::ECS::Invocation

ALIYUN::ECS::Invocation is used to invoke a Cloud Assistant command for one or more ECS instances.

Statement

```
{
  "Type": "ALIYUN::ECS::Invocation",
  "Properties": {
    "Timed": Boolean,
    "Frequency": String,
    "CommandId": String,
    "InstanceIds": List
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|---------|----------|----------|--|------------|
| Timed | Boolean | Erased | Released | Specifies whether to invoke the command on a periodic basis. Default value: false. | None |
| Frequency | String | Yes | Released | The frequency at which the command is invoked. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------|--------|----------|----------|---|---|
| CommandId | String | No | No | The ID of the script. | None |
| InstanceIds | List | Yes | No | The IDs of the instances for which you want to invoke the command. A maximum of 20 instance IDs can be specified. | Valid values: 1025 to 10000. You cannot use the following commonly used port numbers: 2222, 4500, 4510, 4560, 7505, 9000, 9001, and 9002. |

Response parameters

Fn::GetAtt

The execution ID of the Invokeld: command.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "Timed": {
      "Type": "Boolean",
      "Description": "Whether it is timed execution. Default is False.",
      "AllowedValues": [
        "True",
        "true",
        "False",
        "false"
      ]
    },
    "Frequency": {
      "Type": "String",
      "Description": "The frequency of timing execution (the shortest frequency is performed every 1 minute). It is mandatory when Timing is True.The value rule follows the rules of the cron expression. "
    },
    "CommandId": {
      "Type": "String",
      "Description": "The id of command."
    },
    "InstanceIds": {
      "Type": "CommaDelimitedList",
      "Description": "The instance id list. Select up to 20 instances at a time.Instances selected network type must be VPC network, status must be running"
    }
  },
  "Resources": {
    "Invocation": {
      "Type": "ALIYUN::ECS::Invocation",
      "Properties": {
        "Timed": {
          "Ref": "Timed"
        },
        "Frequency": {
          "Ref": "Frequency"
        }
      }
    }
  }
}
```

```

    "Ref": "Frequency"
  },
  "CommandId": {
    "Ref": "CommandId"
  },
  "InstanceIds": {
    "Fn::Split": [
      ",",
      {
        "Ref": "InstanceIds"
      },
      {
        "Ref": "InstanceIds"
      }
    ]
  }
}
},
"Outputs": {
  "Invokeld": {
    "Description": "The id of command execution.",
    "Value": {
      "Fn::GetAtt": [
        "Invocation",
        "Invokeld"
      ]
    }
  }
}
}
}
}
}

```

5.5.1.14. ALIYUN::ECS::JoinSecurityGroup

ALIYUN::ECS::JoinSecurityGroup is used to add one or more ECS instances to a specified security group.

Syntax

```

{
  "Type": "ALIYUN::ECS::JoinSecurityGroup",
  "Properties": {
    "InstanceId": String,
    "InstanceList": List,
    "SecurityGroupId": String,
    "NetworkInterfaceList": List
  }
}

```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|-------------------------------|------------|
| SecurityGroupId | String | Yes | No | The ID of the security group. | None |

| Property | Type | Required | Editable | Description | Constraint |
|----------------------|--------|----------|----------|---|------------|
| InstanceId | String | No | No | The ID of the ECS instance to be added to the security group. | None |
| InstanceIdList | List | No | Yes | The IDs of the ECS instances to be added to the security group. | None |
| NetworkInterfaceList | List | No | Yes | The IDs of the elastic network interfaces (ENIs). | None |

Response parameters

Fn::GetAtt

None

Examples

JSON format

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "SG": {
      "Type": "ALIYUN::ECS::JoinSecurityGroup",
      "Properties": {
        "SecurityGroupId": "sg-m5eagh7rzys2z8sa****",
        "InstanceIdList": [
          "i-m5e505h9bgsio0wy****",
          "i-m5e505hio0wyjc6r****"
        ]
      }
    }
  }
}
```

YAML format

```
ROSTemplateFormatVersion: '2015-09-01'
Resources:
  SG:
    Type: ALIYUN::ECS::JoinSecurityGroup
    Properties:
      SecurityGroupId: sg-m5eagh7rzys2z8sa****
      InstanceIdList:
        - i-m5e505h9bgsio0wy****
        - i-m5e505hio0wyjc6r****
```

5.5.1.15. ALIYUN::ECS::LaunchTemplate

ALIYUN::ECS::LaunchTemplate is used to create an ECS instance launch template.

Syntax

```
{
  "Type": "ALIYUN::ECS::LaunchTemplate",
  "Properties": {
    "LaunchTemplateName": String,
    "VersionDescription": String,
    "ImageId": String,
    "InstanceType": String,
    "SecurityGroupId": String,
    "NetworkType": String,
    "VSwitchId": String,
    "InstanceName": String,
    "Description": String,
    "InternetMaxBandwidthIn": Integer,
    "InternetMaxBandwidthOut": Integer,
    "HostName": String,
    "ZoneId": String,
    "SystemDiskCategory": String,
    "SystemDiskSize": Number,
    "SystemDiskDiskName": String,
    "SystemDiskDescription": String,
    "IoOptimized": String,
    "InternetChargeType": String,
    "UserData": String,
    "KeyPairName": String,
    "RamRoleName": String,
    "AutoReleaseTime": String,
    "SpotStrategy": String,
    "SpotPriceLimit": String,
    "SecurityEnhancementStrategy": String,
    "DiskMappings": List,
    "NetworkInterfaces": List,
    "Tags": List,
    "TemplateTags": List
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|--------------------|--------|----------|----------|---|--|
| LaunchTemplateName | String | Yes | No | The name of the instance launch template. | <ul style="list-style-type: none"> The name must be 2 to 128 characters in length. It must start with a letter but cannot start with <code>http://</code> or <code>https://</code>. It can contain letters, digits, colons (:), underscores (_), and hyphens (-). |

| Property | Type | Required | Editable | Description | Constraint |
|-------------------------|---------|----------|----------|--|--|
| VersionDescription | String | No | No | The description of the version of the instance launch template. | <ul style="list-style-type: none"> The description must be 2 to 128 characters in length. It must start with a letter but cannot start with <code>http://</code> or <code>https://</code>. |
| ImageId | String | No | No | The ID of the image. | None |
| InstanceType | String | No | No | The type of the instance. | None |
| SecurityGroupId | String | No | No | The ID of the security group. | None |
| NetworkType | String | No | No | The network type of the instance. | Valid values: <ul style="list-style-type: none"> classic vpc |
| VSwitchId | String | No | No | The ID of the VSwitch. You must specify this parameter when you create an instance in a VPC. | None |
| InstanceName | String | No | No | The name of the instance. | <ul style="list-style-type: none"> The name must be 2 to 128 characters in length. It must start with a letter but cannot start with <code>http://</code> or <code>https://</code>. |
| Description | String | No | No | The description of the instance. | <ul style="list-style-type: none"> The description must be 2 to 128 characters in length. It must start with a letter but cannot start with <code>http://</code> or <code>https://</code>. |
| InternetMaxBandwidthIn | Integer | No | No | Maximum inbound bandwidth from the Internet. | Valid values: 1 to 200. Unit: Mbit/s. |
| InternetMaxBandwidthOut | Integer | No | No | Maximum outbound bandwidth to the Internet. | Valid values: 0 to 100. Unit: Mbit/s. |

| Property | Type | Required | Editable | Description | Constraint |
|--------------------|--------|----------|----------|--|--|
| HostName | String | No | No | The hostname of the instance. | <p>The name cannot start or end with a period (.) or a hyphen (-). It cannot contain consecutive periods (..) or hyphens (-).</p> <ul style="list-style-type: none"> For Windows instances: <ul style="list-style-type: none"> The name must be 2 to 15 characters in length and can contain letters, digits, and hyphens (-). It cannot only contain digits. For other instances such as Linux instances: <ul style="list-style-type: none"> The name must be 2 to 64 characters in length and can contain letters, digits, and hyphens (-). |
| ZoneId | String | No | No | The ID of the zone where the instance resides. | None |
| SystemDiskCategory | String | No | No | The category of the system disk. | <p>Valid values:</p> <ul style="list-style-type: none"> cloud: the basic disk cloud_efficiency: the ultra disk cloud_ssd: the standard SSD ephemeral_ssd: the local SSD |
| SystemDiskSize | Number | No | No | The size of the system disk. | <p>Valid values: 20 to 500. Unit: GiB.</p> |
| SystemDiskDiskName | String | No | No | The name of the system disk. | <ul style="list-style-type: none"> The name must be 2 to 128 characters in length and can contain letters, digits, colons (:), underscores (_), and hyphens (-). It must start with a letter but cannot start with <code>http://</code> or <code>https://</code>. |

| Property | Type | Required | Editable | Description | Constraint |
|-----------------------|--------|----------|----------|---|---|
| SystemDiskDescription | String | No | No | The description of the system disk. | The description must be 2 to 256 characters in length and cannot start with <code>http://</code> or <code>https://</code> . |
| IoOptimized | String | No | No | Specifies whether the instance is I/O optimized. | Valid values: <ul style="list-style-type: none"> • none • optimized |
| InternetChargeType | String | No | No | The billing method for network usage. | Valid values: <ul style="list-style-type: none"> • PayByBandwidth • PayByTraffic |
| UserData | String | No | No | The user data of the instance. | The data must be encoded in Base64. The maximum size of the raw data is 16 KB. |
| KeyPairName | String | No | No | The AccessKey pair name. | <ul style="list-style-type: none"> • For Windows instances, this parameter is ignored and is empty by default. The Password parameter takes effect even if the KeyPairName parameter is specified. • For Linux instances, the username and password authentication method is disabled by default. |
| RamRoleName | String | No | No | The RAM role name of the instance. | None |
| AutoReleaseTime | String | No | No | The time scheduled for the instance to be automatically released. | Specify the time in the ISO 8601 standard in the <code>vvvv-MM-ddTHH:mm:ssZ</code> format. The time must be in UTC. |

| Property | Type | Required | Editable | Description | Constraint |
|-----------------------------|--------|----------|----------|--|--|
| SpotStrategy | String | No | No | The preemption policy for pay-as-you-go instances. | <p>This parameter takes effect only when the InstanceChargeType parameter is set to PostPaid.</p> <p>Valid values:</p> <ul style="list-style-type: none"> NoSpot: The instance is created as a regular pay-as-you-go instance. SpotWithPriceLimit: The instance to be created is a preemptible instance with a user-defined maximum hourly price. SpotAsPriceGo: The instance to be created is a preemptible instance whose price is based on the market price at the time of purchase. |
| SpotPriceLimit | String | No | No | The maximum hourly price of the instance. | A maximum of three decimal places can be specified. |
| SecurityEnhancementStrategy | String | No | No | Specifies whether to enable security hardening. | <p>Valid values:</p> <ul style="list-style-type: none"> Active: enables security hardening. Deactive: disables security hardening. |
| DiskMappings | List | No | No | The list of data disks. | A maximum of 16 data disks can be specified. |
| NetworkInterfaces | List | No | No | The list of elastic network interfaces (ENIs). | A maximum of eight ENIs can be specified. |
| Tags | List | No | No | The tags of the instance, security group, disks, and ENIs. | A maximum of 20 tags can be specified. |
| TemplateTags | List | No | No | The tags of the launch template. | A maximum of 20 tags can be specified. |

DiskMappings syntax

```
"DiskMappings": [
  {
    "Category": String,
    "DiskName": String,
    "Description": String,
    "SnapshotId": String,
    "Size": String,
    "Encrypted": String,
    "DeleteWithInstance": String
  }
]
```

DiskMappings properties

| Property | Type | Required | Editable | Description | Constraint |
|-------------|--------|----------|----------|--|--|
| Category | String | No | No | The category of the data disk. | Valid values: <ul style="list-style-type: none"> cloud: the basic disk cloud_efficiency: the ultra disk cloud_ssd: the standard SSD ephemeral_ssd: the local SSD |
| DiskName | String | No | No | The name of the data disk. | <ul style="list-style-type: none"> The name must be 2 to 128 characters in length. It must start with a letter but cannot start with <code>http://</code> or <code>https://</code>. It can contain letters, digits, colons (:), underscores (_), and hyphens (-). |
| Description | String | No | No | The description of the data disk. | The description must be 2 to 256 characters in length and cannot start with <code>http://</code> or <code>https://</code> . |
| SnapshotId | String | No | No | The ID of the snapshot used to create the data disk. | None |

| Property | Type | Required | Editable | Description | Constraint |
|--------------------|---------|----------|----------|---|---|
| Size | String | No | No | The size of the system disk. | <ul style="list-style-type: none"> Valid values when the Category parameter is set to cloud: 5 to 2000. Valid values when the Category parameter is set to cloud_efficiency: 20 to 32768. Valid values when the Category parameter is set to cloud_ssd: 20 to 32768. Valid values when the Category parameter is set to ephemeral_ssd: 5 to 800. Unit: GiB. |
| Encrypted | Boolean | No | No | Specifies whether to encrypt the data disk. | None |
| DeleteWithInstance | Boolean | No | No | Specifies whether to release the data disk when the instance is released. | None |

NetworkInterfaces syntax

```
"NetworkInterfaces": [
  {
    "PrimaryIpAddress": String,
    "VSwitchId": String,
    "SecurityGroupId": String,
    "NetworkInterfaceName": String,
    "Description": String
  }
]
```

NetworkInterfaces properties

| Property | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|--|------------|
| PrimaryIpAddress | String | No | No | The primary private IP address of the ENI. | None |
| VSwitchId | String | No | No | The ID of the VSwitch to which the ENI belongs. | None |
| SecurityGroupId | String | No | No | The ID of the security group to which the ENI belongs. | None |

| Property | Type | Required | Editable | Description | Constraint |
|----------------------|--------|----------|----------|-----------------------------|---|
| NetworkInterfaceName | String | No | No | The name of the ENI. | None |
| Description | String | No | No | The description of the ENI. | The description must be 2 to 256 characters in length and cannot start with <code>http://</code> or <code>https://</code> . |

Tags syntax

```
"Tags": [
  {
    "Value": String,
    "Key": String
  }
]
```

Tags properties

| Property | Type | Required | Editable | Description | Constraint |
|----------|--------|----------|----------|-------------|------------|
| key | String | Yes | No | None | None |
| value | String | No | No | None | None |

TemplateTags syntax

```
"TemplateTags": [
  {
    "Value": String,
    "Key": String
  }
]
```

TemplateTags properties

| Property | Type | Required | Editable | Description | Constraint |
|----------|--------|----------|----------|-------------|------------|
| key | String | Yes | No | None | None |
| value | String | No | No | None | None |

Response parameters

Fn::GetAtt

- LaunchTemplateId: the ID of the instance launch template.
- LaunchTemplateName: the name of the instance launch template.
- DefaultVersionNumber: the default version number of the instance launch template.
- LatestVersionNumber: the latest version number of the instance launch template.

Examples

```
r
```

```

{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "Template1": {
      "Type": "ALIYUN::ECS::LaunchTemplate",
      "Properties": {
        "LaunchTemplateName": "MyTemplate",
        "VersionDescription": "Launch template with all properties set",
        "ImageId": "m-2ze9uqi7wo61hwep****",
        "InstanceType": "ecs.n4.small",
        "SecurityGroupId": "sg-2ze8yxgempcdsq3i****",
        "NetworkType": "vpc",
        "VSwitchId": "vsw-2zei67xd9nhcqzxec****",
        "InstanceName": "InstanceName",
        "Description": "Description of template",
        "InternetMaxBandwidthIn": 100,
        "InternetMaxBandwidthOut": 200,
        "HostName": "ttinfo",
        "ZoneId": "cn-beijing-a",
        "SystemDiskCategory": "cloud_ssd",
        "SystemDiskSize": "40",
        "SystemDiskDiskName": "TheSystemDiskName",
        "SystemDiskDescription": "The system disk description",
        "IoOptimized": "optimized",
        "InternetChargeType": "PayByBandwidth",
        "UserData": "dGhpcyBpcyBhIHVzZXIlgZGF0YSBleG1h****",
        "KeyPairName": "ThisIsKeyPair",
        "RamRoleName": "ThisIsRamRole",
        "AutoReleaseTime": "2050-10-01T00:00:00Z",
        "SpotStrategy": "SpotWithPriceLimit",
        "SpotPriceLimit": "100.001",
        "SecurityEnhancementStrategy": "Active",
        "DiskMappings": [
          {
            "Category": "cloud_ssd",
            "Size": 40,
            "SnapshotId": "s-2ze1fr2bipove27b****",
            "Encrypted": true,
            "DiskName": "dataDisk1",
            "Description": "I am data disk 1",
            "DeleteWithInstance": true
          },
          {
            "Category": "cloud_efficiency",
            "Size": 20,
            "SnapshotId": "s-2ze4k0w8b33mlsq****",
            "Encrypted": false,
            "DiskName": "dataDisk2",
            "Description": "I am data disk 2",
            "DeleteWithInstance": true
          }
        ],
        "NetworkInterfaces": [
          {
            "PrimaryIpAddress": "10.10.1.1",
            "VSwitchId": "vsw-2zetgeiqlemyok9z5****",
            "SecurityGroupId": "sg-2ze8yxgempcdsq3i****",
            "NetworkInterfaceName": "my-eni1",
            "Description": "My eni 1"
          }
        ]
      }
    }
  }
}

```

```

},
"Tags": [
  {
    "Key": "key1",
    "Value": "value1"
  },
  {
    "Key": "key2",
    "Value": "value2"
  }
],
"TemplateTags": [
  {
    "Key": "templateKey1",
    "Value": "templateValue1"
  },
  {
    "Key": "templateKey2",
    "Value": "templateValue2"
  }
]
}
},
"Outputs": {
  "LaunchTemplateId": {
    "Value": {"Fn::GetAtt": ["Template1", "LaunchTemplateId"]}
  },
  "LaunchTemplateName": {
    "Value": {"Fn::GetAtt": ["Template1", "LaunchTemplateName"]}
  },
  "DefaultVersionNumber": {
    "Value": {"Fn::GetAtt": ["Template1", "DefaultVersionNumber"]}
  },
  "LatestVersionNumber": {
    "Value": {"Fn::GetAtt": ["Template1", "LatestVersionNumber"]}
  }
}
}
}

```

5.5.1.16. ALIYUN::ECS::NatGateway

ALIYUN::ECS::NatGateway is used to create a NAT Gateway for a VPC.

Statement

```

{
  "Type": "ALIYUN::ECS::NatGateway",
  "Properties": {
    "DeletionProtection": Boolean,
    "VpcId": String,
    "Description": String,
    "NatGatewayName": String,
    "VSwitchId": String,
    "DeletionForce": Boolean,
    "Spec": String
  }
}

```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------|---------|----------|----------|--|---|
| VpcId | String | Yes | No | The ID of the VPC that you want to create NAT Gateway. | None |
| VSwitchId | String | Yes | No | The ID of the vSwitch in the specified VPC. | None |
| Description | String | Erased | Released | The description of the NAT Gateway. | The description must be 2 to 256 characters in length. |
| NatGatewayName | String | Erased | Released | The name of the NAT Gateway. | The name must be 2 to 128 characters in length and can contain letters, digits, underscores (_), periods (.), and hyphens (-). It must start with a letter. |
| Spec | String | Erased | Released | The type of the NAT Gateway. | Valid values: Small, Middle, and Large. |
| DeletionProtection | Boolean | Erased | Released | Indicates whether deletion protection is enabled. Default value: false. | None |
| DeletionForce | Boolean | Erased | Released | Specifies whether to forcibly delete SNAT and DNAT entries in the Gateway and unbind EIP from the NAT gateway. Default value: false. | Valid values: 1025 to 10000. You cannot use the following commonly used port numbers: 2222, 4500, 4510, 4560, 7505, 9000, 9001, and 9002. |

Response parameters

Fn::GetAtt

- ForwardTableId: the ID of the port forwarding table.
- The ID of the SNAT tableId:SNat source network address translation table.
- NatGatewayId: the unique ID of the Nat gateway.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "NatGateway": {
      "Type": "ALIYUN::ECS::NatGateway",
      "Properties": {
        "NatGatewayName": "nat_gateway_1",
        "Description": "my nat gateway",
        "VpcId": "vpc-25o8s****",
        "VSwitchId": "vsw-25rc1****",
        "Spec": "Small"
      }
    }
  },
  "Outputs": {
    "NatGatewayId": {
      "Value": {"Fn::GetAttr": ["NatGateway", "NatGatewayId"]}
    },
    "ForwardTableId": {
      "Value": {"Fn::GetAttr": ["NatGateway", "ForwardTableId"]}
    },
    "SNatTableId": {
      "Value": {"Fn::GetAttr": ["NatGateway", "SNatTableId"]}
    }
  }
}
```

5.5.1.17. ALIYUN::ECS::NetworkInterface

ALIYUN::ECS::NetworkInterface is used to create an elastic network interface (ENI).

Statement

```
{
  "Type": "ALIYUN::ECS::NetworkInterface",
  "Properties": {
    "Description": String,
    "SecurityGroupId": String,
    "PrimaryIpAddress": String,
    "ResourceGroupId": String,
    "VSwitchId": String,
    "NetworkInterfaceName": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|---|------------|
| ResourceGroupId | String | Yes | Released | The ID of the resource group to which the instance belongs. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|---|------------|
| SecurityGroupId | String | No | Yes | The ID of the security group to which the instance belongs. The security group and the instance must be in the same VPC. | None |
| VSwitchId | String | No | No | The ID of the VSwitch in the VPC. | None |
| Description | String | Yes | True | The description of the ENI. It can contain 2 to 256 English letters or Chinese character. It cannot start with <code>http://</code> and <code>https://</code> the beginning. This parameter is empty by default. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|----------------------|--------|----------|----------|--|---|
| NetworkInterfaceName | String | Yes | True | The name of the ENI. The name must be 2 to 128 characters in length. Must start with an uppercase or lowercase letter, and cannot start with <code>http://</code> and <code>https://</code> the beginning. It can contain letters, digits, colons (:), underscores (_), and hyphens (-). Default value: null. | None |
| PrimaryIpAddress | String | Yes | Released | The primary private IP address of the ENI. The specified IP address must be available within the CIDR block of the VSwitch. If this parameter is not specified, an available IP address in the VSwitch CIDR block is assigned at random. | Valid values: 1025 to 10000. You cannot use the following commonly used port numbers: 2222, 4500, 4510, 4560, 7505, 9000, 9001, and 9002. |

Response parameters

Fn::GetAtt

- NetworkInterfaceId: the ID of the ENI.
- The MAC address of the MacAddress: Elastic Network Interface.
- The private IP address of the PrivateIpAddress: Elastic Network Interface.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "Description": {
      "Type": "String",
      "Description": "Description of your ENI. It is a string of [2, 256] English or Chinese characters."
    },
    "SecurityGroupId": {
      "Type": "String",
      "Description": "The ID of the security group that the ENI joins. The security group and the ENI must be in a same VPC."
    },
    "VSwitchId": {
      "Type": "String",
      "Description": "VSwitch ID of the specified VPC. Specifies the switch ID for the VPC."
    },
    "NetworkInterfaceName": {
      "Type": "String",
      "Description": "Name of your ENI. It is a string of [2, 128] Chinese or English characters. It must begin with a letter and can contain numbers, underscores (_), colons (:), or hyphens (-)."
    },
    "PrimaryIpAddress": {
      "Type": "String",
      "Description": "The primary private IP address of the ENI. The specified IP address must have the same Host ID as the VSwitch. If no IP addresses are specified, a random network ID is assigned for the ENI."
    }
  },
  "Resources": {
    "EniInstance": {
      "Type": "ALIYUN::ECS::NetworkInterface",
      "Properties": {
        "Description": {
          "Ref": "Description"
        },
        "SecurityGroupId": {
          "Ref": "SecurityGroupId"
        },
        "VSwitchId": {
          "Ref": "VSwitchId"
        },
        "NetworkInterfaceName": {
          "Ref": "NetworkInterfaceName"
        },
        "PrimaryIpAddress": {
          "Ref": "PrimaryIpAddress"
        }
      }
    }
  },
  "Outputs": {
    "NetworkInterfaceId": {
      "Description": "ID of your Network Interface.",
      "Value": {
        "Fn::GetAtt": [
          "EniInstance",
          "NetworkInterfaceId"
        ]
      }
    }
  }
}
```

```

    },
  }
}
    
```

5.5.1.18. ALIYUN::ECS::NetworkInterfaceAttachment

ALIYUN::ECS::NetworkInterfaceAttachment is used to attach an elastic network interface (ENI) to an instance in a VPC.

Statement

```

{
  "Type": "ALIYUN::ECS::NetworkInterfaceAttachment",
  "Properties": {
    "Instanceld": String,
    "NetworkInterfaceId": String
  }
}
    
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------|--------|----------|----------|---|---|
| Instanceld | String | No | No | The ID of the RDS instance. | None |
| NetworkInterfaceId | String | No | No | The IDs of the elastic network interfaces (ENIs). | Valid values: 1025 to 10000. You cannot use the following commonly used port numbers: 2222, 4500, 4510, 4560, 7505, 9000, 9001, and 9002. |

Response parameters

Fn::GetAtt

NetworkInterfaceId: the ID of the ENI.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "Instanceld": {
      "Type": "String",
      "Description": "ECS instance id"
    },
    "NetworkInterfaceId": {
      "Type": "String",
      "Description": "Network interface id"
    }
  },
  "Resources": {
    "EniAttachment": {
      "Type": "ALIYUN::ECS::NetworkInterfaceAttachment",
      "Properties": {
        "Instanceld": {
          "Ref": "Instanceld"
        },
        "NetworkInterfaceId": {
          "Ref": "NetworkInterfaceId"
        }
      }
    }
  },
  "Outputs": {
    "NetworkInterfaceId": {
      "Description": "ID of your Network Interface.",
      "Value": {
        "Fn::GetAtt": [
          "EniAttachment",
          "NetworkInterfaceId"
        ]
      }
    }
  }
}
```

5.5.1.19. ALIYUN::ECS::NetworkInterfacePermission

ALIYUN::ECS::NetworkInterfacePermission is used to grant an account the permission to attach an elastic network interface (ENI) to an instance.

Syntax

```
{
  "Type": "ALIYUN::ECS::NetworkInterfacePermission",
  "Properties": {
    "NetworkInterfaceId": String,
    "AccountId": String,
    "Permission": String
  }
}
```

Properties

| Name | Type | Required | Editable | Description | Validity |
|--------------------|--------|----------|----------|--|----------|
| NetworkInterfaceId | String | Yes | No | The ID of the ENI. | None |
| AccountId | String | Yes | No | The ID of the account. | None |
| Permission | String | Yes | No | The permission granted to the account. | None |

Response parameters

Fn::GetAtt

- NetworkInterfacePermissionId: the ID of the ENI permission.

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "AccountId": {
      "Type": "String",
      "Description": "the account id"
    },
    "Permission": {
      "Type": "String",
      "Description": "the permission",
      "Default": "InstanceAttach"
    },
    "NetworkInterfaceId": {
      "Type": "String",
      "Description": "Network interface id"
    }
  },
  "Resources": {
    "EniPermission": {
      "Type": "ALIYUN::ECS::NetworkInterfacePermission",
      "Properties": {
        "AccountId": {
          "Ref": "AccountId"
        },
        "Permission": {
          "Ref": "Permission"
        },
        "NetworkInterfaceId": {
          "Ref": "NetworkInterfaceId"
        }
      }
    }
  },
  "Outputs": {
    "NetworkInterfacePermissionId": {
      "Description": "the network interface permission id",
      "Value": {
        "Fn::GetAtt": [
          "EniPermission",
          "NetworkInterfacePermissionId"
        ]
      }
    }
  }
}
```

5.5.1.20. ALIYUN::ECS::Route

ALIYUN::ECS::Route is used to create a custom route.

Syntax

```
{
  "Type": "ALIYUN::ECS::Route",
  "Properties": {
    "DestinationCidrBlock": String,
    "RouteTableId": String,
    "NextHopId": String,
    "NextHopType": String,
    "RouteId": String,
    "NextHopList": List
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|----------------------|--------|----------|----------|---|--|
| DestinationCidrBlock | String | Yes | No | The destination Classless Inter-Domain Routing (CIDR) block of the route entry. | None |
| RouteTableId | String | Yes | No | The ID of the route table. | None |
| NextHopId | String | No | No | The ID of the next-hop instance of the route entry. | The route is a non-ECMP route. |
| RouteId | String | Yes | No | The ID of the route. | None |
| NextHopType | String | No | No | The type of the next hop. | Default value: Instance. Valid values: <ul style="list-style-type: none"> Instance Tunnel HaVip RouterInterface |

| Property | Type | Required | Editable | Description | Constraint |
|-------------|------|----------|----------|---|---|
| NextHopList | List | No | No | The list of next hops of the route entry. | <p>You must specify the NextHopType and NextHopId parameters to specify the next hops.</p> <ul style="list-style-type: none"> If you specify the NextHopList parameter, the route is an ECMP route. The list contains two to four next hops of the ECMP route entry. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin: 5px 0;"> <p> Note The NextHopList parameter can be specified only when the route entry belongs to a VRouter. In addition, the next hops must be the router interfaces pointing to the connected VBRs.</p> </div> <ul style="list-style-type: none"> If you do not specify the NextHopList parameter, the route is a non-ECMP route. |

NextHopList syntax

```
"NextHopList": [
  {
    "NextHopId": String,
    "NextHopType": String
  }
]
```

NextHopList properties

| Property | Type | Required | Editable | Description | Constraint |
|-------------|--------|----------|----------|---|---|
| NextHopId | String | Yes | No | The ID of the next-hop instance of the route entry. | None |
| NextHopType | String | No | No | The type of the next hop. | <p>Default value: RouterInterface. Valid values:</p> <ul style="list-style-type: none"> Instance Tunnel HaVip RouterInterface |

Response parameters

Fn::GetAtt

None

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "ECSRoute": {
      "Type": "ALIYUN::ECS::Route",
      "Properties": {
        "RouteId": "vrt-25mz0****",
        "RouteTableId": "vtb-25oud****",
        "DestinationCidrBlock": "172.16.XX.XX/24",
        "NextHopId": "i-25xzy****"
      }
    }
  }
}
```

5.5.1.21. ALIYUN::ECS::SNatEntry

ALIYUN::ECS::SNatEntry is used to configure a NAT Gateway table in a source network address translation.

Statement

```
{
  "Type": "ALIYUN::ECS::SNatEntry",
  "Properties": {
    "SNatTableId": String,
    "SNatIp": String,
    "SourceVSwitchId": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------|--------|----------|----------|---|---|
| SNatTableId | String | Retained | Yes | The ID source network address translation table. | None |
| SNatIp | String | Retained | Yes | The public IP address used to source network address translation. | The public IP address must be NAT Gateway in the bandwidth plan. It cannot exist in both the forwarding table and the SNAT table. |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|--|---|
| SourceVSwitchId | String | Retained | Yes | The ID of the VSwitch that accesses the Internet through the SNAT function of NAT Gateway. | Valid values: 1025 to 10000. You cannot use the following commonly used port numbers: 2222, 4500, 4510, 4560, 7505, 9000, 9001, and 9002. |

Response parameters

Fn::GetAttr

SNatEntryId: the table entry ID in the source network address translation table.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "SNatEntry": {
      "Type": "ALIYUN::ECS::SNatEntry",
      "Properties": {
        "SNatTableId": "stb-3er41****",
        "SourceVSwitchId": "vsw-25rc1****",
        "SNatIp": "101.201.XX.XX"
      }
    }
  },
  "Outputs": {
    "SNatEntryId": {
      "Value": {"Fn::GetAttr": ["SNatEntry", "SNatEntryId"]}
    }
  }
}
```

5.5.1.22. ALIYUN::ECS::SecurityGroup

ALIYUN::ECS::SecurityGroup is used to create a security group.

Statement

```
{
  "Type": "ALIYUN::ECS::SecurityGroup",
  "Properties": {
    "VpcId": String,
    "Description": String,
    "SecurityGroupName": String,
    "Tags": List,
    "SecurityGroupEgress": List,
    "SecurityGroupIngress": List,
    "ResourceGroupId": String,
    "SecurityGroupType": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|----------------------|--------|----------|----------|---|---|
| ResourceGroupId | String | Yes | Released | The ID of the resource group to which the instance belongs. | None |
| VpcId | String | Yes | Released | The ID of the VPC. | None |
| Description | String | Yes | Released | The description of the new security group. | The description must be 2 to 256 characters in length. |
| Tags | List | Erased | Released | The tags of the security group. | A maximum of 20 tags can be specified. |
| SecurityGroupName | String | Yes | Released | The name of the new security group. | Default value: empty. <ul style="list-style-type: none"> The name must be 2 to 128 characters in length Must start with an uppercase or lowercase letter, and cannot start with <code>http://</code> and <code>https://</code> the beginning. It must start with a letter and cannot start with <code>http://</code> or <code>https://</code>. |
| SecurityGroupEgress | List | Erased | Released | The outbound access rules of the security group. | None |
| SecurityGroupIngress | List | Erased | Released | The inbound access rules of the security group. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------------|--------|----------|----------|-------------------------------------|---|
| SecurityGroupType | String | Yes | Released | The type of the new security group. | Valid values: <ul style="list-style-type: none"> normal (basic security group) enterprise (Advanced Security Group) |

Tags syntax

```
"Tags": [
  {
    "Value": String,
    "Key": String
  }
]
```

Tags properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|-------------|------------|
| Key | String | No | No | None | None |
| Value | String | Yes | Released | None | None |

SecurityGroupEgress syntax

```
"SecurityGroupEgress": [
  {
    "Description": String,
    "PortRange": String,
    "SecurityGroupId": String,
    "NicType": String,
    "Priority": Integer,
    "DestGroupId": String,
    "DestCidrIp": String,
    "Policy": String,
    "IpProtocol": String,
    "DestGroupOwnerAccount": String,
    "DestGroupOwnerId": String,
    "Ipv6DestCidrIp": String
  }
]
```

SecurityGroupEgress properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------|--------|----------|----------|---|--|
| Description | String | Yes | Released | The description of the security group rule. | The description must be 1 to 512 characters in length. |

| Parameter | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|--|---|
| DestGroupOwnerId | String | Yes | Released | The ID of the Alibaba Cloud account that owns the destination security group. This parameter is used to grant the current security group access to security groups in another Alibaba Cloud account. | If neither the DestGroupOwnerId parameter nor the DestGroupOwnerIdAccount parameter is specified, the current security group is granted access to other security groups in the same Alibaba Cloud account. If the DestCidrIp parameter is specified, the DestGroupOwnerId parameter is ignored. |
| IpProtocol | String | No | No | The Internet protocol over which the listener will forward requests. | Valid values: <ul style="list-style-type: none"> • TCP • udp • icmp • GRE • All A value of all specifies that all the four protocols are supported. |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|---|---|
| PortRange | String | Yes | Released | The range of port numbers corresponding to the Internet protocol. | <p>The range of destination ports corresponding to the transport layer protocol. Valid values:</p> <ul style="list-style-type: none"> When the IpProtocol parameter is set to tcp or udp, the port number range is 1 to 65535. Separate the starting port and the ending port with a forward slash (/). Correct example: 1/200. Incorrect example: 200/1. When the IpProtocol parameter is set to icmp, the port number range is -1/-1, indicating that all ports are available. When the IpProtocol parameter is set to gre, the port number range is -1/-1, indicating that all ports are available. When the IpProtocol parameter is set to all, the port number range is -1/-1. |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|---|--|
| SecurityGroupId | String | Yes | Released | The ID of the security group for which to create the outbound access rules. | None |
| NicType | String | Yes | Released | The network type of the instance. Valid values: | Valid values: <ul style="list-style-type: none"> internet intranet Default value: internet. |
| Priority | String | Optional | Released | The priority of the authorization policy. | Valid values: 1 to 100. Default value: 1 |
| DestGroupId | String | Yes | Released | The ID of the destination security group within the same region. | You must specify either the DestGroupId parameter or the DestCidrIp parameter. If both parameters are specified, the system authorizes the destination CIDR block specified by the DestCidrIp parameter. If the DestGroupId parameter is specified, but the DestCidrIp parameter is not, you must set the NicType parameter to intranet. |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------------|--------|----------|----------|--|--|
| DestCidrIp | String | Yes | Released | The source IPv4 CIDR block. | The value must be in CIDR format. The default value is 0.0.0.0/0, indicating that access from any IP addresses is allowed. Examples of other supported formats include 10.159.XX.XX/12. Only IPv4 addresses are supported. |
| Policy | String | Yes | Released | The authorization policy. | Valid values: <ul style="list-style-type: none"> accept: grants access drop: denies access Default value: accept. |
| DestGroupOwnerAccount | String | Yes | Released | The Alibaba Cloud account of the destination security group when you grant security group permissions across accounts. | None |
| Ipv6DestCidrIp | String | Yes | Released | The destination IPv6 CIDR block. | IPv6 addresses in CIDR format are supported. You can only specify the IP addresses for ECS instances in VPCs. |

SecurityGroupIngress syntax

```

"SecurityGroupIngress": [
{
  "SourceGroupOwnerId": String,
  "SourceGroupOwnerAccount": String,
  "Description": String,
  "PortRange": String,
  "SecurityGroupId": String,
  "NicType": String,
  "Ipv6SourceCidrIp": String,
  "Priority": Integer,
  "SourceGroupId": String,
  "Policy": String,
  "IpProtocol": String,
  "SourcePortRange": String,
  "SourceCidrIp": String
}
]

```

SecurityGroupIngress properties

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------|--------|----------|----------|--|---|
| SourceGroupOwnerId | String | Yes | Released | The ID of the Alibaba Cloud account that owns the source security group. | None |
| Description | String | Yes | Released | The description of the security group rule. | The description must be 1 to 512 characters in length. |
| IpProtocol | String | No | No | The Internet protocol over which the listener will forward requests. | Valid values: tcp, udp, icmp, gre, and all. A value of all specifies that all the four protocols are supported. |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|---|--|
| PortRange | String | Yes | Released | The range of port numbers corresponding to the Internet protocol. | <p>The range of destination ports corresponding to the transport layer protocol. Valid values:</p> <ul style="list-style-type: none"> When the IpProtocol parameter is set to tcp or udp, the port number range is 1 to 65535. Separate the starting port and the ending port with a forward slash (/). Correct example: 1/200. Incorrect example: 200/1. When the IpProtocol parameter is set to icmp, the port number range is -1/-1, indicating that all ports are available. When the IpProtocol parameter is set to gre, the port number range is -1/-1, indicating that all ports are available. When the IpProtocol parameter is set to all, the port number range is -1/-1, indicating that all ports are available. |

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------------------|--------|----------|----------|--|---|
| SourceGroupId | String | Yes | Released | The ID of the source security group within the same region. | You must specify either the SourceGroupId parameter or the SourceCidrIp parameter. If both parameters are specified, the system authorizes the source CIDR block specified by the SourceCidrIp parameter. If the SourceGroupId parameter is specified, but the SourceCidrIp parameter is not, you must set the NicType parameter to intranet. |
| SecurityGroupId | String | Yes | Released | The ID of the security group for which you want to create the inbound access rule. | None |
| NicType | String | Yes | Released | The network type of the instance. Valid values: | Valid values: <ul style="list-style-type: none"> internet intranet Default value: internet. |
| SourceGroupOwnerAccount | String | Yes | Released | The Alibaba Cloud account of the destination security group when you grant security group permissions across accounts. | None |
| Priority | String | Optional | Released | The priority of the authorization policy. | Valid values: 1 to 100. Default value: 1 |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------|--------|----------|----------|-----------------------------|--|
| SourceCidrIp | String | Yes | Released | The source IPv4 CIDR block. | The value must be in CIDR format. The default value is 0.0.0.0/0, indicating that access from any IP addresses is allowed. Examples of other supported formats include 10.159.XX.XX/12. Only IPv4 CIDR blocks are supported. |
| Policy | String | Yes | Released | The authorization policy. | Valid values: <ul style="list-style-type: none"> • accept: accepts the request. • drop: access is denied. Default value: accept. |

| Parameter | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|--|--|
| SourcePortRange | String | Yes | Released | The range of source ports relevant to transport layer protocols. | <p>Valid values:</p> <ul style="list-style-type: none"> When the IpProtocol parameter is set to tcp or udp, the port number range is 1 to 65535. Separate the starting port and the ending port with a forward slash (/). Correct example: 1/200. Incorrect example: 200/1. When the IpProtocol parameter is set to icmp, the port number range is -1/-1, indicating that all values are valid. When the IpProtocol parameter is set to gre, the port number range is -1/-1, indicating that all ports are available. The IpProtocol value is all:-1/-1. |
| Ipv6SourceCidrIp | String | Yes | Released | The source IPv6 CIDR block. IPv6 addresses in CIDR format are supported. | You can only specify the IP addresses of ECS instances in VPCs. |

Response parameters

Fn::GetAtt

SecurityGroupId: the ID of the new security group.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "SG": {
      "Type": "ALIYUN::ECS::SecurityGroup",
      "Properties": {
        "SecurityGroupName": {
          "Ref": "SecurityGroupName"
        },
        "SecurityGroupIngress": [
          {
            "SourceCidrIp": "0.0.0.0/0",
            "IpProtocol": "all",
            "NicType": "internet",
            "PortRange": "-1/-1",
            "Priority": 1
          },
          {
            "SourceCidrIp": "0.0.0.0/0",
            "IpProtocol": "all",
            "NicType": "intranet",
            "PortRange": "-1/-1",
            "Priority": 1
          }
        ],
        "SecurityGroupEgress": [
          {
            "IpProtocol": "all",
            "DestCidrIp": "0.0.0.0/0",
            "NicType": "internet",
            "PortRange": "-1/-1",
            "Priority": 1
          },
          {
            "IpProtocol": "all",
            "DestCidrIp": "0.0.0.0/0",
            "NicType": "intranet",
            "PortRange": "-1/-1",
            "Priority": 1
          }
        ],
        "VpcId": {
          "Ref": "Vpc"
        }
      }
    }
  },
  "Outputs": {
    "SecurityGroupId": {
      "Value": {"Fn::GetAtt": ["SG", "SecurityGroupId"]}
    }
  }
}
```

5.5.1.23. ALIYUN::ECS::SecurityGroupClone

ALIYUN::ECS::SecurityGroupClone is used to clone a security group.

Syntax

```
{
  "Type": "ALIYUN::ECS::SecurityGroupClone",
  "Properties": {
    "DestinationRegionId": String,
    "VpcId": String,
    "Description": String,
    "SecurityGroupName": String,
    "SourceSecurityGroupId": String,
    "ResourceGroupId": String,
    "NetworkType": String,
    "SecurityGroupType": String
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|-----------------------|--------|----------|----------|---|---|
| ResourceGroupId | String | No | No | The ID of the resource group to which the instance belongs. | None |
| SourceSecurityGroupId | String | Yes | No | The ID of the source security group. | Only applicable security group rules are copied to the new security group. The security group rules are selected based on the network type of the new security group. |
| NetworkType | String | No | No | The network type of the new security group. | Set the value to Classic. |
| VpcId | String | No | No | The ID of the VPC to which the new security group belongs. | The NetworkType parameter is ignored if both the VpcId and NetworkType parameters are specified. |
| Description | String | No | No | The description of the new security group. | The description must be 2 to 256 characters in length. It cannot start with http:// or https://. |

| Property | Type | Required | Editable | Description | Constraint |
|---------------------|--------|----------|----------|--|---|
| SecurityGroupName | String | No | No | The name of the new security group. | This parameter is empty by default. The name must be 2 to 128 characters in length and can contain letters, digits, periods (.), underscores (_), and hyphens (-). It must start with a letter and cannot start with http:// or https://. |
| DestinationRegionId | String | No | No | The ID of the destination region where the new security group resides. | Default value: CURRENT. |
| SecurityGroupType | String | No | No | The type of the new security group. | Valid values: normal and enterprise. A value of normal specifies a basic security group. A value of enterprise specifies an advanced security group. |

Response parameters

Fn::GetAtt

SecurityGroupId: the ID of the new security group.

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "SecurityGroupClone": {
      "Type": "ALIYUN::ECS::SecurityGroupClone",
      "Properties": {
        "SourceSecurityGroupId": {
          "Ref": "SourceSecurityGroupId"
        },
        "VpcId": {
          "Ref": "VpcId"
        }
      },
      "Description": {
        "Ref": "Description"
      },
      "SecurityGroupName": {
```

```

    "Ref": "SecurityGroupName"
  },
  "DestinationRegionId": {
    "Ref": "DestinationRegionId"
  },
  "NetworkType": {
    "Ref": "NetworkType"
  }
}
},
"Parameters": {
  "SourceSecurityGroupId": {
    "Type": "String",
    "Description": "Source security group ID is used to copy properties to clone new security group. If the NetworkType and VpcId is not specified, the same security group will be cloned. If NetworkType or VpcId is specified, only proper security group rules will be cloned."
  },
  "VpcId": {
    "Type": "String",
    "Description": "Physical ID of the VPC."
  },
  "Description": {
    "Type": "String",
    "Description": "Description of the security group, [2, 256] characters. Do not fill or empty, the default is empty."
  },
  "SecurityGroupName": {
    "Type": "String",
    "Description": "Display name of the security group, [2, 128] English or Chinese characters, must start with a letter or Chinese in size, can contain numbers, '_' or '.', '-'"
  },
  "DestinationRegionId": {
    "Default": "CURRENT",
    "Type": "String",
    "Description": "Clone security group to the specified region. Default to current region."
  },
  "NetworkType": {
    "Type": "String",
    "Description": "Clone new security group as classic network type. If the VpcId is specified, the value will be ignored.",
    "AllowedValues": [
      "Classic"
    ]
  }
}
},
"Outputs": {
  "SecurityGroupId": {
    "Description": "Generated security group id of new security group.",
    "Value": {
      "Fn::GetAtt": [
        "SecurityGroupClone",
        "SecurityGroupId"
      ]
    }
  }
}
}
}

```

5.5.1.24. ALIYUN::ECS::SecurityGroupEgress

ALIYUN::ECS::SecurityGroupEgress is used to create an outbound access rule for a security group.

Syntax

```
{
  "Type": "ALIYUN::ECS::SecurityGroupEgress",
  "Properties": {
    "SecurityGroupId": String,
    "IpProtocol": String,
    "PortRange": String,
    "DestGroupId": String,
    "DestGroupOwnerAccount": String,
    "DestCidrIp": String,
    "Policy": String,
    "Priority": String,
    "NicType": String,
    "Ipv6DestCidrIp": String
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|------------|--------|----------|----------|-------------------------------|---|
| IpProtocol | String | Yes | No | The transport layer protocol. | Valid values: <ul style="list-style-type: none"> • tcp • udp • icmp • gre • all <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p> Note The value all indicates that all the four protocols are supported.</p> </div> |

| Property | Type | Required | Editable | Description | Constraint |
|-----------------|---------|----------|----------|--|--|
| PortRange | String | Yes | No | The range of destination port numbers corresponding to the transport layer protocol. | <ul style="list-style-type: none"> Valid values when IpProtocol is set to tcp or udp: 1 to 65535. Separate the starting and ending port numbers with a forward slash (/). Correct example: 1/200. Incorrect example: 200/1. Set the value to -1/-1 when IpProtocol is set to icmp. Set the value to -1/-1 when IpProtocol is set to gre. Set the value to -1/-1 when IpProtocol is set to all. |
| SecurityGroupId | String | No | No | The ID of the source security group. | None |
| NicType | String | No | No | The type of the network interface controller (NIC). | Default value: internet. Valid values: <ul style="list-style-type: none"> internet intranet If the DestGroupId parameter is specified, but the DestCidrIp parameter is not, set the value to intranet. |
| Priority | Integer | No | No | The priority of the security group rule. | Valid values: 1 to 100 Default value: 1. |

| Property | Type | Required | Editable | Description | Constraint |
|-----------------------|--------|----------|----------|---|--|
| DestGroupId | String | No | No | The ID of the destination security group for which you want to set access permissions. | You must specify at least one of the DestGroupId and DestCidrIp parameters. If the DestGroupId parameter is specified, but the DestCidrIp parameter is not, set the NicType value to intranet. If both the DestGroupId and DestCidrIp parameters are specified, the DestCidrIp parameter takes precedence. |
| DestCidrIp | String | No | No | The destination CIDR block. | Only IPv4 CIDR blocks are supported. |
| Policy | String | No | No | The authorization policy. | Default value: accept. Valid values: <ul style="list-style-type: none"> accept: allows access. drop: denies access. |
| DestGroupOwnerAccount | String | No | No | The Alibaba Cloud account that manages the destination security group when you set a security group rule across accounts. | If you specify neither of the DestGroupOwnerId parameter nor the DestGroupOwnerId parameter, the access permission is configured on another security group managed by your account. If you specify the DestCidrIp parameter, the DestGroupOwnerAccount parameter is ignored. |

| Property | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|---|---|
| Description | String | No | Yes | The description of the security group rule. | The description must be 1 to 512 characters in length. |
| DestGroupOwnerId | String | No | No | The ID of the Alibaba Cloud account that manages the destination security group when you set a security group rule across accounts. | If you specify neither of the DestGroupOwnerId parameter nor the DestGroupOwnerId parameter, the access permission is configured on another security group managed by your account. If you specify the DestCidrIp parameter, the DestGroupOwnerId parameter is ignored. |
| Ipv6DestCidrIp | String | No | No | The destination IPv6 CIDR block. | IPv6 addresses in the CIDR format are supported. You can specify only the IP addresses of ECS instances in VPCs. |

Response parameters

Fn::GetAtt

None

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "SG": {
      "Type": "ALIYUN::ECS::SecurityGroupEgress",
      "Properties": {
        "SecurityGroupId": "sg-25bow****",
        "IpProtocol": "tcp",
        "PortRange": "65535/65535",
        "DestCidrIp": "0.0.0.0/0"
      }
    }
  }
}
```

5.5.1.25. ALIYUN::ECS::SecurityGroupIngress

ALIYUN::ECS::SecurityGroupIngress is used to create an inbound access rule for a security group.

Syntax

```
{
  "Type": "ALIYUN::ECS::SecurityGroupIngress",
  "Properties": {
    "SourceGroupOwnerId": String,
    "Description": String,
    "PortRange": String,
    "SecurityGroupId": String,
    "NicType": String,
    "Ipv6SourceCidrIp": String,
    "Priority": Integer,
    "SourceGroupId": String,
    "Policy": String,
    "IpProtocol": String,
    "SourcePortRange": String,
    "SourceCidrIp": String
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|------------|--------|----------|----------|------------------------|---|
| IpProtocol | String | Yes | No | The Internet protocol. | Valid values: tcp, udp, icmp, gre, and all. A value of all specifies that all the four protocols are supported. |

| Property | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|---|--|
| PortRange | String | Yes | No | The range of destination ports relevant to transport layer protocols. | <ul style="list-style-type: none"> When the IpProtocol parameter is set to tcp or udp, the port number range is 1 to 65535. Separate the starting port and the ending port with a forward slash (/). Correct example: 1/200. Incorrect example: 200/1. When the IpProtocol parameter is set to icmp, the port number range is -1/-1. When the IpProtocol parameter is set to gre, the port number range is -1/-1. When the IpProtocol parameter is set to all, the port number range is -1/-1. |

| Property | Type | Required | Editable | Description | Constraint |
|-------------------------|--------|----------|----------|--|--|
| SourceGroupId | String | No | No | The ID of the source security group for which you want to set access permissions. | You must specify at least one of the SourceGroupId and SourceCidrIp parameters. If the SourceGroupId parameter is specified, but the SourceCidrIp parameter is not, the NicType parameter must be set to intranet. If both the SourceGroupId and SourceCidrIp parameters are specified, the SourceCidrIp value is used by default. |
| SecurityGroupId | String | No | No | The ID of the security group for which you want to create the inbound access rule. | None |
| NicType | String | No | No | The network type of the instance. | Valid values: <ul style="list-style-type: none"> internet intranet Default value: internet. |
| SourceGroupOwnerAccount | String | No | No | The Alibaba Cloud account that manages the source security group when you set a security group rule across accounts. | If neither the SourceGroupOwnerAccount parameter nor the SourceGroupOwnerIid parameter is specified, the access permission is configured for another security group managed by your account. If the SourceCidrIp parameter is specified, this parameter is ignored. |

| Property | Type | Required | Editable | Description | Constraint |
|--------------------|---------|----------|----------|--|--|
| Priority | Integer | No | No | The priority of the security group rule. | Valid values: 1 to 100. Default value: 1. |
| SourceCidrIp | String | No | No | The source IPv4 CIDR block. | Only IPv4 CIDR blocks are supported. |
| Policy | String | No | No | The access control policy. | Valid values: <ul style="list-style-type: none"> accept: grants access. drop: denies access. Default value: accept. |
| SourceGroupOwnerId | String | No | No | The ID of the Alibaba Cloud account that manages the source security group when you set a security group rule across accounts. | If neither the SourceGroupOwnerId parameter nor the SourceGroupOwnerAccount parameter is specified, the access permission is configured for another security group managed by your account. If the SourceCidrIp parameter is specified, this parameter is ignored. |
| Description | String | No | Yes | The description of the security group rule. | The description must be 1 to 512 characters in length. |

| Property | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|--|--|
| SourcePortRange | String | No | No | The range of source ports relevant to transport layer protocols. | <ul style="list-style-type: none"> When the IpProtocol parameter is set to tcp or udp, the port number range is 1 to 65535. Separate the starting port and the ending port with a forward slash (/). Correct example: 1/200. Incorrect example: 200/1. When the IpProtocol parameter is set to icmp, the port number range is -1/-1. When the IpProtocol parameter is set to gre, the port number range is -1/-1. When the IpProtocol parameter is set to all, the port number range is -1/-1. |
| Ipv6SourceCidrIp | String | No | No | The range of source IPv6 addresses. | CIDR blocks and IPv6 addresses are supported. You can only specify the IP addresses of ECS instances in VPCs. |

Response parameters

Fn::GetAtt

None

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "SG": {
      "Type": "ALIYUN::ECS::SecurityGroupIngress",
      "Properties": {
        "SecurityGroupId": "sg-25bow****",
        "IpProtocol": "tcp",
        "PortRange": "65535/65535",
        "SourceCidrIp": "0.0.0.0/0"
      }
    }
  }
}
```

5.5.1.26. ALIYUN::ECS::Snapshot

ALIYUN::ECS::Snapshot is used to create a disk Snapshot.

Statement

```
{
  "Type": "ALIYUN::ECS::Snapshot",
  "Properties": {
    "SnapshotName": String,
    "Timeout": Integer,
    "Description": String,
    "DiskId": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|---|------------|
| DiskId | String | No | No | The ID of the disk for which you want to create the snapshot. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------|--------|----------|----------|---|--|
| SnapshotName | String | Yes | Released | The name of the snapshot. | It must be 2 to 128 characters in length. And can contain letters, digits, underscores (_), and hyphens (-). It cannot start with auto. Snapshot names starting with auto are reserved for automatic snapshots. It cannot start with <code>http://</code> or <code>https://</code> . |
| Timeout | String | Optional | Released | The timeout period that is specified for the snapshot creation request. | If this parameter is set, the timeout period to create a resource stack is prolonged. If the snapshot is not created within the specified time period, the entire resource stack fails to be created. You must set the timeout period according to the disk size and data amount. Valid values: 200 to 1440. Unit: minute. The default value is 200 minutes. |
| Description | String | Yes | Released | The description of the snapshot. | The length must be 2 to 256 characters in length. This parameter is empty by default. It cannot start with <code>http://</code> or <code>https://</code> . |

Response parameters

Fn::GetAtt

SnapshotId: the ID of the snapshot.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "Snapshot": {
      "Type": "ALIYUN::ECS::Snapshot",
      "Properties": {
        "DiskId": "d-2zedgvuvu8cylvrd*****"
      }
    }
  },
  "Outputs": {
    "SnapshotId": {
      "Value": {
        "Fn::GetAtt": [
          "Snapshot",
          "SnapshotId"
        ]
      }
    }
  }
}
```

5.5.1.27. ALIYUN::ECS::SSHKeyPair

ALIYUN::ECS::SSHKeyPair is used to create or import an SSH key pair to an ECS instance.

Statement

```
{
  "Type": "ALIYUN::ECS::SSHKeyPair",
  "Properties": {
    "ResourceGroupId": String,
    "KeyPairName": String,
    "PublicKeyBody": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|---|------------|
| ResourceGroupId | String | Yes | Released | The ID of the resource group to which the instance belongs. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|---------------|--------|----------|----------|---|---|
| KeyPairName | String | No | No | The globally unique name of the SSH key pair. | The name must be 2 to 128 characters in length and can contain letters, digits, periods (.), underscores (_), and hyphens (-). It cannot start with <code>http://</code> or <code>https://</code> . |
| PublicKeyBody | String | Yes | Released | Specifies the SSH public key to import. | Valid values: 1025 to 10000. You cannot use the following commonly used port numbers: 2222, 4500, 4510, 4560, 7505, 9000, 9001, and 9002. |

Response parameters

Fn::GetAtt

- **KeyPairFingerPrint**: the fingerprint of the key pair. The message-digest algorithm 5 (MD5) is used based on the public key fingerprint format defined in RFC 4716.
- **PrivateKeyBody**: the private key of the key pair. An unencrypted RSA private key must be encoded using PEM and must be in the PKCS#8 format. The private key of a key pair can only be obtained at the time of its creation. If you import an existing public key, no private key information will be available.
- **KeyPairName**: the globally unique name of the SSH key pair.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "SSHKeyPair": {
      "Type": "ALIYUN::ECS::SSHKeyPair",
      "Properties": {
        "KeyPairName": "ssh_key_pair_v1"
      }
    }
  },
  "Outputs": {
    "KeyPairName": {
      "Value": {
        "Fn::GetAtt": [
          "SSHKeyPair",
          "KeyPairName"
        ]
      }
    },
    "PrivateKeyBody": {
      "Value": {
        "Fn::GetAtt": [
          "SSHKeyPair",
          "PrivateKeyBody"
        ]
      }
    },
    "KeyPairFingerPrint": {
      "Value": {
        "Fn::GetAtt": [
          "SSHKeyPair",
          "KeyPairFingerPrint"
        ]
      }
    }
  }
}
```

5.5.1.28. ALIYUN::ECS::SSHKeyPairAttachment

ALIYUN::ECS::SSHKeyPairAttachment is used to bind an SSH key pair to an ECS instance.

Statement

```
{
  "Type": "ALIYUN::ECS::SSHKeyPairAttachment",
  "Properties": {
    "Instancelds": List,
    "KeyPairName": String
  }
}
```

Properties

| Parameter | Type | Required or Not | Editable | Description | Constraint |
|-------------|--------|-----------------|----------|--|---|
| InstanceIds | List | Retained | Yes | The IDs of the ECS instances with which you want to associate the EIP. | Separate the IDs with a comma (.). Only Linux instances are supported. |
| KeyPairName | String | No | No | The name of the SSH key pair. | Valid values: 1025 to 10000. You cannot use the following commonly used port numbers: 2222, 4500, 4510, 4560, 7505, 9000, 9001, and 9002. |

Response parameters

Fn::GetAtt

None

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "SSHKeyPairAttachment": {
      "Type": "ALIYUN::ECS::SSHKeyPairAttachment",
      "Properties": {
        "KeyPairName": "ssh_key_pair_v1",
        "InstanceIds": [
          'i-2zeiofnh20hj**** has been added * ',
          'i-2zebt3kfvxm2**** has two records **'
        ]
      }
    }
  }
}
```

5.5.1.29. ALIYUN::ECS::VPC

ALIYUN::ECS::VPC is used to create a VPC.

Statement

```
{
  "Type": "ALIYUN::ECS::VPC",
  "Properties": {
    "Description": String,
    "Ipv6CidrBlock": String,
    "EnableIpv6": Boolean,
    "ResourceGroupId": String,
    "VpcName": String,
    "CidrBlock": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|---|--|
| ResourceGroupId | String | Yes | Released | The ID of the resource group to which the instance belongs. | None |
| VpcName | String | Yes | True | The name of the VPC. | <ul style="list-style-type: none"> The name must be 2 to 128 characters in length Must start with english letters or starts with a Chinese character. It cannot start with <code>http://</code> or <code>https://</code>. It must start with a letter and cannot start with <code>http://</code> or <code>https://</code>. |
| CidrBlock | String | Yes | True | The CIDR block of the VPC. | Valid values: 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16. |
| Description | String | Yes | True | The description of the VPC. | The description must be 2 to 256 characters in length. It cannot start with <code>http://</code> or <code>https://</code> . |

| Parameter | Type | Required | Editable | Description | Constraint |
|---------------|---------|----------|----------|---|--|
| Ipv6CidrBlock | String | Yes | Released | The IPv6 CIDR block of the VPC. | None |
| EnableIpv6 | Boolean | No. | True | Specifies whether to enable an IPv6 CIDR block. | Valid values: <ul style="list-style-type: none"> • true • false Default value: false. |

Response parameters

Fn::GetAtt

- VpcId: The VPC ID allocated by the system.
- VRouterId: the ID of the vRouter.
- RouteTableId: the ID of the routing table.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "EcsVpc": {
      "Type": "ALIYUN::ECS::VPC",
      "Properties": {
        "CidrBlock": "172.16.0.0/12",
        "VpcName": "vpc-test-del"
      }
    }
  },
  "Outputs": {
    "VpcId": {
      "Value": {
        "Fn::GetAtt": [
          "EcsVpc",
          "VpcId"
        ]
      }
    },
    "VRouterId": {
      "Value": {
        "Fn::GetAtt": [
          "EcsVpc",
          "VRouterId"
        ]
      }
    },
    "RouteTableId": {
      "Value": {
        "Fn::GetAtt": [
          "EcsVpc",
          "RouteTableId"
        ]
      }
    }
  }
}
```

5.5.1.30. ALIYUN::ECS::VSwitch

ALIYUN::ECS::VSwitch is used to create a VSwitch.

Statement

```
{
  "Type": "ALIYUN::ECS::VSwitch",
  "Properties": {
    "VSwitchName": String,
    "VpcId": String,
    "Description": String,
    "Ipv6CidrBlock": Integer,
    "ZoneId": String,
    "CidrBlock": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|---------------|--------|----------|----------|---|---|
| VpcId | String | No | No | The ID of the VPC where a vSwitch is to be created | None |
| ZoneId | String | No | No | The ID of the zone where the instance resides. | None |
| VSwitchName | String | Yes | True | The name of the VSwitch. | <ul style="list-style-type: none"> The name must be 2 to 128 characters in length It must start with a letter. Cannot <code>http://</code> or <code>https://</code> the beginning. It must start with a letter and cannot start with <code>http://</code> or <code>https://</code>. |
| CidrBlock | String | No | No | The CIDR block of the VSwitch. | The VSwitch CIDR block must be a subset of the CIDR block assigned to the VPC where the VSwitch resides and not be used by other VSwitches. |
| Description | String | Yes | True | The description of the vSwitch. | The description must be 2 to 256 characters in length. It cannot start with <code>http://</code> or <code>https://</code> . |
| Ipv6CidrBlock | String | Optional | Released | The IPv6 CIDR block of the VSwitch. You can customize the last eight bits of the IPv6 CIDR block. | Valid values: 0 to 255. The value is a decimal integer. By default, the prefix of the IPv6 CIDR block of the VSwitch is set to /64. |

Response parameters

Fn::GetAtt

- VSwitchId: indicates the vSwitch ID allocated by the system.
- CidrBlock: the IPv4 CIDR block of the vSwitch.
- Ipv6CidrBlock: the IPv6 CIDR block of the vSwitch.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "VpcName": {
      "Type": "String"
    },
    "VSwitch1CidrBlock": {
      "Type": "String",
      "Default": "172.16.100.0/24"
    },
    "VSwitch2CidrBlock": {
      "Type": "String",
      "Default": "172.16.80.0/24"
    }
  },
  "Resources": {
    "EcsVpc": {
      "Type": "ALIYUN::ECS::VPC",
      "Properties": {
        "CidrBlock": "172.16.0.0/12",
        "VpcName": {"Ref": "VpcName"},
      },
    },
    "VSwitch1": {
      "Type": "ALIYUN::ECS::VSwitch",
      "Properties": {
        "ZoneId": "cn-beijing-a",
        "CidrBlock": {"Ref": "VSwitch1CidrBlock"},
        "VpcId": {"Fn::GetAtt": [ "EcsVpc", "VpcId" ] },
        "VSwitchName": "create_vpc_vswitch_sg1"
      }
    },
    "VSwitch2": {
      "Type": "ALIYUN::ECS::VSwitch",
      "Properties": {
        "ZoneId": "cn-beijing-a",
        "CidrBlock": {"Ref": "VSwitch2CidrBlock"},
        "VpcId": {"Fn::GetAtt": [ "EcsVpc", "VpcId" ] },
        "VSwitchName": "create_vpc_vswitch_sg2"
      }
    },
    "SG_VSwitch1": {
      "Type": "ALIYUN::ECS::SecurityGroup",
      "Properties": {
        "SecurityGroupName": "app_mall",
        "Description": "this is created by heat",
        "VpcId": {"Fn::GetAtt": [ "EcsVpc", "VpcId" ] }
      },
    },
    "Outputs": {
      "SecurityGroupId": {
```

```
    "Value":{"get_attr":["SG_VSwitch1","SecurityGroupId"]}
  }
}
},
"SG_VSwitch1_InRule":{
  "Type":"ALIYUN::ECS::SecurityGroupIngress",
  "Properties":{
    "SecurityGroupId":{"Fn::GetAtt":["SG_VSwitch1","SecurityGroupId"]},
    "IpProtocol":"tcp",
    "PortRange":"1/65535",
    "SourceCidrIp":{"Ref":"VSwitch2CidrBlock"}
  }
},
"SG_VSwitch1_OutRule":{
  "Type":"ALIYUN::ECS::SecurityGroupEgress",
  "Properties":{
    "SecurityGroupId":{"Fn::GetAtt":["SG_VSwitch1","SecurityGroupId"]},
    "IpProtocol":"tcp",
    "PortRange":"1/65535",
    "DestCidrIp":{"Ref":"VSwitch2CidrBlock"}
  }
},
"SG_VSwitch2":{
  "Type":"ALIYUN::ECS::SecurityGroup",
  "Properties":{
    "SecurityGroupName":"app_mall",
    "Description":"this is created by heat",
    "VpcId":{"Fn::GetAtt":["EcsVpc","VpcId"]}
  }
},
"SG_VSwitch2_InRule":{
  "Type":"ALIYUN::ECS::SecurityGroupIngress",
  "Properties":{
    "SecurityGroupId":{"Fn::GetAtt":["SG_VSwitch2","SecurityGroupId"]},
    "IpProtocol":"tcp",
    "PortRange":"1/65535",
    "SourceCidrIp":{"Ref":"VSwitch1CidrBlock"}
  }
},
"SG_VSwitch2_OutRule":{
  "Type":"ALIYUN::ECS::SecurityGroupEgress",
  "Properties":{
    "SecurityGroupId":{"Fn::GetAtt":["SG_VSwitch2","SecurityGroupId"]},
    "IpProtocol":"tcp",
    "PortRange":"1/65535",
    "DestCidrIp":{"Ref":"VSwitch1CidrBlock"}
  }
}
}
}
```

5.5.2. ESS

5.5.2.1. ALIYUN::ESS::AlarmTask

ALIYUN::ESS::AlarmTask is used to create a metric-based alarm task.

Syntax

```
{
  "Type": "ALIYUN::ESS::AlarmTask",
  "Properties": {
    "Statistics": String,
    "Name": String,
    "EvaluationCount": Integer,
    "Period": Integer,
    "MetricType": String,
    "ComparisonOperator": String,
    "Dimensions": List,
    "ScalingGroupID": String,
    "AlarmAction": List,
    "Threshold": Number,
    "MetricName": String,
    "GroupID": Integer,
    "Description": String
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|--------------------|---------|----------|----------|---|--|
| Statistics | String | No | No | The method used to calculate monitoring data. The statistics must be appropriate for the metric chosen. | Valid values: Average, Minimum, and Maximum. Default value: Average. |
| Name | String | No | Yes | The name of the alarm rule. | None |
| EvaluationCount | Integer | No | No | The number of consecutive times that the threshold must be exceeded before an alarm is triggered. | Default value: 3. Minimum value: 1. |
| Period | Integer | No | No | The metric query period, which must be appropriate for the metric chosen. Unit: seconds. | Valid values: 60, 120, 300, and 900. Default value: 300. |
| MetricType | String | No | No | The metric type. | Valid values: system and custom. |
| ComparisonOperator | String | No | No | The alarm comparison operator used to define a condition in the alarm rule. | Valid values: <=, <, >, and >=. |
| Dimensions | List | No | No | The list of instances associated with the alarm rule. | You must include at least one instance in the list. |

| Property | Type | Required | Editable | Description | Constraint |
|----------------|---------|----------|----------|---|---|
| ScalingGroupId | String | Yes | No | The ID of the scaling group. | None |
| AlarmAction | List | Yes | Yes | The list of alarm actions. | You must include one to five alarm actions in the list. |
| Threshold | Number | Yes | No | The alarm threshold, which must be a numeric value. | None |
| MetricName | String | Yes | No | The metric name of a service. For more information, see the metrics defined for each service. | None |
| GroupId | Integer | No | No | The group ID. | None |
| Description | String | No | Yes | The description of the alarm task. | None |

Dimensions syntax

```
"Dimensions": [
  {
    "DimensionKey": String,
    "DimensionValue": String
  }
]
```

Dimensions properties

| Property | Type | Required | Editable | Description | Constraint |
|----------------|--------|----------|----------|-------------|------------|
| DimensionValue | String | Yes | No | None | None |
| DimensionKey | String | Yes | No | None | None |

Response parameters

Fn::GetAtt

AlarmTaskId: the ID of the alarm task.

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "ComparisonOperator": {
      "Type": "String",
      "Description": "Comparison Operator",
      "AllowedValues": [
        ">=",
        "<=",
        ">",

```

```
"<"
]
},
"Description": {
  "Type": "String",
  "Description": "Description"
},
"ScalingGroupId": {
  "Type": "String",
  "Description": "The ID of the scaling group."
},
"MetricType": {
  "Type": "String",
  "Description": "Metric Type",
  "AllowedValues": [
    "system",
    "custom"
  ]
},
"EvaluationCount": {
  "Type": "Number",
  "Description": "Evaluation Count",
  "MinValue": 1
},
"Period": {
  "Type": "Number",
  "Description": "Period",
  "AllowedValues": [
    60,
    120,
    300,
    900
  ]
},
"Dimensions": {
  "Type": "CommaDelimitedList",
  "Description": "Dimensions",
  "MinLength": 1
},
"Statistics": {
  "Type": "String",
  "Description": "Statistics",
  "AllowedValues": [
    "Average",
    "Minimum",
    "Maximum"
  ]
},
"Name": {
  "Type": "String",
  "Description": "Name"
},
"GroupId": {
  "Type": "Number",
  "Description": "Group Id"
},
"MetricName": {
  "Type": "String",
  "Description": "Metric Name"
},
}
```

```

"AlarmAction": {
  "Type": "CommaDelimitedList",
  "Description": "Alarm Actions",
  "MinLength": 1,
  "MaxLength": 5
},
"Threshold": {
  "Type": "Number",
  "Description": "Threshold"
}
},
"Resources": {
  "AlarmTask": {
    "Type": "ALIYUN::ESS::AlarmTask",
    "Properties": {
      "ComparisonOperator": {
        "Ref": "ComparisonOperator"
      },
      "Description": {
        "Ref": "Description"
      },
      "ScalingGroupId": {
        "Ref": "ScalingGroupId"
      },
      "MetricType": {
        "Ref": "MetricType"
      },
      "EvaluationCount": {
        "Ref": "EvaluationCount"
      },
      "Period": {
        "Ref": "Period"
      },
      "Dimensions": {
        "Fn::Split": [
          ",",
          {
            "Ref": "Dimensions"
          }
        ],
        {
          "Ref": "Dimensions"
        }
      ]
    }
  },
  "Statistics": {
    "Ref": "Statistics"
  },
  "Name": {
    "Ref": "Name"
  },
  "GroupId": {
    "Ref": "GroupId"
  },
  "MetricName": {
    "Ref": "MetricName"
  },
  "AlarmAction": {
    "Fn::Split": [
      ",",
      {

```

```

    "Ref": "AlarmAction"
  },
  {
    "Ref": "AlarmAction"
  }
]
},
"Threshold": {
  "Ref": "Threshold"
}
}
},
"Outputs": {
  "AlarmTaskId": {
    "Description": "The alarm task ID",
    "Value": {
      "Fn::GetAtt": [
        "AlarmTask",
        "AlarmTaskId"
      ]
    }
  }
}
}
}
}

```

5.5.2.2. ALIYUN::ESS::AlarmTaskEnable

ALIYUN::ESS::AlarmTaskEnable is used to start an alarm task. You can call this operation to enable alarm tasks when the task is stopped.

Statement

```

{
  "Type": "ALIYUN::ESS::AlarmTaskEnable",
  "Properties": {
    "AlarmTaskId": String,
    "Enable": Boolean
  }
}

```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------|--------|----------|----------|--------------------------------|------------|
| AlarmTaskId | String | No | No | The ID of the monitoring task. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|---|---|
| Enable | String | Retained | Yes | Specifies whether to enable the alarm task. | Valid values: 1025 to 10000. You cannot use the following commonly used port numbers: 2222, 4500, 4510, 4560, 7505, 9000, 9001, and 9002. |

Response parameters

Fn::GetAtt

None

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "Enable": {
      "Type": "Boolean",
      "Description": "Enable alarm task or not",
      "AllowedValues": [
        "True",
        "true",
        "False",
        "false"
      ]
    },
    "AlarmTaskId": {
      "Type": "String",
      "Description": "The id of alarm task."
    }
  },
  "Resources": {
    "AlarmTaskEnable": {
      "Type": "ALIYUN::ESS::AlarmTaskEnable",
      "Properties": {
        "Enable": {
          "Ref": "Enable"
        }
      },
      "AlarmTaskId": {
        "Ref": "AlarmTaskId"
      }
    }
  },
  "Outputs": {}
}
```

5.5.2.3. ALIYUN::ESS::LifecycleHook

ALIYUN::ESS::LifecycleHook is used to create a lifecycle hook for a scaling group.

Syntax

```
{
  "Type": "ALIYUN::ESS::LifecycleHook",
  "Properties": {
    "LifecycleHookName": String,
    "NotificationArn": String,
    "HeartbeatTimeout": Integer,
    "NotificationMetadata": String,
    "ScalingGroupId": String,
    "DefaultResult": String,
    "LifecycleTransition": String
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|-------------------|--------|----------|----------|---|--|
| LifecycleHookName | String | No | Yes | The name of the lifecycle hook. Each lifecycle hook name must be unique within a scaling group. | <p>The name must be 2 to 40 characters in length and can contain letters, digits, underscores (_), hyphens (-), and periods (.). It must start with a letter or digit.</p> <p>The default name is the ID of the lifecycle hook.</p> |
| NotificationArn | String | No | Yes | The Alibaba Cloud Resource Name (ARN) of the notification target that Auto Scaling uses to notify you when an instance is in the transition state for the lifecycle hook. | <p>This target can be either an MNS queue or an MNS topic. The format of the parameter value is <code>acs:ess:{region}:{account-id}:{resource-relative-id}</code>.</p> <ul style="list-style-type: none"> region : the region where the scaling group resides. account-id : the ID of the Apsara Stack tenant account. <p>Examples:</p> <ul style="list-style-type: none"> MNS queue: <code>acs:ess:{region}:{account-id}:queue/{queuename}</code> MNS topic: <code>acs:ess:{region}:{account-id}:topic/{topicname}</code> |

| Property | Type | Required | Editable | Description | Constraint |
|----------------------|---------|----------|----------|--|--|
| HeartbeatTimeout | Integer | No | Yes | The waiting period before the lifecycle hook times out. When the lifecycle hook times out, the scaling group performs the action specified by the DefaultResult parameter. Unit: seconds. | Valid values: 30 to 21600. Default value: 600. |
| NotificationMetadata | String | No | Yes | The fixed string to include when Auto Scaling sends a notification about the wait state of a scaling activity. Auto Scaling sends the specified NotificationMetadata parameter value along with the notification message so that you can easily categorize notifications. The NotificationMetadata parameter is valid only after you set the NotificationArn parameter. | The parameter value cannot exceed 128 characters in length. |
| ScalingGroupId | String | Yes | No | The ID of the scaling group. | None |
| DefaultResult | String | No | Yes | The action that the scaling group takes when the lifecycle hook times out. If the scaling group has multiple lifecycle hooks and one of them is terminated when the DefaultResult parameter is set to ABANDON during a scale-in event, the remaining lifecycle hooks in the same scaling group will also be terminated. Otherwise, the scaling activity will proceed normally after the waiting period expires and continue with the action specified by the DefaultResult parameter. | Valid values: <ul style="list-style-type: none"> CONTINUE: The scaling group continues the scale-in or scale-out event. ABANDON: The scaling group releases the created ECS instances if the scaling activity type is scale-out or removes the ECS instances to be scaled in if the scaling activity type is scale-in. Default value: CONTINUE. |

| Property | Type | Required | Editable | Description | Constraint |
|---------------------|--------|----------|----------|---|--|
| LifecycleTransition | String | Yes | Yes | The type of scaling activity to which the lifecycle hook applies. | Valid values: <ul style="list-style-type: none"> SCALE_OUT: scale-out events of the scaling group. SCALE_IN: scale-in events of the scaling group. |

Response parameters

Fn::GetAtt

LifecycleHookId: the ID of the lifecycle hook.

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "LifecycleHookName": {
      "Type": "String",
      "Description": "The name of the lifecycle hook. Each name must be unique within a scaling group. The name must be 2 to 40 characters in length and can contain letters, numbers, Chinese characters, and special characters including underscores (_), hyphens (-) and periods (.).\nDefault value: Lifecycle Hook ID",
      "AllowedPattern": "^[a-zA-Z0-9\\u4e00-\\u9fa5][-_a-zA-Z0-9\\u4e00-\\u9fa5]{1,63}$"
    },
    "NotificationArn": {
      "Type": "String",
      "Description": "The Alibaba Cloud Resource Name (ARN) of the notification target that Auto Scaling will use to notify you when an instance is in the transition state for the lifecycle hook. This target can be either an MNS queue or an MNS topic. The format of the parameter value is acs:ess:{region}:{account-id}:{resource-relative-id}.\nregion: the region to which the scaling group locates\naccount-id: Alibaba Cloud ID\nFor example:\nMNS queue: acs:ess:{region}:{account-id}:queue/{queueName}\nMNS topic: acs:ess:{region}:{account-id}:topic/{topicName}",
      "AllowedPattern": "^[acs:ess:([a-zA-Z0-9-]+):(\\d+):(queue|topic)/([a-zA-Z0-9][a-zA-Z0-9-]{0,255})$]",
      "MaxLength": 300
    },
    "ScalingGroupId": {
      "Type": "String",
      "Description": "The ID of the scaling group."
    },
    "LifecycleTransition": {
      "Type": "String",
      "Description": "The scaling activities to which lifecycle hooks apply Value range:\n SCALE_OUT: scale-out event\n SCALE_IN: scale-in event",
      "AllowedValues": [
        "SCALE_OUT",
        "SCALE_IN"
      ]
    },
    "HeartbeatTimeout": {
      "Type": "Number",
      "Description": "The time, in seconds, that can elapse before the lifecycle hook times out. If the lifecycle hook times out, the scaling group performs the default action (DefaultResult). The range is from 30 to 21,600 seconds. The default value is 600 seconds.\nYou can prevent the lifecycle hook from timing out by calling the RecordLifecycleActionHeartbeat operation. You can also terminate the lifecycle action by calling the CompleteLifecycleAction operation.",
      "MinValue": 30,
      "MaxValue": 21600
    }
  }
}
```

```

},
  "NotificationMetadata": {
    "Type": "String",
    "Description": "The fixed string that you want to include when Auto Scaling sends a message about the wait state of the scaling activity to the notification target. The length of the parameter can be up to 128 characters. Auto Scaling will send the specified NotificationMetadata parameter along with the notification message so that you can easily categorize your notifications. The NotificationMetadata parameter will only take effect after you specify the NotificationArn parameter.",
    "MaxLength": 128
  },
  "DefaultResult": {
    "Type": "String",
    "Description": "The action that the scaling group takes when the lifecycle hook times out. Value range:\n CONTINUE: the scaling group continues with the scale-in or scale-out process.\n ABANDON: the scaling group stops any remaining action of the scale-in or scale-out event.\nDefault value: CONTINUE\nIf the scaling group has multiple lifecycle hooks and one of them is terminated by the DefaultResult=ABANDON parameter during a scale-in event (SCALE_IN), the remaining lifecycle hooks under the same scaling group will also be terminated. Otherwise, the action following the wait state is the next action, as specified in the parameter DefaultResult, after the last lifecycle event under the same scaling group.",
    "AllowedValues": [
      "CONTINUE",
      "ABANDON"
    ]
  }
},
"Resources": {
  "LifecycleHook": {
    "Type": "ALIYUN::ESS::LifecycleHook",
    "Properties": {
      "LifecycleHookName": {
        "Ref": "LifecycleHookName"
      },
      "NotificationArn": {
        "Ref": "NotificationArn"
      },
      "ScalingGroupId": {
        "Ref": "ScalingGroupId"
      },
      "LifecycleTransition": {
        "Ref": "LifecycleTransition"
      },
      "HeartbeatTimeout": {
        "Ref": "HeartbeatTimeout"
      },
      "NotificationMetadata": {
        "Ref": "NotificationMetadata"
      },
      "DefaultResult": {
        "Ref": "DefaultResult"
      }
    }
  }
},
"Outputs": {
  "LifecycleHookId": {
    "Description": "The lifecycle hook ID",
    "Value": {
      "Fn::GetAtt": [
        "LifecycleHook",
        "LifecycleHookId"
      ]
    }
  }
}

```

```

    }
  }
}
}

```

5.5.2.4. ALIYUN::ESS::ScalingConfiguration

ALIYUN::ESS::ScalingConfiguration is used to create a scaling configuration for a scaling group.

Statement

```

{
  "Type": "ALIYUN::ESS::ScalingConfiguration",
  "Properties": {
    "PasswordInherit": Boolean,
    "DiskMappings": List,
    "RamRoleName": String,
    "IoOptimized": String,
    "InternetChargeType": String,
    "KeyPairName": String,
    "InstanceId": String,
    "InstanceTypes": List,
    "ImageId": String,
    "ResourceGroupId": String,
    "SpotStrategy": String,
    "InstanceType": String,
    "SystemDiskCategory": String,
    "SystemDiskSize": Integer,
    "SystemDiskAutoSnapshotPolicyId": String,
    "InternetMaxBandwidthOut": Integer,
    "InstanceName": String,
    "InternetMaxBandwidthIn": Integer,
    "ScalingConfigurationName": String,
    "UserData": String,
    "DeploymentSetId": String,
    "SecurityGroupId": String,
    "SpotPriceLimit": Number,
    "HpcClusterId": String,
    "ScalingGroupId": String,
    "SpotPriceLimitForInstanceType": Map,
    "TagList": List
  }
}

```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|---|------------|
| ResourceGroupId | String | Yes | True | The ID of the resource group to which the instance belongs. | None |
| DeploymentSetId | String | Yes | Released | The ID of the deployment set. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------------------|--------|----------|----------|---|--|
| HpcClusterId | String | Yes | Released | The ID of the E-HPC cluster to which the instance belongs. | None |
| ScalingGroupId | String | No | No | The ID of the scaling group to which the scaling configuration belongs. | None |
| DiskMappings | List | No. | True | The disks to be attached to created instances. | A maximum of 16 disks can be attached to each instance. |
| InternetChargeType | String | Yes | True | The billing method for Internet usage. | Valid values: <ul style="list-style-type: none"> PayByBandwidth PayByTraffic: pay-by-traffic Default value: PayByTraffic |
| InternetMaxBandwidthIn | String | Optional | Released | The maximum inbound bandwidth from the Internet. | Unit: Mbit/s. Valid values: 1 to 100. Default value: 100 |
| InternetMaxBandwidthOut | String | No. | True | The maximum outbound bandwidth to the Internet. | Valid values: <ul style="list-style-type: none"> Pay-by-bandwidth: 0 to 100. Default value: 0. Pay-by-data-transfer: 1 to 200. This parameter is required. Unit: Mbit/s. |
| InstanceId | String | Yes | Released | The instance ID of the scaling configuration. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------|--------|----------|----------|--|--|
| SystemDiskCategory | String | Yes | True | The category of the system disk. | Valid values: <ul style="list-style-type: none"> cloud: indicates a basic disk. cloud_efficiency: indicates an ultra disk. cloud_ssd: indicates a standard SSD. ephemeral_ssd: indicates a local SSD. cloud_essd: enhanced SSD (ESSD) Default value: cloud for Generation I instance types that are not I/O optimized, default value: cloud_efficiency. |
| ImageId | String | Yes | True | The ID of the image used to start the instance. You can use a public image, a custom image, or an Alibaba Cloud Marketplace image. | None |
| InstanceType | String | Yes | True | The specification of the instance. | None |
| SecurityGroupId | String | Yes | True | The ID of the security group to which the instance belongs. | None |
| IoOptimized | String | Yes | True | Specifies whether the created instances are I/O optimized. | Valid values: <ul style="list-style-type: none"> none (non-I/O optimized) optimized Default value: none. |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------------|--------|----------|----------|---|---|
| ScalingConfigurationName | String | Yes | True | The name of the scaling configuration. | <ul style="list-style-type: none"> The name must be 2 to 64 characters in length. It can contain letters, digits, underscores (_), hyphens (-), and periods (.). It must start with a letter or digit. The name of the scaling configuration must be unique within a scaling group. If this parameter is not specified, the value of scalingconfigurationid is used. |
| KeyPairName | String | Yes | True | The name of the key pair that is bound to the instance. | <ul style="list-style-type: none"> This parameter is ignored if the instance type is Windows and the default value is null. If the instance type is Linux, password logon is disabled by default. |
| RamRoleName | String | Yes | True | The RAM role name of the instance. | You can use RAM API ListRoles you can call this operation to query the RAM role name of an instance. |
| SystemDiskSize | String | No. | True | The size of the system disk. Unit: GB. | <p>Valid values: 40 to 500. Unit: GB.</p> <p>If a custom image is used to create a system disk, the system disk size must be larger than the size of the custom image.</p> |
| UserData | String | Yes | True | The user data that you pass when you create the instance. | The user can encode up to 16KB in size. You do not need to perform Base64 encoding. Special characters must be escaped with a backslash (\). |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|---------|----------|----------|--|--|
| InstanceTypes | List | No. | True | The instance types from which ECS instances can be created. If you specify InstanceTypes, InstanceType is invalid. | Up to 10 instance types can be configured in a scaling configuration. The priority of each instance type is decreased in the order of its list elements. Auto Scaling creates instances in order of priority. If an instance of the highest priority type cannot be created, Auto Scaling will create an instance of the next highest priority type. |
| PasswordInherit | Boolean | No. | True | Specifies whether to use the preconfigured password of the specified image. | To use this parameter, ensure that a password is configured for the specified image. |
| TagList | List | No. | True | The tags of the instance. | Tags must be specified as key-value pairs. You can specify a maximum of five Tag groups in the format of {"key1": "value1", "key2": "value2", ... "key5": "value5"} . The key can contain a maximum of 64 characters. alivun . http:// or https:// the beginning. If you use tags, the key cannot be an empty string. The value must be 0 to 128 characters in length. |
| SpotStrategy | String | Yes | True | The preemption policy for pay-as-you-go instances. | Valid values: <ul style="list-style-type: none"> NoSpot (pay-as-you-go instance) SpotWithPriceLimit (a preemptible instance with a maximum price) SpotAsPriceGo (the SpotAsPriceGo parameter that is set automatically based on the actual market price.) Default value: NoSpot. |
| InstanceName | String | Yes | True | The name of the instance created based on the current scaling configuration. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------------------|--------|----------|----------|---|--|
| SpotPriceLimit | Number | No. | True | The maximum hourly price of the instance. | A maximum of three decimal places can be specified. This parameter takes effect only when the SpotStrategy parameter is set to SpotWithPriceLimit. The value of this parameter can be overwritten by the value of the SpotPriceLimitForInstanceType parameter. |
| SpotPriceLimitForInstanceType | Map | No. | True | Preemptible instance type and bid of the instance. | The format is {"<instance_type_1>": {"price_limit_1": ...}, {"<instance_type_10>": {"price_limit_10": ...}}. This parameter takes effect only when the SpotStrategy parameter is set to SpotWithPriceLimit. You can set up to 10 instance groups and prices. |
| SystemDiskAutoSnapshotPolicyId | String | Yes | True | The ID of the automatic snapshot policy applied to the data disk. | None |

DiskMappings syntax

```
"DiskMappings": [
{
  "Category": String,
  "Device": String,
  "SnapshotId": String,
  "Size": String,
  "Encrypted": String,
  "KMSKeyId": String,
  "Description": String,
  "DiskName": String
}
]
```

DiskMappings properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|------|----------|----------|-------------|------------|
|-----------|------|----------|----------|-------------|------------|

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------|--------|----------|----------|-----------------------------------|---|
| Size | String | No | No | The size of the data disk. | <p>Valid values:</p> <ul style="list-style-type: none"> cloud: 5 to 2000 cloud_efficiency: 20 to 32768 cloud_ssd: 20 to 32768 cloud_essd: 20 to 32768. ephemeral_ssd: 5 to 800 <p>The value of this parameter must be greater than or equal to that of the snapshot specified by SnapshotId.</p> <p>Unit: GiB.</p> |
| Category | String | Yes | Released | The type of the data disk. | <p>Valid values: cloud, cloud_efficiency, cloud_ssd, ephemeral_ssd, and cloud_essd. For I/O optimized instances, the default value is cloud_efficiency. For non-I/O optimized instances, the default value is cloud.</p> |
| DiskName | String | Yes | Released | The name of the data disk. | <p>The name must be 2 to 128 characters in length. It can contain letters, digits, colons (:), underscores (_), and hyphens (-). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code> the beginning.</p> |
| Description | String | Yes | Released | The description of the data disk. | <p>The description must be 2 to 256 characters in length. Cannot <code>http://</code> or <code>https://</code> the beginning.</p> |
| Device | String | Yes | Released | The device name of the data disk. | <p>By default, the system automatically assigns a value for this parameter when the ECS instance is created. The value starts from <code>/dev/xvdb</code> and ends at <code>/dev/xvdz</code>.</p> |

| Parameter | Type | Required | Editable | Description | Constraint |
|------------|--------|----------|----------|--|--|
| SnapshotId | String | Yes | Released | The ID of the snapshot used to create the data disk. | If this parameter is specified, the Size parameter will be ignored, and the Size of the created disk will be the Size of the specified snapshot. If the snapshot was created on or before July 15, 2013, calling the snapshot is denied and InvalidSnapshot.TooOld is displayed in the response parameter. |
| Encrypted | String | Yes | Released | Specifies whether to encrypt the data disk. | Default value: false. |
| KMSKeyId | String | Yes | Released | The KMS key ID for data disk N. | Valid values: 1025 to 10000. You cannot use the following commonly used port numbers: 2222, 4500, 4510, 4560, 7505, 9000, 9001, and 9002. |

Response parameters

Fn::GetAtt

ScalingConfigurationId: the ID of the scaling configuration. This ID is a globally unique identifier (GUID) generated by the system.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "ScalingConfiguration": {
      "Type": "ALIYUN::ESS::ScalingConfiguration",
      "Properties": {
        "ImageId": "ubuntu1404_64_20G_aliaegis_2015****.vhd",
        "InstanceType": "ecs.t1.small",
        "InstanceId": "i-25xhh****",
        "InternetChargeType": "PayByTraffic",
        "InternetMaxBandwidthIn": 1,
        "InternetMaxBandwidthOut": 20,
        "SystemDisk_Category": "cloud",
        "ScalingGroupId": "bwhtvpcBcKYac9fe3vd0****",
        "SecurityGroupId": "sg-25zwc****",
        "DiskMappings": [
          {
            "Size": 10
          },
          {
            "Category": "cloud",
            "Size": 10
          }
        ]
      }
    },
    "Outputs": {
      "ScalingConfiguration": {
        "Value": {"get_attr": ["ScalingConfigurationId"]}
      }
    }
  }
}
```

5.5.2.5. ALIYUN::ESS::ScalingGroup

ALIYUN::ESS::ScalingGroup is used to create a scaling group. A scaling group is a group of ECS instances that are dynamically scaled based on the configured scenario. A scaling group does not take effect immediately after it is created. You must use ALIYUN::ESS::ScalingGroupEnable to enable the scaling group to trigger scaling rules and execute scaling activities.

Syntax

```
{
  "Type": "ALIYUN::ESS::ScalingGroup",
  "Properties": {
    "MultiAZPolicy": String,
    "DesiredCapacity": Integer,
    "NotificationConfigurations": List,
    "ProtectedInstances": List,
    "LaunchTemplateId": String,
    "LaunchTemplateVersion": String,
    "ScalingGroupName": String,
    "VSwitchIds": List,
    "DefaultCooldown": Integer,
    "MinSize": Integer,
    "GroupDeletionProtection": Boolean,
    "MaxSize": Integer,
    "Instanceld": String,
    "VSwitchId": String,
    "LoadBalancerIds": List,
    "StandbyInstances": List,
    "RemovalPolicies": List,
    "HealthCheckType": String,
    "DBInstancelds": List
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|----------|---------|----------|----------|---|---|
| MinSize | Integer | Yes | Yes | The minimum number of ECS instances in the scaling group. | Valid values: 0 to 1000. When the number of ECS instances in the scaling group is less than the MinSize value, Auto Scaling automatically creates ECS instances until the number of instances is equal to the MinSize value. |

| Property | Type | Required | Editable | Description | Constraint |
|------------------|---------|----------|----------|--|--|
| MaxSize | Integer | Yes | Yes | The maximum number of ECS instances in the scaling group. | Valid values: 0 to 1000. When the number of ECS instances in the scaling group is greater than the MaxSize value, Auto Scaling removes ECS instances from the scaling group until the number of instances is equal to the MaxSize value. |
| ScalingGroupName | String | No | Yes | The display name of the scaling group. | <ul style="list-style-type: none"> The name must be 2 to 40 characters in length and can contain letters, digits, underscores (_), hyphens (-), and periods (.). It must start with an uppercase letter, lowercase letter, or digit. The name must be unique to an Alibaba Cloud account in a region. The default value is the ID of the scaling group. |
| LaunchTemplateId | String | No | Yes | The ID of the instance launch template from which the scaling group obtains launch configurations. | None |

| Property | Type | Required | Editable | Description | Constraint |
|-----------------------|--------|----------|----------|--|---|
| LaunchTemplateVersion | String | No | Yes | The version of the instance launch template. | Valid values: <ul style="list-style-type: none"> • The fixed template version number. • Default: The default template version is always used. • Latest: The latest template version is always used. |
| RemovalPolicies | List | No | Yes | The list of one or more policies that are used to remove ECS instances from the scaling group. | Default value: OldestScalingConfiguration or OldestInstance. Valid values: <ul style="list-style-type: none"> • OldestInstance: removes the ECS instance that is added to the scaling group at the earliest point in time. • NewestInstance: removes the ECS instance that is added to the scaling group at the latest point in time. • OldestScalingConfiguration: removes the ECS instance that is created based on the earliest scaling configuration. |
| VSwitchId | String | No | No | The ID of the vSwitch. | None |

| Property | Type | Required | Editable | Description | Constraint |
|-----------------|---------|----------|----------|--|--|
| LoadBalancerIds | List | No | Yes | The ID of the Server Load Balancer (SLB) instance. | This value can be a JSON array that contains up to five SLB instance IDs. Separate multiple IDs with commas (,). |
| DefaultCooldown | Integer | No | Yes | The cooldown time after a scaling activity (adding or removing ECS instances) is executed. | <ul style="list-style-type: none"> Valid values: 0 to 86400. Unit: seconds. Default value: 300. During the cooldown time, the scaling group executes only scaling activities that are triggered by Cloud Monitor event-triggered tasks. |
| DBInstanceIds | List | No | Yes | The list of one or more ApsaraDB RDS instance IDs. | This value can be a JSON array that contains up to eight ApsaraDB RDS instance IDs. Separate multiple IDs with commas (,). |

| Property | Type | Required | Editable | Description | Constraint |
|---------------|--------|----------|----------|---|---|
| VSwitchIds | List | No | No | The list of one or more vSwitch IDs. | <p>You can specify a maximum of five vSwitch IDs. If you specify this parameter, the VSwitchId parameter is ignored. vSwitches are sorted in descending order of priority. When an ECS instance cannot be created in the zone where the vSwitch with the highest priority resides, the system uses the vSwitch with the next highest priority to create the ECS instance.</p> |
| MultiAZPolicy | String | No | No | The ECS instance scaling policy for the multi-zone scaling group. | <p>Valid values:</p> <ul style="list-style-type: none"> • PRIORITY: ECS instances are scaled based on the specified vSwitch. When an ECS instance cannot be created in the zone where the vSwitch with the highest priority resides, the system uses the vSwitch with the next highest priority to create the ECS instance. • BALANCE: ECS instances are distributed evenly in multiple zones specified in the scaling group. • COST_OPTIMIZ |

| Property | Type | Required | Editable | Description | ED: ECS Constraint instances are |
|----------------------------|--------|----------|----------|---|--|
| | | | | | created based on the unit price of vCPUs, from low to high. Preemptible instances are created first when preemptible instance types are specified for the scaling configuration. Pay-as-you-go instances are automatically created when no preemptible instances are available due to issues such as insufficient ECS resources. |
| NotificationConfigurations | List | No | Yes | The notification configurations for event and resource changes. | None |
| ProtectedInstances | List | No | Yes | The number of protected ECS instances in the scaling group. | Maximum value: 1000. |
| StandbyInstances | List | No | Yes | The number of ECS instances that are in the standby state in the scaling group. | Maximum value: 1000. |
| HealthCheckType | String | No | Yes | The health check type. | Valid values: <ul style="list-style-type: none"> • ECS • NONE |

| Property | Type | Required | Editable | Description | Constraint |
|-------------------------|---------|----------|----------|---|---|
| GroupDeletionProtection | Boolean | No | Yes | Specifies whether to enable deletion protection for the scaling group. | Default value: false. Valid values: <ul style="list-style-type: none"> true: enables deletion protection for the scaling group. In this case, you cannot delete the scaling group. false: disables deletion protection for the scaling group. |
| DesiredCapacity | Integer | No | Yes | The expected number of ECS instances in the scaling group. The scaling group automatically keeps the number of ECS instances at the expected value. | The number of ECS instances must be greater than the MinSize value and less than the MaxSize value. |
| InstanceId | String | No | No | The ID of the ECS instance from which the scaling group obtains configuration information to create scaling configurations. | None |

Response parameters

Fn::GetAtt

ScalingGroupId: the ID of the scaling group. This ID is a globally unique identifier (GUID) that is generated by the system.

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "ScalingGroup": {
      "Type": "ALIYUN::ESS::ScalingGroup",
      "Properties": {
        "MaxSize": 1,
        "MinSize": 1,
        # "ScalingGroupName": "HeatCreatedR****",
        # "DefaultCooldown": 500,
        # "RemovalPolicy_1": "",
        # "RemovalPolicy_2": "",
      }
    }
  },
  "Outputs": {
    "ScalingGroup": {
      "Value": {"Fn::GetAtt": ["ScalingGroup", "ScalingGroupId"]}
    }
  }
}
```

5.5.2.6. ALIYUN::ESS::ScalingGroupEnable

ALIYUN::ESS::ScalingGroupEnable is used to enable a scaling group.

Syntax

```
{
  "Type": "ALIYUN::ESS::ScalingGroupEnable",
  "Properties": {
    "ScalingConfigurationId": String,
    "ScalingRuleArisExecuteVersion": Integer,
    "ScalingRuleAris": List,
    "ScalingGroupId": String,
    "RemoveInstanceIds": List,
    "InstanceIds": List
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|------------------------|--------|----------|----------|---|------------|
| ScalingGroupId | String | Yes | No | The ID of the scaling group. | None |
| ScalingConfigurationId | String | No | No | The ID of the scaling configuration to be activated in the scaling group. | None |

| Property | Type | Required | Editable | Description | Constraint |
|-------------------------------|---------|----------|----------|---|--|
| InstanceIds | List | No | Yes | The IDs of ECS instances to be added to the enabled scaling group. | A maximum of 20 instance IDs can be specified. |
| ScalingRuleArisExecuteVersion | Integer | No | Yes | The version of the identifier for the scaling rule to be executed. If you change this property, all scaling rules specified by ScalingRuleAris will be executed once. | Minimum value: 0. |
| ScalingRuleAris | List | No | Yes | The unique identifiers of scaling rules in the scaling group. Invalid unique identifiers are not displayed in the query results and no errors are reported. | A maximum of 10 scaling rule identifiers can be specified. |
| RemoveInstanceIds | List | No | Yes | The IDs of ECS instances to be deleted. | A maximum of 1,000 instance IDs can be specified. |

Response parameters

Fn::GetAtt

- LifecycleState: the status of the scaling group.
- ScalingInstances: the instances that are automatically created in the scaling group.
- ScalingGroupId: the ID of the scaling group.
- ScalingRuleArisExecuteResultInstancesRemoved: the instances that are removed from the scaling group by executing the scaling rules specified by ScalingRuleAris.
- ScalingRuleArisExecuteResultNumberOfAddedInstances: the number of instances that are added to the scaling group by executing the scaling rules specified by ScalingRuleAris.
- ScalingInstanceDetails: the instance scaling details.
- ScalingRuleArisExecuteErrorInfo: the error information about the execution of the scaling rules specified by ScalingRuleAris.
- ScalingRuleArisExecuteResultInstancesAdded: the instances that are added to the scaling group by executing the scaling rules specified by ScalingRuleAris.

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "ScalingGroupEnable": {
      "Type": "ALIYUN::ESS::ScalingGroupEnable",
      "Properties": {
        "ScalingGroupId": "r0HUqbJ411cc2eQw8bU****",
        "ScalingConfigurationId": "bJLLfdexm77Ldsyptmel****",
        "InstanceIds": "",
      }
    }
  },
  "Outputs": {
    "ScalingGroupEnable": {
      "Value": {"Fn::GetAtt": ["ScalingGroupEnable", "LifecycleState"]}
    }
  }
}
```

5.5.2.7. ALIYUN::ESS::ScalingRule

ALIYUN::ESS::ScalingRule is used to create a scaling rule.

Syntax

```
{
  "Type": "ALIYUN::ESS::ScalingRule",
  "Properties": {
    "AdjustmentValue": Integer,
    "Cooldown": Integer,
    "ScalingGroupId": String,
    "AdjustmentType": String,
    "ScalingRuleName": String
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|-----------------|---------|----------|----------|--|--|
| AdjustmentValue | Integer | No | Yes | The number of ECS instances to add or release when scaling occurs. The number of ECS instances to be adjusted in a single scaling activity cannot exceed 500. | Valid values in different adjustment modes: <ul style="list-style-type: none"> QuantityChangeInCapacity: -500 to 500. PercentChangeInCapacity: -100 to 10000. TotalCapacity: 0 to 1000. |
| Cooldown | Integer | No | Yes | The cooldown period of the scaling rule. Unit: seconds. | Valid values: 0 to 86400. This parameter is empty by default. |

| Property | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|--|---|
| ScalingGroupId | String | Yes | No | The ID of the scaling group to which the scaling rule belongs. | None |
| AdjustmentType | String | Yes | Yes | The adjustment mode of the scaling rule. | Valid values: <ul style="list-style-type: none"> QuantityChangeInCapacity: adds or removes a specified number of ECS instances. PercentChangeInCapacity: adds or removes a specified proportion of ECS instances. TotalCapacity: adds or removes ECS instances to ensure that the current scaling group has a specified number of ECS instances. |
| ScalingRuleName | String | No | Yes | The display name of the scaling rule. | The name must be 2 to 40 characters in length and can contain letters, digits, underscores (_), hyphens (-), and periods (.). It must start with a letter or digit. The name of a scaling rule must be unique within the scaling group that it belongs to. The default value is the ID of the scaling rule. |

Response parameters

Fn::GetAtt

- ScalingRuleAri: the unique identifier of the scaling rule.
- ScalingRuleId: the ID of the scaling rule. It is a globally unique identifier (GUID) generated by the system.

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "ScalingRule": {
      "Type": "ALIYUN::ESS::ScalingRule",
      "Properties": {
        "ScalingRuleName": {
          "Ref": "ScalingRuleName"
        },
        "Cooldown": {
          "Ref": "Cooldown"
        }
      },
      "ScalingGroupId": {
```

```

    "Ref": "ScalingGroupId"
  },
  "AdjustmentType": {
    "Ref": "AdjustmentType"
  },
  "AdjustmentValue": {
    "Ref": "AdjustmentValue"
  }
}
},
"Parameters": {
  "ScalingRuleName": {
    "AllowedPattern": "^[a-zA-Z0-9\\u4e00-\\u9fa5][_\\.a-zA-Z0-9\\u4e00-\\u9fa5]{1,63}$",
    "Type": "String",
    "Description": "Name shown for the scaling group, which is a string containing 2 to 40 English or Chinese characters. It must begin with a number, a letter (case-insensitive) or a Chinese character and can contain numbers, \"_\", \"-\" or \".\". The account name in the same scaling group is unique in the same region. If this parameter value is not specified, the default value is ScalingRuleId."
  },
  "Cooldown": {
    "Type": "Number",
    "Description": "Cool-down time of a scaling rule. Value range: [0, 86,400], in seconds. The default value is empty.",
    "MaxValue": 86400,
    "MinValue": 0
  },
  "ScalingGroupId": {
    "Type": "String",
    "Description": "ID of the scaling group of a scaling rule."
  },
  "AdjustmentType": {
    "Type": "String",
    "Description": "Adjustment mode of a scaling rule. Optional values:\n- QuantityChangeInCapacity: It is used to increase or decrease a specified number of ECS instances.\n- PercentChangeInCapacity: It is used to increase or decrease a specified proportion of ECS instances.\n- TotalCapacity: It is used to adjust the quantity of ECS instances in the current scaling group to a specified value.",
    "AllowedValues": [
      "QuantityChangeInCapacity",
      "PercentChangeInCapacity",
      "TotalCapacity"
    ]
  },
  "AdjustmentValue": {
    "Type": "Number",
    "Description": "Adjusted value of a scaling rule. Value range:\n- QuantityChangeInCapacity: [-500, 500]\n- PercentChangeInCapacity: [-100, 10000]\n- TotalCapacity: [0, 1000]",
    "MaxValue": 10000,
    "MinValue": -500
  }
},
"Outputs": {
  "ScalingRuleAri": {
    "Description": "Unique identifier of a scaling rule.",
    "Value": {
      "Fn::GetAtt": [
        "ScalingRule",
        "ScalingRuleAri"
      ]
    }
  }
}
}

```

```
"ScalingRuleId": {
  "Description": "ID of a scaling rule, generated by the system and globally unique.",
  "Value": {
    "Fn::GetAtt": [
      "ScalingRule",
      "ScalingRuleId"
    ]
  }
}
}
```

5.5.2.8. ALIYUN::ESS::ScheduledTask

ALIYUN::ESS::ScheduledTask is used to create a scheduled task based on input parameters.

Syntax

```
{
  "Type": "ALIYUN::ESS::ScheduledTask",
  "Properties": {
    "TaskEnabled": Boolean,
    "Description": String,
    "ScheduledTaskName": String,
    "LaunchExpirationTime": Integer,
    "LaunchTime": String,
    "RecurrenceEndTime": String,
    "RecurrenceType": String,
    "RecurrenceValue": String,
    "ScheduledAction": String
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|-------------|---------|----------|----------|---|--|
| TaskEnabled | Boolean | No | Yes | Specifies whether to start the scheduled task. <ul style="list-style-type: none"> true: starts the scheduled task. false: stops the scheduled task. Default value: true. | None |
| Description | String | No | Yes | The description of the scheduled task. | The description must be 2 to 200 characters in length. |

| Property | Type | Required | Editable | Description | Constraint |
|----------------------|---------|----------|----------|---|---|
| ScheduledTaskName | String | No | Yes | The display name of the scheduled task. | <p>The name must be 2 to 40 characters in length and can contain letters, digits, underscores (_), hyphens (-), and periods (.). It must start with a letter or digit.</p> <p>This parameter must be unique in a region and under an Apsara Stack tenant account.</p> <p>The default value is the ID of the scheduled scaling task.</p> |
| LaunchExpirationTime | Integer | No | Yes | <p>The time period during which a failed scheduled task is retried.</p> <p>Unit: seconds. Default value: 600.</p> | Valid values: 0 to 21600. |
| LaunchTime | String | Yes | Yes | <p>The time at which the scheduled task is triggered.</p> <p>Specify the time in the ISO 8601 standard in the YYYY-MM-DDThh:mmZ format. The time must be in UTC.</p> <p>If the RecurrenceType parameter is specified, the task is executed each day at the time specified by LaunchTime.</p> <p>If the RecurrenceType parameter is not specified, the task is only executed once at the date and time specified by LaunchTime.</p> <p>You cannot enter a point in time later than 90 days from the date of scheduled task creation or modification.</p> | None |

| Property | Type | Required | Editable | Description | Constraint |
|-------------------|--------|----------|----------|--|---|
| RecurrenceEndTime | String | No | Yes | <p>The end time after which the scheduled task will not be repeated.</p> <p>Specify the time in the ISO 8601 standard in the YYYY-MM-DDThh:mmZ format. The time must be in UTC.</p> <p>You cannot enter a point in time later than 90 days from the date of scheduled task creation or modification.</p> <p>If you set RecurrenceEndTime, you must also set both RecurrenceType and RecurrenceValue.</p> | None |
| RecurrenceType | String | No | Yes | <p>The interval that the scheduled task is repeated at.</p> | <p>Valid values:</p> <ul style="list-style-type: none"> • Daily: The scheduled task is executed once every specified number of days. • Weekly: The scheduled task is executed on each specified day of a week. • Monthly: The scheduled task is executed on each specified day of a month. • Cron: The scheduled task is executed based on the specified Cron expression. <p>If you set RecurrenceType, you must also set both RecurrenceEndTime and RecurrenceValue.</p> |

| Property | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|--|--|
| RecurrenceValue | String | No | Yes | Specifies how often the scheduled task recurs. | <ul style="list-style-type: none"> Daily: indicates the interval of days that the scheduled task is repeated on. You can enter a single value ranging from 1 to 31. Weekly: indicates which days of the week that the scheduled task is repeated on. You can enter multiple values separated by commas (.). The values 0 to 6 correspond to the days of the week in sequence from Sunday to Saturday. Monthly: indicates which days of the month that the scheduled task is repeated on. You can enter two values ranging from 1 to 31. The format is A-B. B must be greater than or equal to A. Cron: indicates a user-defined Cron expression that the scheduled task is repeated on. A Cron expression is written in UTC time and consists of five fields: minute, hour, day of month (date), month, and day of week. The expression can contain wildcard characters including commas (,), question marks (?), hyphens (-), asterisks (*), number signs (#), forward slashes (/), and the L and W characters. <p>If you set RecurrenceValue, you must also set both RecurrenceEndTime and RecurrenceType.</p> |
| ScheduledAction | String | Yes | Yes | <p>The operations to be performed when the scheduled task is triggered.</p> <p>When you set this parameter, you must also enter the unique identifier of the scaling rule.</p> | The parameter value can be up to 200 characters in length. |

Response parameters

Fn::GetAtt

ScheduledTaskId: the ID of the scheduled task. This ID is a globally unique identifier (GUID) generated by the system.

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "ScheduledTask": {
      "Type": "ALIYUN::ESS::ScheduledTask",
      "Properties": {
        "TaskEnabled": "true",
        "Description": "scheduledtask",
        "ScheduledTaskName": "task1",
        "LaunchTime": "2014-08-17T16:52Z",
        "RecurrenceEndTime": "2014-08-17T16:55Z",
        "RecurrenceType": "Daily",
        "RecurrenceValue": "1",
        "ScheduledAction": "ari:acs:ess:cn-qingdao:1344371:scalingRule/cCBpdYdQuBe2cUxOdu6piOk"
      }
    }
  },
  "Outputs": {
    "ScheduledTaskId": {
      "Value": {
        "Fn::GetAtt": [
          "ScheduledTask",
          "ScheduledTaskId"
        ]
      }
    }
  }
}
```

5.5.3. OSS

5.5.3.1. ALIYUN::OSS::Bucket

ALIYUN::OSS::Bucket is used to create an OSS bucket.

Syntax

```
{
  "Type": "ALIYUN::OSS::Bucket",
  "Properties": {
    "AccessControl": String,
    "RefererConfiguration": Map,
    "ServerSideEncryptionConfiguration": Map,
    "CORSConfiguration": Map,
    "Tags": Map,
    "LoggingConfiguration": Map,
    "LifecycleConfiguration": Map,
    "StorageClass": String,
    "DeletionForce": Boolean,
    "WebsiteConfiguration": Map,
    "Policy": Map,
    "BucketName": String
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|------------------------|---------|----------|----------|---|--|
| BucketName | String | Yes | No | The name of the bucket. | <ul style="list-style-type: none"> The name must be 3 to 63 characters in length and can contain lowercase letters, digits, and hyphens (-). It must start and end with a lowercase letter or digit. |
| AccessControl | String | No | No | The access control policy. | Valid values: private, public-read, and public-read-write. |
| CORSConfiguration | Map | No | No | The configuration of cross-origin resource sharing for objects in the bucket. | None |
| LifecycleConfiguration | Map | No | No | The lifecycle configuration for objects in the bucket. | None |
| LoggingConfiguration | Map | No | No | The logging configuration. | None |
| RefererConfiguration | Map | No | No | The hotlinking protection configuration. | None |
| DeletionForce | Boolean | No | No | Specifies whether to forcibly delete objects from an OSS bucket | Valid values: <ul style="list-style-type: none"> true false Default value: false. |

| Property | Type | Required | Editable | Description | Constraint |
|-----------------------------------|--------|----------|----------|---|--|
| WebsiteConfiguration | Map | No | No | The information used to configure the bucket as a static website. | None |
| ServerSideEncryptionConfiguration | Map | No | No | The server-side encryption rules. | None |
| Tags | Map | No | No | The tags of the bucket. Tags exist as key-value pairs. | <ul style="list-style-type: none"> A maximum of 20 tags can be specified. A tag key must be 1 to 64 bytes in length and cannot start with <code>http://</code>, <code>https://</code>, or <code>Aliyun</code>. A tag value can be up to 128 bytes in length and must be encoded in UTF-8. |
| StorageClass | String | No | No | The type of the bucket. | Valid values: Standard, IA, and Archive. |
| Policy | Map | No | No | The bucket policy configuration. | None |

CORSConfiguration syntax

```
"CORSConfiguration": {
  "CORSRule": [
    {
      "AllowedHeader": String,
      "AllowedMethod": List,
      "AllowedOrigin": List,
      "ExposeHeader": List,
      "MaxAgeSeconds": Integer
    }
  ]
}
```

CORSConfiguration properties

| Property | Type | Required | Editable | Description | Constraint |
|----------|------|----------|----------|---|------------|
| CORSRule | List | No | No | The rules that define cross-origin resource sharing of objects in the bucket. | None |

| Property | Type | Required | Editable | Description | Constraint |
|---------------|---------|----------|----------|---|---|
| AllowedHeader | String | No | No | The allowed cross-origin request headers. | Valid values: *, Cache-Control, Content-Language, Content-Type, Expires, Last-Modified, and Pragma. |
| AllowedMethod | List | No | No | The allowed cross-origin request methods. | Valid values: *, GET, PUT, POST, DELETE, and HEAD. |
| AllowedOrigin | List | No | No | The origins from which cross-origin requests are allowed. | None |
| ExposeHeader | List | No | No | The response headers for allowed access requests from applications. | Asterisks (*) cannot be used as wildcard characters. |
| MaxAgeSeconds | Integer | No | No | The period of time that the browser can cache the response of a preflight (OPTIONS) request to a specific resource. | None |

LifecycleConfiguration syntax

```
"LifecycleConfiguration": {
  "Rule": [
    {
      "ID": String,
      "Prefix": String,
      "Status": String,
      "Expiration": Map,
      "AbortMultipartUpload": Map
    }
  ]
}
```

LifecycleConfiguration properties

| Property | Type | Required | Editable | Description | Constraint |
|----------------------|--------|----------|----------|--|--|
| Rule | List | No | No | The lifecycle rule. | None |
| ID | String | No | No | The unique ID of the rule. | The ID can be up to 255 characters in length. When this parameter is empty or not specified, OSS generates a unique rule ID. |
| Prefix | String | No | No | The prefix to which the rule applies. | The rule takes effect only on objects that have a matching prefix. |
| Status | String | No | No | Specifies whether to enable or disable the rule. | Valid values: Enable and Disable. |
| Expiration | Map | No | No | The expiration attributes of the rule for the specified object. | None |
| AbortMultipartUpload | Map | No | No | The expiration attributes of the multipart upload tasks that are not complete. | None |

Expiration syntax

```
"Expiration":{
  "Days": Number,
  "CreatedBeforeDate": String
}
```

Expiration properties

| Property | Type | Required | Editable | Description | Constraint |
|-------------------|--------|----------|----------|--|---|
| Days | Number | No | No | The number of days since the object was last modified after which the rule will take effect. | When the number of days since the object was last modified exceeds the specified number of days, the object is deleted. If you set the Days parameter to 30, objects that were last modified on January 1, 2016 are deleted by the backend application on January 31, 2016. |
| CreatedBeforeDate | String | No | No | The date before which the rule takes effect. | Specify the time in the ISO 8601 standard. The time must be UTC 00:00. Example: 2002-10-11T00:00:00.000Z. |

AbortMultipartUpload syntax

```
"AbortMultipartUpload": {
  "CreatedBeforeDate": String,
  "Days": Number
}
```

AbortMultipartUpload properties

| Property | Type | Required | Editable | Description | Constraint |
|-------------------|--------|----------|----------|--|---|
| Days | Number | No | No | The number of days since the object was last modified after which the rule will take effect. | When the number of days since the object was last modified exceeds the specified number of days, the object is deleted. If you set the Days parameter to 30, objects that were last modified on January 1, 2016 are deleted by the backend application on January 31, 2016. |
| CreatedBeforeDate | String | No | No | The date before which the rule takes effect. | Specify the time in the ISO 8601 standard. The time must be UTC 00:00. Example: 2002-10-11T00:00:00.000Z. |

LoggingConfiguration syntax

```
"LoggingConfiguration": {
  "TargetBucket": String,
  "TargetPrefix": String
}
```

LoggingConfiguration properties

| Property | Type | Required | Editable | Description | Constraint |
|--------------|--------|----------|----------|--|------------|
| TargetBucket | String | No | No | The storage space for storing access logs. | None |
| TargetPrefix | String | No | No | The prefix of the names of saved access log files. | None |

WebsiteConfiguration syntax

```
"WebsiteConfiguration":{
  "IndexDocument": String,
  "ErrorDocument": String
}
```

WebsiteConfiguration properties

| Property | Type | Required | Editable | Description | Constraint |
|---------------|--------|----------|----------|--|------------|
| IndexDocument | String | No | No | The default homepage for a static website. | None |
| ErrorDocument | String | No | No | The default error page for a static website. | None |

RefererConfiguration syntax

```
"RefererConfiguration":{
  "AllowEmptyReferer": String,
  "RefererList": List
}
```

RefererConfiguration properties

| Property | Type | Required | Editable | Description | Constraint |
|-------------------|--------|----------|----------|---|------------|
| AllowEmptyReferer | String | No | No | Specifies whether the Referer field can be left empty in an access request. | None |
| RefererList | List | No | No | The referer whitelist. OSS allows requests whose Referer field values are in the referer whitelist. | None |

ServerSideEncryptionConfiguration syntax

```
"ServerSideEncryptionConfiguration":{
  "KMSMasterKeyID": String,
  "SSEAlgorithm": String
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|----------------|--------|----------|----------|--|--|
| KMSMasterKeyID | String | No | No | The ID of the customer master key. | The key ID is required only when the SSEAlgorithm value is KMS and the specified key is used for encryption. |
| SSEAlgorithm | String | Yes | No | The default server-side encryption method. | Valid values: KMS and AES256. |

Response parameters

Fn::GetAtt

- Name: the bucket name, which must be globally unique.
- DomainName: the public domain name of the specified bucket.
- InternalDomainName: the internal domain name of the specified bucket.

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "Bucket": {
      "Type": "ALIYUN::OSS::Bucket",
      "Properties": {
        "AccessControl": "private",
        "BucketName": "roctest",
        "WebsiteConfiguration": {
          "IndexDocument": "index1.html",
          "ErrorDocument": "error404.html"
        },
        "LoggingConfiguration": {
          "TargetBucket": "cos-mirror",
          "TargetPrefix": "test404"
        }
      },
      "CORSConfiguration": {
        "CORSRule": [
          {
            "AllowedHeader": ["*"],
            "AllowedMethod": ["GET", "PUT"],
            "AllowedOrigin": ["*"],
            "ExposeHeader": ["Date"],
            "MaxAgeSeconds": 3600
          }
        ]
      },
      "LifecycleConfiguration": {
        "Rule": [
          {
            "ID": "deleteRule",
            "Prefix": "test/",
            "Status": "Enabled",
            "Expiration": {
              "Days": 2
            }
          },
          {
            "AbortMultipartUpload": {
              "CreatedBeforeDate": "2014-10-11T00:00:00.000Z"
            }
          }
        ]
      },
      "RefererConfiguration": {
        "AllowEmptyReferer": true,
        "RefererList": ["http://www.aliyun.com", "https://www?.aliyuncs.com"]
      }
    }
  },
  "Outputs": {
    "Name": {
      "Value": {"Fn::GetAtt": ["Bucket", "Name"]}
    },
    "DomainName": {
      "Value": {"Fn::GetAtt": ["Bucket", "DomainName"]}
    }
  }
}
```

5.5.4. RDS

5.5.4.1. ALIYUN::RDS::Account

ALIYUN::RDS::Account is used to create a database management Account.

Statement

```
{
  "Type": "ALIYUN::RDS::Account",
  "Properties": {
    "AccountDescription": String,
    "DBInstanceID": String,
    "AccountPassword": String,
    "AccountType": String,
    "AccountName": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------|--------|----------|----------|---------------------------------------|--|
| AccountDescription | String | Yes | True | The description of the account. | The name must be 2 to 256 characters in length. It can contain digits, letters, underscores (_), and hyphens (-); but must start with a letter. |
| DBInstanceID | String | No | No | The ID of the RDS instance. | None |
| AccountPassword | String | No | No | The password of the database account. | The password must be 8 to 32 characters in length. |
| AccountType | String | Yes | Released | The type of the database account. | Valid values: <ul style="list-style-type: none"> Normal: indicates a standard account. Super: indicates a privileged account. Default value: Normal. |

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------|--------|----------|----------|-----------------------------------|---|
| AccountName | String | No | No | The name of the database account. | Valid values: 1025 to 10000. You cannot use the following commonly used port numbers: 2222, 4500, 4510, 4560, 7505, 9000, 9001, and 9002. |

Response parameters

Fn::GetAtt

AccountName: the name of the database account.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "Account": {
      "Type": "ALIYUN::RDS::Account",
      "Properties": {
        "AccountDescription": {
          "Ref": "AccountDescription"
        },
        "DBInstanceID": [
          "Ref": "DBInstanceID"
        ],
        "AccountPassword": {
          "Ref": "AccountPassword"
        },
        "AccountType": {
          "Ref": "AccountType"
        },
        "AccountName": {
          "Ref": "AccountName"
        }
      }
    },
    "Parameters": {
      "AccountDescription": {
        "Type": "String",
        "Description": "Account remarks.\nIt cannot begin with http:// or https://.\nIt must start with a Chinese character or English letter.\nIt can include Chinese and English characters/letters, underscores (_), hyphens (-), and digits.\nThe length may be 2-256 characters."
      },
      "DBInstanceID": [
        "Type": "String",
        "Description": "RDS instance ID."
      ],
      "AccountPassword": {
        "MinLength": 8,
        "Type": "String",
        "Description": "The account password for the database instance. It may consist of letters, digits, or underlines, with a"
      }
    }
  }
}
```

```
length of 8 to 32 characters.",
  "MaxLength": 32
},
"AccountType": {
  "Default": "Normal",
  "Type": "String",
  "Description": "Privilege type of account.\nNormal: Common privilege.\nSuper: High privilege. And the default value is Normal.\nThis parameter is valid for MySQL 5.5/5.6 only.\nMySQL 5.7, SQL Server 2012/2016, PostgreSQL, and PPAS each can have only one initial account. Other accounts are created by the initial account that has logged on to the database.",
  "AllowedValues": ["Normal", "Super"]
},
"AccountName": {
  "Type": "String",
  "Description": "Account name, which must be unique and meet the following requirements:\nStart with a letter;\nConsist of lower-case letters, digits, and underscores (_);\nContain no more than 16 characters.\nFor other invalid characters, see Forbidden keywords table."
}
},
"Outputs": {
  "AccountName": {
    "Description": "Account name",
    "Value": {
      "Fn::GetAtt": ["Account", "AccountName"]
    }
  }
}
}
```

5.5.4.2. ALIYUN::RDS::AccountPrivilege

ALIYUN::RDS::AccountPrivilege is used to grant database access permissions to accounts.

Statement

```
{
  "Type": "ALIYUN::RDS::AccountPrivilege",
  "Properties": {
    "AccountPrivilege": String,
    "DBInstanceId": String,
    "DBName": String,
    "AccountName": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|------|----------|----------|-------------|------------|
|-----------|------|----------|----------|-------------|------------|

| Parameter | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|--|---|
| AccountPrivilege | String | No | Yes | The permissions of the database account. | Valid values: <ul style="list-style-type: none"> • ReadWrite: has read and write permissions on the database. • ReadOnly: The account has read-only permission on the database. • DDLOnly: The account can run only data definition language (DDL) commands in the database. This is applicable to MySQL and MariaDB. • DMLOnly: The account can run only data manipulation language (DML) commands in the database. This is applicable to MySQL and MariaDB. • DBOwner: The account has full permissions on the database. This is applicable to SQL Server. |
| DBInstanceid | String | No | No | The ID of the RDS instance. | None |
| DBName | String | No | No | The name of the database. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------|--------|----------|----------|--------------------------|---|
| AccountName | String | No | No | The name of the account. | Valid values: 1025 to 10000. You cannot use the following commonly used port numbers: 2222, 4500, 4510, 4560, 7505, 9000, 9001, and 9002. |

Response parameters

Fn::GetAtt

None

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "AccountPrivilege": {
      "Type": "ALIYUN::RDS::AccountPrivilege",
      "Properties": {
        "AccountPrivilege": {
          "Ref": "AccountPrivilege"
        },
        "DBInstanceId": [
          "Ref": "DBInstanceId"
        ],
        "DBName": {
          "Ref": "DBName"
        },
        "AccountName": {
          "Ref": "AccountName"
        }
      }
    },
    "Parameters": {
      "AccountPrivilege": {
        "Type": "String",
        "Description": "RDS account privilege",
        "AllowedValues": ["ReadOnly", "ReadWrite", "DDLOnly", "DMLOnly", "DBOwner"]
      },
      "DBInstanceId": {
        "Type": "String",
        "Description": "RDS instance ID."
      },
      "DBName": {
        "Type": "String",
        "Description": "RDS database name"
      },
      "AccountName": {
        "Type": "String",
        "Description": "RDS account name."
      }
    },
    "Outputs": {}
  }
}
```

5.5.4.3. ALIYUN::RDS::DBInstance

ALIYUN::RDS::DBInstance is used to create an ApsaraDB RDS instance.

Syntax

```
{
  "Type": "ALIYUN::RDS::DBInstance",
  "Properties": {
    "Engine": String,
    "MultiAZ": Boolean,
    "VpcId": String,
    "DBMappings": List,
    "DBInstanceDescription": String,
    "ConnectionMode": String,
    "MasterUsername": String,
    "MasterUserPassword": String,
    "ZoneId": String,
    "DBInstanceNetType": String,
    "DBInstanceStorage": Integer,
    "VSwitchId": String,
    "AllocatePublicConnection": Boolean,
    "EngineVersion": String,
    "PreferredBackupTime": String,
    "DBInstanceClass": String,
    "SecurityIPList": String,
    "BackupRetentionPeriod": Integer,
    "PrivateIpAddress": String,
    "PreferredBackupPeriod": List,
    "PeriodType": String,
    "PayType": String,
    "Period": Integer,
    "ResourceGroupId": String
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|---|--|
| ResourceGroupId | String | No | No | The ID of the resource group. | None |
| Engine | String | Yes | No | The database engine that the instance runs. | Valid values: <ul style="list-style-type: none"> MySQL SQLServer PostgreSQL PPAS |

| Property | Type | Required | Editable | Description | Constraint |
|-------------------|---------|----------|----------|---|---|
| DBInstanceStorage | Integer | Yes | Yes | The storage capacity of the instance. | <ul style="list-style-type: none"> Valid values when Engine is set to MySQL: 5 to 1000. Valid values when Engine is set to SQLServer: 10 to 1000. Valid values when Engine is set to PostgreSQL: 5 to 2000. Valid values when Engine is set to PPAS: 5 to 2000. Unit: GB. <div style="border: 1px solid #ccc; background-color: #e0f2f1; padding: 5px; margin-top: 10px;"> <p> Note This value must be in 5 GB increments.</p> </div> |
| EngineVersion | String | Yes | No | The version of the database engine. | <ul style="list-style-type: none"> Valid values when Engine is set to MySQL: 5.5, 5.6, 5.7, and 8.0. Set the value to 2008r2 when Engine is set to SQLServer. Set the value to 9.4 when Engine is set to PostgreSQL. Set the value to 9.3 when Engine is set to PPAS. |
| DBInstanceClass | String | Yes | Yes | The instance type. | Valid values: <ul style="list-style-type: none"> rds.mys2.large rds.mss1.large rds.pg.s1.small |
| SecurityIPList | String | Yes | Yes | The whitelist of IP addresses that are allowed to access all databases in the instance. | <ul style="list-style-type: none"> Separate multiple IP addresses with commas (,). Each IP address in the whitelist must be unique. A maximum of 1,000 IP addresses can be specified. The 0.0.0.0/0 format is supported. You can specify IP addresses in the 10.23.XX.XX format and CIDR blocks in the 10.23.XX.XX/24 format. In 10.23.XX.XX/24, /24 indicates the length of the prefix in the CIDR block, and the prefix length can range from 1 to 32. 0.0.0.0/0 indicates that no access restriction is applied. |
| MultiAZ | Boolean | No | No | Specifies whether the instance can be deployed across multiple zones. | None |
| VpcId | String | No | No | The ID of the VPC. | None |

| Property | Type | Required | Editable | Description | Constraint |
|--------------------------|---------|----------|----------|--|--|
| DBMappings | List | No | No | The list of one or more databases to be created in the instance. | None |
| DBInstanceDescription | String | No | No | The description of the instance. | <ul style="list-style-type: none"> The description must be 2 to 256 characters in length and can contain letters, digits, underscores (_), and hyphens (-). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code>. |
| ConnectionMode | String | No | No | The connection mode of the instance. | Valid values: <ul style="list-style-type: none"> Performance: standard connection mode Safty: safe connection mode If you do not specify this parameter, the system assigns a connection mode. |
| MasterUsername | String | No | No | The name of the database account. | The name must be unique. The name can be up to 16 characters in length and can contain letters, digits, and underscores (_). |
| MasterUserPassword | String | No | No | The password of the database account. | The password must be 6 to 32 characters in length and can contain letters, digits, and underscores (_). |
| ZoneId | String | No | No | The zone ID of the instance. | None |
| DBInstanceNetType | String | No | No | The network type of the instance. | Default value: Intranet. Valid values: <ul style="list-style-type: none"> Internet Intranet |
| VSwitchId | String | No | No | The ID of the vSwitch in the specified VPC. | None |
| AllocatePublicConnection | Boolean | No | No | Specifies whether to apply for a public endpoint for the instance. | None |

| Property | Type | Required | Editable | Description | Constraint |
|-----------------------|--------|----------|----------|---|---|
| PreferredBackupTime | String | No | No | The backup window. | <ul style="list-style-type: none"> Specify the window in the HH:mmZ-HH:mmZ format. Valid values: 00:00Z-01:00Z, 01:00Z-02:00Z, 02:00Z-03:00Z, 03:00Z-04:00Z, 04:00Z-05:00Z, 05:00Z-06:00Z, 06:00Z-07:00Z, 07:00Z-08:00Z, 08:00Z-09:00Z, 09:00Z-10:00Z, 10:00Z-11:00Z, 11:00Z-12:00Z, 12:00Z-13:00Z, 13:00Z-14:00Z, 14:00Z-15:00Z, 15:00Z-16:00Z, 16:00Z-17:00Z, 17:00Z-18:00Z, 18:00Z-19:00Z, 19:00Z-20:00Z, 20:00Z-21:00Z, 21:00Z-22:00Z, 22:00Z-23:00Z, and 23:00Z-24:00Z. |
| BackupRetentionPeriod | Number | No | No | The number of days for which backup files can be retained. | Valid values: 7 to 30. Unit: days. Default value: 7. |
| PrivateIPAddresses | String | No | No | The private IP address within the CIDR block of the vSwitch. | If you do not specify this parameter, the system allocates a private IP address. |
| PreferredBackupPeriod | List | No | No | The backup cycle. | Valid values: <ul style="list-style-type: none"> Monday Tuesday Wednesday Thursday Friday Saturday Sunday |
| MasterUserType | String | No | No | The type of the database account. | Default value: Normal. Valid values: <ul style="list-style-type: none"> Normal Super |
| Tags | Map | No | Yes | The list of one or more tags. Each tag consists of a tag key and a tag value. | <ul style="list-style-type: none"> The tag key is required and the tag value is optional. Format example: <code>{"key1":"value1","key2":""}</code>. |
| PeriodType | String | No | No | The unit of the subscription period. | Default value: Month. Valid values: <ul style="list-style-type: none"> Month Year |

| Property | Type | Required | Editable | Description | Constraint |
|----------|---------|----------|----------|--|---|
| PayType | String | No | No | The billing method of the instance. | Valid values: <ul style="list-style-type: none"> PostPaid: pay-as-you-go PrePaid: subscription |
| Period | Integer | No | No | The subscription period of the instance. | <ul style="list-style-type: none"> Valid values when PeriodType is set to Year: 1, 2, and 3. Valid values when PeriodType is set to Month: 1, 2, 3, 4, 5, 6, 7, 8, and 9. |

DBMappings syntax

```
"DBMappings": [
{
  "DBDescription": String,
  "CharacterSetName": String,
  "DBName": String
}
]
```

DBMappings properties

| Property | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|---------------------------|--|
| CharacterSetName | String | Yes | No | The character set. | <ul style="list-style-type: none"> Valid values when Engine is set to MySQL: <ul style="list-style-type: none"> utf8 gbk latin1 utf8mb4 (applicable to versions 5.5 and 5.6) Valid values when Engine is set to SQLServer: <ul style="list-style-type: none"> Chinese_PRC_CI_AS Chinese_PRC_CS_AS SQL_Latin1_General_CP1_CI_AS SQL_Latin1_General_CP1_CS_AS Chinese_PRC_BIN |
| DBName | String | Yes | No | The name of the database. | <p>The name must be unique.</p> <p>The name can be up to 64 characters in length and can contain letters, digits, and underscores (_). It must start with a letter.</p> |

| Property | Type | Required | Editable | Description | Constraint |
|---------------|--------|----------|----------|----------------------------------|--|
| DBDescription | String | No | No | The description of the database. | <ul style="list-style-type: none"> The description must be 2 to 256 characters in length and can contain letters, digits, underscores (_), and hyphens (-). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code>. |

Response parameters

Fn::GetAtt

- DBInstanceID: the ID of the instance.
- InnerPort: the internal port of the instance.
- InnerIPAddress: the internal IP address of the instance.
- InnerConnectionString: the internal endpoint of the instance.
- PublicPort: the public port of the instance.
- PublicConnectionString: the public endpoint of the instance.
- PublicIPAddress: the public IP address of the instance.

Examples

The following example demonstrates how to create an ApsaraDB RDS instance in the classic network:

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "Database": {
      "Type": "ALIYUN::RDS::DBInstance",
      "Properties": {
        "Engine": "MySQL",
        "EngineVersion": "5.6",
        "DBInstanceClass": "rds.mysql.t1.small",
        "DBInstanceStorage": 10,
        "DBInstanceNetType": "Internet",
        "SecurityIPList": "0.0.0.0/0",
        "MasterUsername": "A****",
        "DBMappings": [{
          "DBName": "hope",
          "CharacterSetName": "utf8"
        }]
      }
    }
  },
  "Outputs": {
    "DBInstanceId": {
      "Value": {"get_attr": ["DBInstanceId"]}
    },
    "PublicConnectionString": {
      "Value": {"get_attr": ["ConnectionString"]}
    },
    "PublicPort": {
      "Value": {"get_attr": ["Port"]}
    }
  }
}
```

The following example demonstrates how to create an ApsaraDB RDS instance in a VPC:

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "Database": {
      "Type": "ALIYUN::RDS::DBInstance",
      "Properties": {
        "Engine": "MySQL",
        "EngineVersion": "5.6",
        "DBInstanceClass": "rds.mys2.small",
        "DBInstanceStorage": "10",
        "DBInstanceNetType": "Intranet",
        "SecurityIPList": "0.0.0.0/0",
        "VSwitchId": "ttt",
        "VpcId": "myvp*****"
      }
    }
  },
  "Outputs": {
    "DBInstanceId": {
      "Value": {"get_attr": ["DBInstanceId"]}
    },
    "InnerConnectionString": {
      "Value": {"get_attr": ["ConnectionString"]}
    },
    "InnerPort": {
      "Value": {"get_attr": ["Port"]}
    }
  }
}
```

5.5.4.4. ALIYUN::RDS::DBInstanceParameterGroup

ALIYUN::RDS::DBInstanceParameterGroup is used to modify parameters of an ApsaraDB RDS instance.

Syntax

```
{
  "Type": "ALIYUN::RDS::DBInstanceParameterGroup",
  "Properties": {
    "Forcerestart": String,
    "DBInstanceId": String,
    "Parameters": List
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|--------------|--------|----------|----------|--------------------------------------|------------|
| DBInstanceId | String | Yes | No | The ID of the ApsaraDB RDS instance. | None |

| Property | Type | Required | Editable | Description | Constraint |
|--------------|--------|----------|----------|---|---|
| Parameters | List | Yes | No | The list of one or more parameters of the instance. | The parameters and their values must be arranged in the JSON format. The parameter values must be of the string type. Example: {"auto_increment_increment": "1", "character_set_client": "utf8"}. |
| Forcerestart | String | No | No | Specifies whether to forcibly restart the instance. | Default value: false. Valid values: <ul style="list-style-type: none"> • true: The system forcibly restarts the instance. • false: The system does not forcibly restart the instance. |

Response parameters

Fn::GetAtt

None

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "Database": {
      "Type": "ALIYUN::RDS::DBInstance",
      "Properties": {
        "Engine": "MySQL",
        "EngineVersion": "5.6",
        "DBInstanceClass": "rds.mys2.small",
        "DBInstanceStorage": "10",
        "DBInstanceNetType": "Intranet",
        "SecurityIPList": "0.0.0.0/0"
      }
    },
    "DatabaseConfig": {
      "Type": "ALIYUN::RDS::DBInstanceParameterGroup",
      "Properties": {
        "DBInstanceID": {
          "Ref": "Database"
        },
        "Parameters": [
          {
            "Key": "auto_increment_increment",
            "Value": "xxx"
          }
        ]
      }
    }
  },
  "Outputs": {
    "DBInstanceID": {
      "Value": {
        "Fn::GetAtt": [
          "Database",
          "DBInstanceID"
        ]
      }
    }
  }
}
```

5.5.4.5. ALIYUN::RDS::DBInstanceSecurityIps

ALIYUN::RDS::DBInstanceSecurityIps is used to modify the instance whitelist.

Statement

```
{
  "Type": "ALIYUN::RDS::DBInstanceSecurityIps",
  "Properties": {
    "DBInstanceID": String,
    "DBInstanceIPArrayName": String,
    "DBInstanceIPArrayAttribute": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|----------------------------|--------|----------|----------|--|--|
| DBInstanceID | String | No | No | The ID of the RDS instance. | None |
| DBInstanceIPArrayAttribute | String | No | Yes | The attribute of the IP address whitelist. | The console does not display groups labeled with hidden. |
| DBInstanceIPArrayName | String | Yes | Released | The name of the IP address whitelist. | The name can contain only lowercase letters and underscores (_). Default value: Default. |

Response parameters

Fn::GetAtt

SecurityIps: the IP address whitelist after the modification.

Sample request

```

{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "DBInstanceSecurityIps": {
      "Type": "ALIYUN::RDS::DBInstanceSecurityIps",
      "Properties": {
        "DBInstanceIPArrayName": {
          "Ref": "DBInstanceIPArrayName"
        },
        "DBInstanceID": [
          "Ref": "DBInstanceID"
        ],
        "DBInstanceIPArrayAttribute": {
          "Ref": "DBInstanceIPArrayAttribute"
        }
      }
    }
  },
  "Parameters": {
    "DBInstanceIPArrayName": {
      "Type": "String",
      "Description": "Group name of the security ips, only support lower characters and '_'. Advice use a new group name a void effect your database system. If the properties is not specified, it will set to default group, please be careful."
    },
    "DBInstanceID": [
      "Type": "String",
      "Description": "Database instance id to update security ips."
    ],
    "DBInstanceIPArrayAttribute": {
      "Type": "String",
      "Description": "Security ips to add or remove."
    }
  },
  "Outputs": {
    "SecurityIps": {
      "Description": "The security ips of selected database instance.",
      "Value": {
        "Fn::GetAtt": [
          "DBInstanceSecurityIps",
          "SecurityIps"
        ]
      }
    }
  }
}

```

5.5.4.6. ALIYUN::RDS::PrepayDBInstance

ALIYUN::RDS::PrepayDBInstance is used to create a subscription ApsaraDB RDS instance.

Syntax

```

{
  "Type": "ALIYUN::RDS::PrepayDBInstance",
  "Properties": {
    "DBMappings": List,
    "CouponCode": String,
    "MasterUsername": String,
    "PeriodType": String,
    "PayType": String,
    "DBInstanceNetType": String,
    "MasterUserType": String,
    "AutoRenew": Boolean,
    "PreferredBackupTime": String,
    "PrivateIpAddress": String,
    "Engine": String,
    "MultiAZ": Boolean,
    "VpcId": String,
    "ConnectionMode": String,
    "ResourceGroupId": String,
    "VSwitchId": String,
    "BackupRetentionPeriod": Number,
    "Quantity": Number,
    "CommodityCode": String,
    "ZoneId": String,
    "AutoPay": Boolean,
    "EngineVersion": String,
    "DBInstanceClass": String,
    "PreferredBackupPeriod": List,
    "DBInstanceStorage": Integer,
    "DBInstanceDescription": String,
    "Tags": Map,
    "Period": Number,
    "MasterUserPassword": String,
    "AllocatePublicConnection": Boolean
  }
}

```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|--|---|
| ResourceGroupId | String | No | No | The ID of the resource group. | None |
| DBMappings | List | No | No | The list of one or more databases to be created in the instance. | None |
| CouponCode | String | No | No | None | None |
| MasterUsername | String | No | No | The name of the database account. | The name must be unique. The name can be up to 16 characters in length and can contain lowercase letters, digits, and underscores (_). It must start with a lowercase letter. |

| Property | Type | Required | Editable | Description | Constraint |
|---------------------|--------|----------|----------|--|--|
| PeriodType | String | Yes | No | The unit of the subscription period. | Default value: Month. Valid values: <ul style="list-style-type: none"> Year Month |
| DBInstanceNetType | String | No | No | The network type of the instance. | Default value: Intranet. Valid values: <ul style="list-style-type: none"> Internet Intranet |
| MasterUserType | String | No | No | The type of the database account. | Valid values: <ul style="list-style-type: none"> Normal Master |
| PreferredBackupTime | String | No | No | The backup window. | Specify the window in the HH:mmZ-HH:mmZ format. Valid values: 00:00Z-01:00Z, 01:00Z-02:00Z, 02:00Z-03:00Z, 03:00Z-04:00Z, 04:00Z-05:00Z, 05:00Z-06:00Z, 06:00Z-07:00Z, 07:00Z-08:00Z, 08:00Z-09:00Z, 09:00Z-10:00Z, 10:00Z-11:00Z, 11:00Z-12:00Z, 12:00Z-13:00Z, 13:00Z-14:00Z, 14:00Z-15:00Z, 15:00Z-16:00Z, 16:00Z-17:00Z, 17:00Z-18:00Z, 18:00Z-19:00Z, 19:00Z-20:00Z, 20:00Z-21:00Z, 21:00Z-22:00Z, 22:00Z-23:00Z, and 23:00Z-24:00Z. |
| PrivateIpAddress | String | No | No | The private IP address with the CIDR block of the specified vSwitch. | If you do not specify this parameter, the system allocates a private IP address. |
| Engine | String | Yes | No | The database engine that the instance runs. | Valid values: <ul style="list-style-type: none"> MySQL SQLServer PostgreSQL PPAS |

| Property | Type | Required | Editable | Description | Constraint |
|-----------------------|---------|----------|----------|---|--|
| MultiAZ | Boolean | No | No | Specifies whether the instance can be deployed across multiple zones. | None |
| VpcId | String | No | No | The ID of the VPC. | None |
| ConnectionMode | String | No | No | The connection mode of the instance. | Default value: Safty. Valid values: <ul style="list-style-type: none"> • Performance: the standard mode. • Safty: the database proxy mode. If you do not specify this parameter, the system assigns a connection mode. |
| AutoRenew | Boolean | No | No | Specifies whether to enable automatic renewal for the instance. | Valid values: <ul style="list-style-type: none"> • True • False |
| VSwitchId | String | No | No | The ID of the vSwitch in the specified VPC. | None |
| BackupRetentionPeriod | Number | No | No | The number of days for which backup files can be retained. | None |
| Quantity | Number | No | No | The number of instances to be created. | Valid values: 1 to 99. Default value: 1. |
| CommodityCode | String | Yes | No | The commodity code. | Valid values: <ul style="list-style-type: none"> • rds • bards • rords |
| ZoneId | String | No | No | The zone ID of the instance. | None |
| EngineVersion | String | Yes | No | The version of the database engine. | <ul style="list-style-type: none"> • Valid values when Engine is set to MySQL: 5.5 and 5.6. • Set the value to 2008r2 when Engine is set to SQLServer. • Set the value to 9.4 when Engine is set to PostgreSQL. • Set the value to 9.3 when Engine is set to PPAS. |

| Property | Type | Required | Editable | Description | Constraint |
|-----------------------|---------|----------|----------|--|---|
| DBInstanceClass | String | Yes | Yes | The instance type. | Examples: rds.mys2.large, rds.mss1.large, and rds.pg.s1.small. |
| PreferredBackupPeriod | List | No | No | The backup cycle. | Valid values: <ul style="list-style-type: none"> Monday Tuesday Wednesday Thursday Friday Saturday Sunday |
| DBInstanceStorage | Integer | Yes | Yes | The storage capacity of the instance. | <ul style="list-style-type: none"> Valid values when Engine is set to MySQL: 5 to 1000. Valid values when Engine is set to SQLServer: 10 to 1000. Valid values when Engine is set to PostgreSQL or PPAS: 5 to 2000. Unit: GB. This value must be in 5 GB increments. |
| DBInstanceDescription | String | No | No | The description of the instance. | The description must be 2 to 256 characters in length and can contain letters, digits, underscores (_), and hyphens (-). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . |
| Tags | map | No | Yes | The tags of the instance. | None |
| Period | Number | Yes | No | The subscription period of the instance. | <ul style="list-style-type: none"> Valid values when PeriodType is set to Month: 1, 2, 3, 4, 5, 6, 7, 8, and 9. Valid values when PeriodType is set to Year: 1, 2, and 3. |
| MasterUserPassword | String | No | No | The password of the database account. | The password must be 6 to 32 characters in length and can contain letters, digits, and underscores (_). |

| Property | Type | Required | Editable | Description | Constraint |
|--------------------------|---------|----------|----------|--|--|
| AllocatePublicConnection | Boolean | No | No | Specifies whether to apply for a public endpoint for the instance. | None |
| PayType | String | No | No | The billing method of the instance. | Valid values: <ul style="list-style-type: none"> Postpaid: pay-as-you-go Prepaid: subscription |
| AutoPay | Boolean | No | No | Specifies whether to enable automatic payment for the instance. | Default value: False. Valid values: <ul style="list-style-type: none"> True False |

DBMappings syntax

```
"DBMappings": [
  {
    "DBDescription": String,
    "CharacterSetName": String,
    "DBName": String
  }
]
```

DBMappings properties

| Property | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|----------------------------------|--|
| DBDescription | String | No | No | The description of the database. | The description must be 2 to 256 characters in length and can contain letters, digits, underscores (_), and hyphens (-). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . |
| CharacterSetName | String | Yes | No | The character set. | <ul style="list-style-type: none"> Valid values when Engine is set to MySQL: utf8, gbk, latin1, and utf8mb4 (applicable to versions 5.5 and 5.6). Valid values when Engine is set to SQLServer: Chinese_PRC_CI_AS, Chinese_PRC_CS_AS, SQL_Latin1_General_CP1_CI_AS, SQL_Latin1_General_CP1_CS_AS, and Chinese_PRC_BIN. |

| Property | Type | Required | Editable | Description | Constraint |
|----------|--------|----------|----------|---------------------------|---|
| DBName | String | Yes | No | The name of the database. | The name must be unique. It can be up to 64 characters in length and can contain lowercase letters, digits, and underscores (_). It must start with a lowercase letter. |

Response parameters

Fn::GetAtt

- InnerPort: the internal port of the instance.
- OrderId: the order ID of the instance.
- PublicConnectionString: the public endpoint of the instance.
- InnerIPAddress: the internal IP address of the instance.
- DBInstanceCid: the ID of the instance.
- PublicIPAddress: the public IP address of the instance.
- PublicPort: the public port of the instance.
- InnerConnectionString: the internal endpoint of the instance.

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "PeriodType": {
      "Type": "String",
      "Description": "Charge period for created instances.",
      "AllowedValues": [
        "Month",
        "Year"
      ],
      "Default": "Month"
    },
    "PrivateIpAddress": {
      "Type": "String",
      "Description": "The private ip for created instance."
    },
    "DBInstanceNetType": {
      "Type": "String",
      "Description": "Database instance net type, default is Intranet.Internet for public access, Intranet for private access.",
      "AllowedValues": [
        "Internet",
        "Intranet"
      ],
      "Default": "Intranet"
    },
    "AutoRenew": {
      "Type": "Boolean",
      "Description": "Auto renew the prepay instance. If the period type is by year, it will renew by year, else it will renew by month.",
      "AllowedValues": [
        "True",
        "true",
        "False",
        "false"
      ]
    }
  }
}
```

```

    ],
    "Default": false
  },
  "PreferredBackupPeriod": {
    "Type": "CommaDelimitedList",
    "Description": "Automate backups cycle if automated backups are enabled.",
    "AllowedValues": [
      "Monday",
      "Tuesday",
      "Wednesday",
      "Thursday",
      "Friday",
      "Saturday",
      "Sunday"
    ]
  },
  "DBInstanceStorage": {
    "Type": "Number",
    "Description": "Database instance storage size. mysql is [5,1000]. sql server 2008r2 is [10,1000], sql server 2012/2012_web/2016-web is [20,1000]. PostgreSQL and PPAS is [5,2000]. Increased every 5 GB, Unit in GB"
  },
  "CommodityCode": {
    "Type": "String",
    "Description": "The CommodityCode of the order.",
    "AllowedValues": [
      "rds",
      "bards",
      "rords"
    ],
    "Default": "rds"
  },
  "DBMappings": {
    "Type": "CommaDelimitedList",
    "Description": "Database mappings to attach to db instance."
  },
  "MultiAZ": {
    "Type": "Boolean",
    "Description": "Specifies if the database instance is a multiple Availability Zone deployment. ",
    "AllowedValues": [
      "True",
      "true",
      "False",
      "false"
    ],
    "Default": false
  },
  "Engine": {
    "Type": "String",
    "Description": "Database instance engine type. Support MySQL/SQLServer/PostgreSQL/PPAS now.",
    "AllowedValues": [
      "MySQL",
      "SQLServer",
      "PostgreSQL",
      "PPAS"
    ]
  },
  "DBInstanceDescription": {
    "Type": "String",
    "Description": "Description of created database instance."
  }
}

```

```

    },
    "Tags": {
      "Type": "Json",
      "Description": "The tags of an instance.\nYou should input the information of the tag with the format of the Key-Value, such as {\"key1\": \"value1\", \"key2\": \"value2\", ... \"key5\": \"value5\"}.\nAt most 5 tags can be specified.\nKey\n\t can be up to 64 characters in length.\nCannot begin with aliyun.\nCannot begin with http:// or https://.\nCannot be a null string.\nValue\n\t can be up to 128 characters in length.\nCannot begin with aliyun.\nCannot begin with http:// or https://.\nCan be a null string."
    },
    "EngineVersion": {
      "Type": "String",
      "Description": "Database instance version of the relative engine type.Support MySQL: 5.5/5.6/5.7; SQLServer: 2008r2, 2012, 2012_web, 2012_std_ha, 2012_ent_ha, 2016_web, 2016_std_ha, 2016_ent_ha; PostgreSQL:9.4; PPAS: 9.3.",
      "AllowedValues": [
        "5.5",
        "5.6",
        "5.7",
        "2008r2",
        "2012",
        "2012_web",
        "2012_std_ha",
        "2012_ent_ha",
        "2016_web",
        "2016_std_ha",
        "2016_ent_ha",
        "9.4",
        "9.3"
      ]
    },
    "Zoneld": {
      "Type": "String",
      "Description": "selected zone to create database instance. You cannot set the Zoneld parameter if the MultiAZ parameter is set to true."
    },
    "DBInstanceClass": {
      "Type": "String",
      "Description": "Database instance type. Refer the RDS database instance type reference, such as 'rds.mys2.large', 'rds.mss1.large', 'rds.pg.s1.small' etc"
    },
    "AllocatePublicConnection": {
      "Type": "Boolean",
      "Description": "If true, allocate public connection automate.",
      "AllowedValues": [
        "True",
        "true",
        "False",
        "false"
      ]
    },
    "PreferredBackupTime": {
      "Type": "String",
      "Description": "The daily time range during which automated backups are created if automated backups are enabled."
    },
    "AllowedValues": [
      "00:00Z-01:00Z",
      "01:00Z-02:00Z",
      "02:00Z-03:00Z",
      "03:00Z-04:00Z",
      "04:00Z-05:00Z",
      "05:00Z-06:00Z",
      "06:00Z-07:00Z"
    ]
  }
}

```



```

},
"VpcId": {
  "Type": "String",
  "Description": "The VPC id of created database instance. For VPC network, the property is required."
},
"MasterUsername": {
  "Type": "String",
  "Description": "The master user name for the database instance. "
},
"ConnectionMode": {
  "Type": "String",
  "Description": "Connection Mode for database instance,support 'Performance' and 'Safty' mode. Default is RDS system assigns. ",
  "AllowedValues": [
    "Performance",
    "Safty"
  ]
},
"BackupRetentionPeriod": {
  "Type": "Number",
  "Description": "The number of days for which automatic DB backups are retained.",
  "MinValue": 7,
  "MaxValue": 30,
  "Default": 7
},
"Resources": {
  "PrepayDBInstance": {
    "Type": "ALIYUN::RDS::PrepayDBInstance",
    "Properties": {
      "PeriodType": {
        "Ref": "PeriodType"
      },
      "PrivateIpAddress": {
        "Ref": "PrivateIpAddress"
      },
      "DBInstanceNetType": {
        "Ref": "DBInstanceNetType"
      },
      "AutoRenew": {
        "Ref": "AutoRenew"
      },
      "PreferredBackupPeriod": {
        "Fn::Split": [
          ",",
          {
            "Ref": "PreferredBackupPeriod"
          },
          {
            "Ref": "PreferredBackupPeriod"
          }
        ]
      },
      "DBInstanceStorage": {
        "Ref": "DBInstanceStorage"
      },
      "CommodityCode": {
        "Ref": "CommodityCode"
      },
      "DBMappings": {
        "Fn::Split": [

```



```

    },
    "ConnectionMode": {
      "Ref": "ConnectionMode"
    },
    },
    "BackupRetentionPeriod": {
      "Ref": "BackupRetentionPeriod"
    }
  }
}
},
"Outputs": {
  "InnerConnectionString": {
    "Description": "DB instance connection url by Intranet.",
    "Value": {
      "Fn::GetAtt": [
        "PrepayDBInstance",
        "InnerConnectionString"
      ]
    }
  },
  "DBInstancelid": {
    "Description": "The instance id of created database instance.",
    "Value": {
      "Fn::GetAtt": [
        "PrepayDBInstance",
        "DBInstancelid"
      ]
    }
  },
  "InnerIPAddress": {
    "Description": "IP Address for created DB instance of Intranet.",
    "Value": {
      "Fn::GetAtt": [
        "PrepayDBInstance",
        "InnerIPAddress"
      ]
    }
  },
  "PublicConnectionString": {
    "Description": "DB instance connection url by Internet.",
    "Value": {
      "Fn::GetAtt": [
        "PrepayDBInstance",
        "PublicConnectionString"
      ]
    }
  },
  "PublicIPAddress": {
    "Description": "IP Address for created DB instance of Internet.",
    "Value": {
      "Fn::GetAtt": [
        "PrepayDBInstance",
        "PublicIPAddress"
      ]
    }
  },
  "OrderId": {
    "Description": "The order id list of created instance.",
    "Value": {
      "Fn::GetAtt": [
        "PrepayDBInstance"
      ]
    }
  }
}

```


| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|--|------------|
| Count | Number | No. | True | The total number of messages to be received. | None |

Response parameters

Fn::GetAtt

- **Data:** A JSON-serialized dictionary that contains the signal Data after the most recent stack creation or update.
- **LastData:** a JSON-serialized dictionary that contains the signal data before the most recent stack update.
- **JoinedErrorData:** a string consisting of the ErrorData signal data.
- **JoinedLastErrorData:** a string consisting of the LastErrorData signal data.
- **ErrorData:** a JSON-serialized dictionary that contains the error signal data after the most recent stack creation or update.
- **Lasterprotodata:** a JSON-serialized dictionary that contains the error signal data before the most recent stack update.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "WaitCondition": {
      "Type": "ALIYUN::ROS::WaitCondition",
      "Properties": {
        "Handle": {
          "Ref": "WaitConHandle"
        },
        "Timeout": 5,
        "Count": 2
      }
    },
    "WaitConHandle": {
      "Type": "ALIYUN::ROS::WaitConditionHandle"
    }
  },
  "Outputs": {
    "CurlCli": {
      "Value": {
        "Fn::GetAtt": [
          "WaitConHandle",
          "CurlCli"
        ]
      }
    },
    "Data": {
      "Value": {
        "Fn::GetAtt": [
          "WaitCondition",
          "Data"
        ]
      }
    }
  }
}
```

5.5.5.2. ALIYUN::ROS::WaitConditionHandle

ALIYUN::ROS::WaitConditionHandle is used to create an instance that sends and receives messages during UserData execution.

Statement

```
{
  "Type": "ALIYUN::ROS::WaitConditionHandle",
  "Properties": {
    "Count": Integer,
    "Mode": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|---------|----------|----------|---|---|
| Count | Integer | No. | True | The total number of messages to be received. | Default value: -1. |
| Mode | String | Yes | True | If you set this parameter to Increment, all previous signals will be updated before they are deleted. If you set this parameter to Full, no previous signals will be deleted unless the Count parameter is specified. | Valid values: <ul style="list-style-type: none"> Increment Full Default value: Full. |

Response parameters

Fn::GetAtt

- **CurlCli:** A curl Command is generated by the resource. You can use the command to send the UserData execution result or status to Resource Orchestration Service.
- **WindowsCurlCli:** provides Windows with cURL CLI command prefixes and sends a message indicating that the execution is completed or failed. Windows does not support the curl command. Therefore, you must install curl.exe and add it to PATH. You can add `--data-binary "{\"status\": \" success\"}` to indicate success, or by adding `--data-binary "{\"status\": \" failure\"}` to indicate failure.
- **PowerShellCurlCli:** provides PowerShell with cURL CLI command prefixes and sends a message indicating that the execution is completed or failed. Because this cmdlet was introduced in PowerShell 3.0, make sure that the PowerShell version meets this constraint. By `$PSVersionTable.PSVersion` displays the version. You can add `-Body '{"status": "success"}` to indicate success, or by adding `-Body '{"status": "failure"}` to indicate failure.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "Mode": {
      "Type": "String",
      "Description": "If set to Increment, all old signals will be deleted before update. In this mode, WaitCondition.Count should reference an incremental value instead of a full value, such as ScalingGroupEnable.ScalingRuleAriseExecuteResultNumberOfAddedInstances.\n\nIf set to Full, no old signal will be deleted unless Count is set. In this mode, WaitCondition.Count should reference a full value, such as the same value with InstanceGroup.MaxAmount. It is recommended to use this mode with Count.\n\nDefault to Full.",
      "AllowedValues": [
        "Increment",
        "Full"
      ],
      "Default": "Full"
    },
    "Count": {
      "Type": "Number",
      "Description": "There are 3 preconditions that make Count taking effect:\n1.Mode is set to Full.\n2.Count >= 0.\n3.The id of signal is not specified. If so, it will be a self-increasing integer started from 1. For example, the id of the first signal is
```

id of signal is not specified, it will be a self-increasing integer started from 1. For example, the id of the first signal is 1, the id of the second signal is 2, and so on.
 If Count takes effect, signals with id > Count will be deleted before update.
 The default value is -1, which means no effect.
 It is recommended to quote the same value with WaitCondition.Count.",

```

    "Default": -1
  }
},
"Resources": {
  "WaitConditionHandle": {
    "Type": "ALIYUN::ROS::WaitConditionHandle",
    "Properties": {
      "Mode": {
        "Ref": "Mode"
      },
      "Count": {
        "Ref": "Count"
      }
    }
  }
},
"Outputs": {
  "CurlCli": {
    "Description": "Convenience attribute, provides curl CLI command prefix, which can be used for signalling handle completion or failure. You can signal success by adding --data-binary '{\"status\": \"SUCCESS\"}', or signal failure by adding -data-binary '{\"status\": \"FAILURE\"}',",
    "Value": {
      "Fn::GetAtt": [
        "WaitConditionHandle",
        "CurlCli"
      ]
    }
  },
  "WindowsCurlCli": {
    "Description": "Convenience attribute, provides curl CLI command prefix for Windows, which can be used for signalling handle completion or failure. As Windows does not support curl command, you need to install curl.exe and add it to PATH first. You can signal success by adding --data-binary '{\"status\": \"SUCCESS\"}', or signal failure by adding -data-binary '{\"status\": \"FAILURE\"}',",
    "Value": {
      "Fn::GetAtt": [
        "WaitConditionHandle",
        "WindowsCurlCli"
      ]
    }
  },
  "PowerShellCurlCli": {
    "Description": "Convenience attribute, provides curl CLI command prefix for PowerShell, which can be used for signalling handle completion or failure. As this cmdlet was introduced in PowerShell 3.0, ensure the version of PowerShell satisfies the constraint. (Show the version via $PSVersionTable.PSVersion.) You can signal success by adding -Body '{\"status\": \"SUCCESS\"}', or signal failure by adding -Body '{\"status\": \"FAILURE\"}',",
    "Value": {
      "Fn::GetAtt": [
        "WaitConditionHandle",
        "PowerShellCurlCli"
      ]
    }
  }
}
}
}
}
}

```

5.5.5.3. ALIYUN::ROS::Stack

ALIYUN::ROS::Stack is used to create a nested stack. You can have a maximum of five nested levels.

ALIYUN::ROS::Stack is used in a top-level template to nest stacks as resources.

You can add output values from a nested stack contained within the template. You can use Fn::GetAtt together with the logical name of the nested stack and the output name in the Outputs.NestedStackOutputName format.

 **Note** We recommend that you run an update to the Nested stack from the parent stack.

When you apply a template change to update a top-level stack, ROS updates the top-level stack and initiates an update to its nested stacks. Resource orchestration service (ROS) updates resources that have been modified in the nested stack, but does not update resources that have not been modified in the nested stack.

Statement

```
{
  "Type": "ALIYUN::ROS::Stack",
  "Properties": {
    "TemplateURL": String,
    "TimeoutMins": Number,
    "Parameters": Map
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|------|----------|----------|-------------|------------|
|-----------|------|----------|----------|-------------|------------|

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------|--------|----------|----------|--|---|
| TemplateURL | String | No | Yes | <p>The URL of the file containing the template body. The template file can be up to 524,288 bytes in size.</p> <p>The URL must point to a template located on the http or https Web server or Alibaba Cloud OSS bucket.</p> <p>For example: <code>oss://ros/template/demo</code> <code>oss://ros/template/demo?RegionId=cn-hangzhou</code></p> <p>If the region of the OSS bucket is not specified, the RegionId of the stack is used.</p> | The URL can be up to 1,024 bytes in length. |
| TimeoutMins | Number | No. | True | The length of time that ROS will wait for the nested stack to be created or updated. | <p>Unit: minutes.</p> <p>Default value: 60.</p> |

| Parameter | Type | Required | Editable | Description | Constraint |
|------------|------|----------|----------|--|------------|
| Parameters | Map | No. | True | A set of value pairs that represent the parameters passed to ROS when this Nested stack is created. Each parameter has a name corresponding to a parameter defined in the embedded template and the value to which you want to set the parameter. This parameter is required if the nested stack needs input parameters. | None |

Response parameters

Fn::GetAtt

You can use the following code to obtain the output of the nested stack:

```
{
  "Fn::GetAtt": [
    "<nested_stack>",
    "Outputs.<nested_stack_output_name>"
  ]
}
```

When you use `Ref` to reference resources in a nested stack, the Alibaba Cloud Resource Name (ARN) of the nested stack is returned. Example: `arn:acs:ros::cn-hangzhou:12345****:stacks/test-nested-stack-Demo-jzkyq7mn2ykj/e71c1e04-1a57-46fc-b9a4-cf7ce0d3****`

Examples

- The following code provides an example of how to create a VPC, a VSwitch, and a security group in a nested stack and save the output results to the `oss://ros/template/vpc.txt` directory:

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Description": "One VPC, VSwitch, security group.",
  "Parameters": {
    "ZoneId": {
      "Type": "String",
      "Description": "The available zone"
    },
    "SecurityGroupName": {
      "Type": "String",
      "Description": "The security group name",
      "Default": "my-sg-name"
    }
  }
}
```

```

},
"VpcName": {
  "Type": "String",
  "Description": "The VPC name",
  "MinLength": 2,
  "MaxLength": 128,
  "ConstraintDescription": "[2, 128] English or Chinese letters",
  "Default": "my-vpc-name"
},
"VpcCidrBlock": {
  "Type": "String",
  "AllowedValues": [
    "192.168.0.0/16",
    "172.16.0.0/12",
    "10.0.0.0/8"
  ],
  "Default": "10.0.0.0/8"
},
"VSwitchCidrBlock": {
  "Type": "String",
  "Description": "The VSwitch subnet which must be within VPC",
  "Default": "10.0.10.0/24"
},
"UpdateVersion": {
  "Type": "Number",
  "Default": 0
}
},
"Resources": {
  "Vpc": {
    "Type": "ALIYUN::ECS::VPC",
    "Properties": {
      "CidrBlock": {
        "Ref": "VpcCidrBlock"
      },
      "VpcName": {
        "Ref": "VpcName"
      }
    }
  },
  "VSwitch": {
    "Type": "ALIYUN::ECS::VSwitch",
    "Properties": {
      "CidrBlock": {
        "Ref": "VSwitchCidrBlock"
      },
      "ZoneId": {
        "Ref": "ZoneId"
      },
      "VpcId": {
        "Fn::GetAtt": [
          "Vpc",
          "VpcId"
        ]
      }
    }
  },
  "SecurityGroup": {
    "Type": "ALIYUN::ECS::SecurityGroup",
    "Properties": {

```

```

    "SecurityGroupName": {
      "Ref": "SecurityGroupName"
    },
    "VpcId": {
      "Ref": "Vpc"
    }
  },
  "WaitConditionHandle": {
    "Type": "ALIYUN::ROS::WaitConditionHandle",
    "Properties": {
      "UpdateVersion": {
        "Ref": "UpdateVersion"
      }
    }
  },
  "Outputs": {
    "SecurityGroupId": {
      "Value": {
        "Fn::GetAtt": [
          "SecurityGroup",
          "SecurityGroupId"
        ]
      }
    },
    "VpcId": {
      "Value": {
        "Fn::GetAtt": [
          "Vpc",
          "VpcId"
        ]
      }
    },
    "VSwitchId": {
      "Value": {
        "Fn::GetAtt": [
          "VSwitch",
          "VSwitchId"
        ]
      }
    }
  }
}

```

- The following code provides an example of a top-level stack:

```

{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Description": "One ECS instance.",
  "Parameters": {
    "ImageId": {
      "Default": "centos_7",
      "Type": "String",
      "Description": "Image Id, represents the image resource to startup the ECS instance"
    },
    "InstanceType": {
      "Type": "String",
      "Description": "The ECS instance type,",
      "Default": "ecs.xn4.small"
    }
  }
}

```

```

},
"ZoneId": {
  "Type": "String",
  "Description": "The available zone "
},
"InstanceChargeType": {
  "Type": "String",
  "AllowedValues": [
    "PrePaid",
    "PostPaid"
  ],
  "Default": "PostPaid",
  "Description": "The instance charge type"
},
"SecurityGroupName": {
  "Type": "String",
  "Description": "The security group name",
  "Default": "my-sg-name"
},
"NetworkInterfaceName": {
  "Type": "String",
  "Description": "The Network interface name",
  "Default": "my-eni-name"
},
"VpcName": {
  "Type": "String",
  "Description": "The VPC name",
  "MinLength": 2,
  "MaxLength": 128,
  "ConstraintDescription": "[2, 128] English or Chinese letters",
  "Default": "my-vpc-name"
},
"IoOptimized": {
  "AllowedValues": [
    "none",
    "optimized"
  ],
  "Description": "IO optimized, optimized is for the IO optimized instance type",
  "Type": "String",
  "Default": "optimized"
},
"SystemDiskCategory": {
  "AllowedValues": [
    "cloud",
    "cloud_efficiency",
    "cloud_ssd"
  ],
  "Description": "System disk category: average cloud disk(cloud), efficient cloud disk(cloud_efficiency) or SSD cloud disk(cloud_ssd)",
  "Type": "String",
  "Default": "cloud_ssd"
},
"VpcCidrBlock": {
  "Type": "String",
  "AllowedValues": [
    "192.168.0.0/16",
    "172.16.0.0/12",
    "10.0.0.0/8"
  ],
  "Default": "10.0.0.0/8"
}

```

```

    },
    "VSwitchCidrBlock": {
      "Type": "String",
      "Description": "The VSwitch subnet which must be within VPC",
      "Default": "10.0.10.0/24"
    },
    "UpdateVersion": {
      "Type": "Number",
      "Default": 0
    }
  },
  "Resources": {
    "NetworkStack": {
      "Type": "ALIYUN::ROS::Stack",
      "Properties": {
        "TemplateURL": "oss://ros/template/vpc.txt",
        "TimeoutMins": 5,
        "Parameters": {
          "ZoneId": {
            "Ref": "ZoneId"
          },
          "SecurityGroupName": {
            "Ref": "SecurityGroupName"
          },
          "VpcName": {
            "Ref": "VpcName"
          },
          "VpcCidrBlock": {
            "Ref": "VpcCidrBlock"
          },
          "VSwitchCidrBlock": {
            "Ref": "VSwitchCidrBlock"
          }
        },
        "UpdateVersion": {
          "Ref": "UpdateVersion"
        }
      }
    },
    "WebServer": {
      "Type": "ALIYUN::ECS::Instance",
      "Properties": {
        "ImageId": {
          "Ref": "ImageId"
        },
        "InstanceType": {
          "Ref": "InstanceType"
        },
        "InstanceChargeType": {
          "Ref": "InstanceChargeType"
        },
        "SecurityGroupId": {
          "Fn::GetAtt": [
            "NetworkStack",
            "Outputs.SecurityGroupId"
          ]
        }
      }
    },
    "VpcId": {
      "Fn::GetAtt": [
        "NetworkStack",

```

```
    "Outputs.VpcId"
  ]
},
"VSwitchId": {
  "Fn::GetAtt": [
    "NetworkStack",
    "Outputs.VSwitchId"
  ]
},
"IOOptimized": {
  "Ref": "IOOptimized"
},
"ZoneId": {
  "Ref": "ZoneId"
},
"SystemDisk_Category": {
  "Ref": "SystemDiskCategory"
},
"DiskMappings": [
  {
    "Category": "cloud_ssd",
    "Size": 20
  }
]
}
},
"Outputs": {
  "InstanceId": {
    "Value": {
      "Fn::GetAtt": [
        "WebServer",
        "InstanceId"
      ]
    }
  },
  "PublicIp": {
    "Value": {
      "Fn::GetAtt": [
        "WebServer",
        "PublicIp"
      ]
    }
  },
  "SecurityGroupId": {
    "Value": {
      "Fn::GetAtt": [
        "NetworkStack",
        "Outputs.SecurityGroupId"
      ]
    }
  },
  "VpcId": {
    "Value": {
      "Fn::GetAtt": [
        "NetworkStack",
        "Outputs.VpcId"
      ]
    }
  }
},
}
```

```
"VSwitchId": {
  "Value": {
    "Fn::GetAtt": [
      "NetworkStack",
      "Outputs.VSwitchId"
    ]
  }
},
"NetworkStackArn": {
  "Value": {
    "Ref": "NetworkStack"
  }
}
}
```

5.5.6. SLB

5.5.6.1. ALIYUN::SLB::AccessControl

ALIYUN::SLB::AccessControl is used to create an access control list (ACL).

Syntax

```
{
  "Type": "ALIYUN::SLB::AccessControl",
  "Properties": {
    "AddressIPVersion": String,
    "AclName": String,
    "AclEntrys": List
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|--------------------------------|--|
| AddressIPVersion | String | No | No | The Internet protocol version. | Valid values: ipv4 and ipv6. |
| AclName | String | Yes | Yes | The name of the ACL. | None |
| AclEntrys | List | No | No | The list of ACL entries. | A list can contain up to 50 ACL entries. |

AclEntrys syntax

```
"AclEntrys": [
  {
    "comment": String,
    "entry": String
  }
]
```

AclEntrys properties

| Property | Type | Required | Editable | Description | Constraint |
|----------|--------|----------|----------|---|------------|
| comment | String | No | No | The comments on ACL entries. | None |
| entry | String | Yes | No | The authorized IP addresses or CIDR blocks. | None |

Response parameters

Fn::GetAtt

AclId: the ID of the ACL.

Examples

Resource usage example

```

{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "AccessControl": {
      "Type": "ALIYUN::SLB::AccessControl",
      "Properties": {
        "AddressIPVersion": {
          "Ref": "AddressIPVersion"
        },
        "AclName": {
          "Ref": "AclName"
        },
        "AclEntries": {
          "Fn::Split": [",", {
            "Ref": "AclEntries"
          }], {
            "Ref": "AclEntries"
          }
        }
      }
    },
    "Parameters": {
      "AddressIPVersion": {
        "Type": "String",
        "Description": "IP version. Could be \"ipv4\" or \"ipv6\".",
        "AllowedValues": ["ipv4", "ipv6"]
      },
      "AclName": {
        "Type": "String",
        "Description": "The name of the access control list."
      },
      "AclEntries": {
        "Type": "CommaDelimitedList",
        "Description": "A list of acl entries. Each entry can be IP addresses or CIDR blocks. Max length: 50.",
        "MaxLength": 50
      }
    },
    "Outputs": {
      "AclId": {
        "Description": "The ID of the access control list.",
        "Value": {
          "Fn::GetAtt": ["AccessControl", "AclId"]
        }
      }
    }
  }
}

```

Example of combined use of SLB-related resources

```

{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "LoadBalancer": {
      "Type": "ALIYUN::SLB::LoadBalancer",
      "Properties": {
        "LoadBalancerName": "slb-with-listener-and-acl",
        "AddressType": "internet",
        "InternetChargeType": "paybybandwidth",

```

```

    "Bandwidth": 10,
    "VpcId": "vpc-xxxxxxxxxxxxxxxxxxxx",
    "VSwitchId": "vsw-xxxxxxxxxxxxxxxxxxxx"
  }
},
"ACL": {
  "Type": "ALIYUN::SLB::AccessControl",
  "Properties": {
    "AclName": "acl-for-listener",
    "AddressIPVersion": "ipv4",
    "AclEntries": [
      {
        "entry": "192.168.x.x"
      },
      {
        "entry": "10.0.x.x/24",
        "comment": "just comment"
      }
    ]
  }
},
"CreateListener": {
  "Type": "ALIYUN::SLB::Listener",
  "Properties": {
    "LoadBalancerId": {
      "Ref": "LoadBalancer"
    },
    "ListenerPort": "80",
    "BackendServerPort": 8080,
    "Bandwidth": 1,
    "Protocol": "http",
    "HealthCheck": {
      "HealthyThreshold": 3,
      "UnhealthyThreshold": 3,
      "Interval": 2,
      "Timeout": 5
    },
    "Scheduler": "wrr",
    "RequestTimeout": 179,
    "IdleTimeout": 59,
    "AclId": {
      "Ref": "ACL"
    },
    "AclStatus": "on",
    "AclType": "white"
  }
}
},
"Outputs": {
  "LoadBalanceDetails": {
    "Value": {
      "Fn::GetAtt": [
        "LoadBalancer",
        "Listeners"
      ]
    }
  }
}
}
}

```

5.5.6.2. ALIYUN::SLB::BackendServerAttachment

ALIYUN::SLB::BackendServerAttachment is used to add backend servers.

Statement

```
{
  "Type": "ALIYUN::SLB::BackendServerAttachment",
  "Properties": {
    "LoadBalancerId": String,
    "BackendServers": List,
    "BackendServerList": List,
    "BackendServerWeightList": List
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------------|--------|----------|----------|-------------------------------------|---|
| LoadBalancerId | String | No | No | The unique ID of the SLB instance. | None |
| BackendServerList | List | No. | True | The list of backend servers to add. | You can call this operation with LoadBalancerId and BackendServerWeightList. Separate ECS instance IDs with commas (.). This parameter is ignored when the BackendServers parameter is specified. |

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------------------|------|----------|----------|--|---|
| BackendServerWeightList | List | No. | True | The weights of the ECS instances in the BackendServerList, which are specified in order. | If this parameter is not specified, the weight of all ECS instances included in the BackendServerList is 100. When the BackendServerWeightList length is less than BackendServerList, the last value in the BackendServerWeightList is used to weight the remaining ECS instances in the BackendServerList. |
| BackendServers | List | No. | True | The list of backend servers to add. | Only backend servers in the running state can be attached to the SLB instance. |

BackendServers syntax

```
"BackendServers": [
  {
    "ServerId": String,
    "Weight": Integer
  }
]
```

BackendServers properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|---------|----------|----------|---|--|
| ServerId | String | No | Yes | The ID of the ECS instance that acts as a backend server. | The ECS instance must be in the Running state. |
| Weight | Integer | Retained | Yes | The weight of the ECS instance in the SLB instance. | Valid values: 0 to 100. Default value: 100. |

Response parameters

Fn::GetAtt

- BackendServers: the backend servers added to the SLB instance.

- LoadBalancerId: the ID of the SLB instance.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "Attachment2": {
      "Type": "ALIYUN::SLB::BackendServerAttachment",
      "Properties": {
        "LoadBalancerId": "15187200816-cn-beijing-btc-****",
        "BackendServerList": [
          "i-25o0m****",
          "i-25zsk****"
        ],
        "BackendServerWeightList": [
          "20",
          "100"
        ]
      }
    }
  }
}
```

5.5.6.3. ALIYUN::SLB::BackendServerToVServerGroupAddition

ALIYUN::SLB::BackendServerToVServerGroupAddition is used to add backend servers to an existing VServer group.

Statement

```
{
  "Type": "ALIYUN::SLB::BackendServerToVServerGroupAddition",
  "Properties": {
    "BackendServers": List,
    "VServerGroupId": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|----------------|--------|----------|----------|--|------------|
| VServerGroupId | String | No | No | The ID of the VServer group. | None |
| BackendServers | List | Retained | Yes | The list of ECS instances to be added. | None |

BackendServers syntax

```
"BackendServers": [
  {
    "ServerId": String,
    "Port": Integer,
    "Weight": Integer
  }
]
```

BackendServers properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|---------|----------|----------|---|---------------------------|
| ServerId | String | No | Yes | The ID of the ECS instance that acts as a backend server. | None |
| Port | Integer | Retained | Yes | The ECS port number that is listened to in the server load balancer instance. | Valid values: 1 to 65535. |
| Weight | Integer | Retained | Yes | The weight of the ECS instance to be attached to the SLB instance. | Valid values: 0 to 100. |

Response parameters

Fn::GetAtt

VServerGroupId: the ID of the VServer group.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "AttachVServerGroup": {
      "Type": "ALIYUN::SLB::BackendServerToVServerGroupAddition",
      "Properties": {
        "VServerGroupId": "sg-2zenh4ndwrqg14yt0****",
        "BackendServers": [
          {
            "ServerId": "i-25zsk****",
            "Weight": 20,
            "Port": 8080
          },
          {
            "ServerId": "i-25zsk****",
            "Weight": 100,
            "Port": 8081
          }
        ]
      }
    }
  }
}
```

5.5.6.4. ALIYUN::SLB::Certificate

ALIYUN::SLB::Certificate is used to upload a certificate to an SLB instance. Server certificates and CA certificates are supported.

Notice

- You can upload only one CA certificate at a time ("CertificateType": "CA ").
- You can upload only one server certificate and the corresponding private key at a time ("CertificateType": "Server ").

Syntax

```
{
  "Type": "ALIYUN::SLB::Certificate",
  "Properties": {
    "CertificateName": String,
    "Certificate": String,
    "AliCloudCertificateName": String,
    "PrivateKey": String,
    "ResourceGroupId": String,
    "CertificateType": String,
    "AliCloudCertificateId": String
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|----------|------|----------|----------|-------------|------------|
|----------|------|----------|----------|-------------|------------|

| Property | Type | Required | Editable | Description | Constraint |
|-------------------------|--------|----------|----------|---|--|
| ResourceGroupId | String | No | No | The ID of the resource group. | None |
| CertificateName | String | No | Yes | The name of the certificate. | None |
| Certificate | String | Yes | No | The public key of the certificate. | None |
| AliCloudCertificateName | String | No | No | The name of the Alibaba Cloud certificate. | None |
| PrivateKey | String | No | No | The server private key that you want to upload. | None |
| AliCloudCertificateId | String | No | No | The ID of the Alibaba Cloud certificate. | This parameter is required if you use a certificate from Alibaba Cloud SSL Certificates Service. |
| CertificateType | String | No | No | The type of the certificate. | Valid values: Server and CA. |

Response parameters

Fn::GetAtt

- **CertificateId**: the ID of the certificate.
- **Fingerprint**: the fingerprint of the certificate.

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Parameters": {
    "CertificateType": {
      "Type": "String",
      "Description": "The type of the certificate.",
      "AllowedValues": [
        "Server",
        "CA"
      ],
      "Default": "Server"
    },
    "AliCloudCertificateName": {
      "Type": "String",
      "Description": "The name of the Alibaba Cloud certificate."
    },
    "PrivateKey": {
      "Type": "String",
      "Description": "The private key."
    }
  }
}
```

```

},
"CertificateName": {
  "Type": "String",
  "Description": "The name of the certificate."
},
"Certificate": {
  "Type": "String",
  "Description": "The content of the certificate public key."
},
"AliCloudCertificateId": {
  "Type": "String",
  "Description": "The ID of the Alibaba Cloud certificate."
}
},
"Resources": {
  "Certificate": {
    "Type": "ALIYUN::SLB::Certificate",
    "Properties": {
      "CertificateType": {
        "Ref": "CertificateType"
      },
      "AliCloudCertificateName": {
        "Ref": "AliCloudCertificateName"
      },
      "PrivateKey": {
        "Ref": "PrivateKey"
      },
      "CertificateName": {
        "Ref": "CertificateName"
      },
      "Certificate": {
        "Ref": "Certificate"
      },
      "AliCloudCertificateId": {
        "Ref": "AliCloudCertificateId"
      }
    }
  }
},
"Outputs": {
  "Fingerprint": {
    "Description": "The fingerprint of the certificate.",
    "Value": {
      "Fn::GetAtt": [
        "Certificate",
        "Fingerprint"
      ]
    }
  },
  "CertificateId": {
    "Description": "The ID of the certificate.",
    "Value": {
      "Fn::GetAtt": [
        "Certificate",
        "CertificateId"
      ]
    }
  }
}
}
}
}

```

5.5.6.5. ALIYUN::SLB::DomainExtension

ALIYUN::SLB::DomainExtension is used to create a domain extension for an SLB instance.

Statement

```
{
  "Type": "ALIYUN::SLB::DomainExtension",
  "Properties": {
    "Domain": String,
    "ListenerPort": Integer,
    "ServerCertificateId": String,
    "LoadBalancerId": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|---------------------|---------|----------|----------|---|---------------------------|
| Domain | String | No | No | The custom domain name. | None |
| ListenerPort | Integer | Yes | No | The frontend port used by the HTTPS listener of the SLB instance. | Valid values: 1 to 65535. |
| ServerCertificateId | String | No | Yes | The ID of the certificate corresponding to the domain name. | None |
| LoadBalancerId | String | No | No | The ID of the SLB instance. | None |

Response parameters

Fn::GetAtt

- DomainExtensionId: the ID of the created domain extension.
- ListenerPort: The frontend port used by the SLB instance.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "DomainExtension": {
      "Type": "ALIYUN::SLB::DomainExtension",
      "Properties": {
        "Domain": "*.example1.com",
        "ListenerPort": "443",
        "ServerCertificateId": "123157908552****_166f8204689_1714763408_70998****",
        "LoadBalancerId": "lb-bp1o94dp5i6earr9g****"
      }
    }
  },
  "Outputs": {
    "DomainExtensionId": {
      "Value": {
        "Fn::GetAtt": [
          "DomainExtension",
          "DomainExtensionId"
        ]
      }
    },
    "ListenerPort": {
      "Value": {
        "Fn::GetAtt": [
          "DomainExtension",
          "ListenerPort"
        ]
      }
    }
  }
}
```

5.5.6.6. ALIYUN::SLB::Listener

ALIYUN::SLB::Listener is used to create a listener for an SLB instance.

Statement

```
{
  "Type": "ALIYUN::SLB::Listener",
  "Properties": {
    "MasterSlaveServerGroupId": String,
    "AclStatus": String,
    "Protocol": String,
    "AclId": String,
    "ServerCertificateId": String,
    "HealthCheck": Map,
    "RequestTimeout": Integer,
    "IdleTimeout": Integer,
    "ListenerPort": Integer,
    "HttpConfig": Map,
    "Bandwidth": Integer,
    "AclType": String,
    "BackendServerPort": Integer,
    "Scheduler": String,
    "LoadBalancerId": String,
    "CACertificateId": String,
    "Persistence": Map,
    "VServerGroupId": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------------|--------|----------|----------|---|---|
| MasterSlaveServerGroupId | String | Yes | Released | The ID of the active/standby server group. | None |
| AclStatus | String | Yes | Released | Specifies whether to enable access control on the listener. | Valid values: <ul style="list-style-type: none"> on off Default value: off. |
| AclId | String | Yes | Released | The ID of the access control list (ACL) to which the listener is bound. This parameter is required when the AclStatus parameter is set to on. | None |
| | | | | The type of the ACL. Valid values: white and black. <ul style="list-style-type: none"> white: specifies the ACL as a whitelist. Only | |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|---|---|
| AclType | String | Yes | Released | <p>addresses or CIDR blocks specified in the ACL are forwarded. Whitelists are applicable to scenarios where you want an application to only be accessed from specific IP addresses. Configuring a whitelist poses risks to your services. After a whitelist is configured, only the IP addresses specified in the whitelist are able to access the SLB listener. If a whitelist is enabled without any IP addresses specified, the SLB listener will not forward any requests.</p> <ul style="list-style-type: none"> • white: specifies the ACL as a whitelist. Requests from the IP addresses or CIDR blocks specified in the ACL are forwarded. Whitelists are applicable to scenarios where you want an application to only be accessed from specific IP addresses. • black: specifies the ACL as a blacklist. Requests from the IP addresses or CIDR blocks specified in the ACL are not forwarded. Blacklists are applicable to scenarios where you want an application to only be denied access from specific IP addresses. | <p>Valid values:</p> <ul style="list-style-type: none"> • White • Black |

| Parameter | Type | Required | Editable | If a blacklist is Description enabled | Constraint |
|--------------|---------|----------|----------|---|---|
| | | | | without any IP addresses specified, the SLB listener will forward all requests. This parameter is required when the AclStatus parameter is set to on. | |
| Protocol | String | No | No | The Internet protocol over which the listener will forward requests. | Valid values: <ul style="list-style-type: none"> • http • https • tcp • udp |
| ListenerPort | Integer | Yes | No | The frontend port used by the SLB instance. | Valid values: 1 to 65535. |

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------------|---------|----------|----------|--|--|
| Bandwidth | Integer | Yes | No | The peak bandwidth of the listener. Unit: Mbit/s. | <ul style="list-style-type: none"> Valid values: -1 and 1 to 1000. For an SLB instance that is connected to the Internet and billed by fixed bandwidth, this parameter cannot be set to -1, and the sum of peak bandwidth values assigned to different listeners cannot exceed the Bandwidth value specified when the SLB instance is created. For an SLB instance that is connected to the Internet and billed by traffic, this parameter can be set to -1. Unit: Mbit/s. |
| BackendServerPort | Integer | Yes | No | The backend port used by the SLB instance. | Valid values: 1 to 65535. |
| LoadBalancerId | String | No | No | The ID of the SLB instance. | None |
| HealthCheck | Map | Erased | Released | The health check settings of the listener. | None |
| Persistence | Map | Erased | Released | The persistence properties. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|---------------------|--------|----------|----------|---|--|
| Scheduler | String | Yes | Released | The algorithm used to direct traffic to individual servers. | Valid values: <ul style="list-style-type: none"> wrr wlc Default value: wrr |
| CACertificateId | String | Yes | Released | The ID of the CA certificate. | Only valid for HTTPS |
| ServerCertificateId | String | Yes | Released | The ID of the server certificate. | This parameter is required and valid only for HTTPS listeners. |
| VServerGroupId | String | Yes | Released | The ID of the VServer group. | None |
| RequestTimeout | String | Optional | Released | The request timeout period. Unit: seconds. | Valid values: 1 to 180. |
| IdleTimeout | String | Optional | Released | The idle connection timeout period. Unit: seconds. | Valid values: 1 to 60. |
| HttpConfig | Map | Erased | Released | The HTTP configurations. | None |

HealthCheck syntax

```
"HealthCheck": {
  "Domain": String,
  "Interval": Integer,
  "URI": String,
  "HttpCode": String,
  "HealthyThreshold": Integer,
  "Timeout": Integer,
  "UnhealthyThreshold": Integer,
  "Port": Integer
}
```

HealthCheck properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|------|----------|----------|-------------|------------|
|-----------|------|----------|----------|-------------|------------|

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|--|---|
| Domain | String | Yes | Released | The domain name used for health checks. | <ul style="list-style-type: none"> The value can be <code>\$_ip</code>, a custom string, or an empty string. A custom string must be 1 to 80 characters in length and can contain only letters, digits, hyphens (-), and periods (.). When this parameter is set to <code>\$_ip</code> or left empty, the SLB instance uses the private IP addresses of backend servers as the domain names for health checks. |
| Interval | String | Optional | Released | The time interval between consecutive health checks. Unit: seconds. | Valid values: 1 to 5. Unit: seconds. |

| Parameter | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|--|--|
| URI | String | Yes | Released | The URI used for health checks. | <ul style="list-style-type: none"> The URI must be 1 to 80 characters in length and can contain letters, digits, hyphens (-), forward slashes (/), periods (.), percent signs (%), question marks (?), number signs (#), and ampersands (&). It must start with a forward slash (/). |
| HttpCode | String | Yes | Released | The HTTP status code that indicates a positive health status of the backend servers. | <ul style="list-style-type: none"> Valid values: http_2xx, http_3xx, http_4xx, and http_5xx. Separate multiple HTTP status codes with commas (,). <p>Default value: http_2xx</p> |
| HealthyThreshold | String | Optional | Released | The threshold used to determine that the backend servers are healthy. This value indicates the number of consecutive successful health checks required before the health status of a backend server can be changed from fail to success. | Valid values: 1 to 10. |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------|--------|----------|----------|--|--|
| Timeout | String | Optional | Released | The length of time to wait for the response from a health check. Unit: seconds. | Valid values: 1 to 50.  Notice This parameter is valid only when its value is greater than or equal to that of the Interval parameter. Otherwise, this parameter will be overridden by the Interval value. |
| UnhealthyThreshold | String | Optional | Released | The threshold used to determine that the backend servers are unhealthy. This value indicates the number of consecutive failed health checks required before the health status of a backend server can be changed from success to fail. | Valid values: 1 to 10. |
| Port | String | Optional | Released | The port used for health checks. | Valid values: 0 to 65535. |

Persistence syntax

```
"Persistence": {
  "PersistenceTimeout": Integer,
  "CookieTimeout": Integer,
  "XForwardedFor": String,
  "Cookie": String,
  "StickySession": String,
  "StickySessionType": String
}
```

Persistence properties

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------|--------|----------|----------|--|---|
| StickySession | String | No | No | Specifies whether to enable session persistence. | Valid values: <ul style="list-style-type: none"> • on • off |
| PersistenceTimeout | String | Optional | Released | The maximum duration that the client can be connected to the server. Unit: seconds. | Valid values: 0 to 1000. The default value is 0, which indicates that connection persistence is disabled. Unit: seconds. |
| CookieTimeout | String | Optional | Released | The maximum duration the cookie can be retained before it expires. Unit: seconds. | This parameter is required when the StickySession parameter is set to on and the StickySessionType parameter is set to insert. Valid values: 1 to 86400. Unit: seconds. |
| XForwardedFor | String | Yes | Released | Specifies whether to use the X-Forwarded-For header field to obtain the real IP address of the client. | Valid values: <ul style="list-style-type: none"> • on • off Default value: on |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|--|---|
| Cookie | String | Yes | Released | The cookie configured on the backend server. | <ul style="list-style-type: none"> The parameter value must be 1 to 200 characters in length and follow the RFC 2965 standard. It can contain only ASCII characters. It cannot contain commas (,), semicolons (;), or spaces, and cannot start with a dollar sign (\$). This parameter is required when the StickySession parameter is set to on and the StickySession Type parameter is set to server. |

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------------|--------|----------|----------|------------------------------------|---|
| StickySessionType | String | Yes | Released | The method for processing cookies. | <ul style="list-style-type: none"> Valid values: insert and server. When this parameter is set to insert, SLB adds a cookie to the first response from the backend server. Then, the next request contains the cookie and the listener distributes the request to the same backend server. When this parameter is set to server, SLB overwrites the original cookie when a new cookie is set. The next time the client carries the new cookie to access SLB, the listener distributes the request to the previously recorded backend server. This parameter is required when the StickySession parameter is set to on. This parameter is ignored when the StickySession parameter is set to off. |

HttpConfig syntax

```
"HttpConfig": {  
  "ForwardPort": Integer,  
  "ListenerForward": String  
}
```

HttpConfig properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|--|---|
| ForwardPort | String | Optional | Released | The port used to redirect HTTP requests to HTTPS. | Valid values: 1 to 65535. Default value: 443. |
| ListenerForward | String | No | No | Specifies whether to enable HTTP-to-HTTPS redirection. | Valid values: • on • off Default value: off. |

Response parameters

Fn::GetAtt

- LoadBalancerId: the unique ID of the SLB instance.
- ListenerPortsAndProtocol: an array consisting of the ports and protocols used by the SLB listener.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "LoadBalancer": {
      "Type": "ALIYUN::SLB::LoadBalancer",
      "Properties": {
        "LoadBalancerName": "createdByHeat",
        "AddressType": "internet",
        "InternetChargeType": "paybybandwidth"
      }
    },
    "CreateListener": {
      "Type": "ALIYUN::SLB::Listener",
      "Properties": {
        "LoadBalancerId": {"Ref": "LoadBalancer"},
        "ListenerPort": "8094",
        "BackendServerPort": 8080,
        "Bandwidth": 1,
        "Protocol": "http",
        "HealthCheck": {
          "HealthyThreshold": 3,
          "UnhealthyThreshold": 3,
          "Interval": 2,
          "Timeout": 5,
          "HttpCode": "http_2xx,http_3xx,http_4xx,http_5xx"
        },
        "Scheduler": "wrr",
        "Persistence": {
          "PersistenceTimeout": 1,
          "XForwardedFor": "on",
          "StickySession": "on",
          "StickySessionType": "insert",
          "CookieTimeout": 10,
          "Cookie": "1"
        }
      }
    }
  },
  "Outputs": {
    "LoadBalanceDetails": {
      "Value": {"Fn::GetAtt": ["LoadBalancer", "LoadBalancerId"]}
    }
  }
}
```

5.5.6.7. ALIYUN::SLB::LoadBalancer

ALIYUN::SLB::LoadBalancer is used to create an SLB instance.

Statement

```
{
  "Type": "ALIYUN::SLB::LoadBalancer",
  "Properties": {
    "DeletionProtection": Boolean,
    "AddressType": String,
    "Tags": List,
    "InternetChargeType": String,
    "Bandwidth": Integer,
    "SlaveZoneId": String,
    "ResourceGroupId": String,
    "AutoPay": Boolean,
    "VpcId": String,
    "PricingCycle": String,
    "LoadBalancerName": String,
    "Duration": Number,
    "VSwitchId": String,
    "LoadBalancerSpec": String,
    "MasterZoneId": String,
    "PayType": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------|---------|----------|----------|--|---|
| ResourceGroupId | String | Yes | Released | The ID of the resource group to which the RDS instance belongs. | None |
| DeletionProtection | Boolean | Erased | Released | Specifies whether to enable deletion protection to prevent the SLB instance from being deleted by mistake. | Valid values: <ul style="list-style-type: none"> • true • false |
| VpcId | String | Yes | Released | The ID of the VPC. | None |
| SlaveZoneId | String | Yes | Released | The ID of the secondary zone to which the SLB instance belongs. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------|--------|----------|----------|---|--|
| Bandwidth | String | Optional | Released | The peak bandwidth of SLB instances that are connected to the Internet and billed by fixed bandwidth. | For SLB instances that are connected to the Internet and billed by fixed Bandwidth, this parameter is valid only when the Bandwidth parameter of the SLB Listener is specified. For Internet instances whose billing type is to pay by traffic, we recommend that you set the peak Bandwidth through the Listener parameter. In this case, this parameter is ignored. Valid values: 1 to 1000. Unit: Mbps. Default value: 1 VPC-type instances are billed by traffic. |
| AddressType | String | Yes | Released | The address type of the SLB instance. | Valid values: <ul style="list-style-type: none"> internet intranet Default value: internet. |
| VSwitchId | String | Yes | Released | The ID of the vSwitch in the VPC. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------|--------|----------|----------|--|--|
| LoadBalancerName | String | Yes | Released | The name of the SLB instance to be created. | A custom string. The name must be 1 to 80 characters in length and can contain letters, digits, hyphens (-), forward slashes (/), periods (.), and underscores (_). If this parameter is not specified, the system assigns a value. |
| InternetChargeType | String | Yes | Released | The billing method of SLB instances that are connected to the Internet. | Valid values: <ul style="list-style-type: none"> • paybybandwidth • paybytraffic Default value: paybytraffic. |
| MasterZoneId | String | Yes | Released | The ID of the primary zone to which the SLB instance belongs. | None |
| Tags | List | Erased | Released | The tags to be attached to the SLB instance. The tags are listed in JSON format. Each tag consists of a TagKey and a TagValue. | A maximum of five tags can be attached to an SLB instance. |
| LoadBalancerSpec | String | Yes | Released | The type of the SLB instance. | Valid values: <ul style="list-style-type: none"> • slb.s1.small • slb.s2.small • slb.s2.medium • slb.s3.small • slb.s3.medium • slb.s3.large The available types vary by region. |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------|---------|----------|----------|---|--|
| AutoPay | Boolean | Erased | Released | Specifies whether to automatically pay for subscription SLB instances that are connected to the Internet. | Valid values: <ul style="list-style-type: none"> • true • false Default value: false. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note This parameter is valid only on the China site (aliyun.com). </div> |
| PayType | String | Yes | Released | The billing method of the SLB instance. | Valid values: <ul style="list-style-type: none"> • PayOnDemand • PrePay |
| PricingCycle | String | Yes | Released | The billing cycle of subscription SLB instances that are connected to the Internet. | Valid values: <ul style="list-style-type: none"> • month • year <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note This parameter is valid only on the China site (aliyun.com). </div> |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|--------|----------|----------|---|---|
| Duration | Number | Erased | Released | The subscription period of subscription SLB instances that are connected to the Internet. | <p>Valid values:</p> <ul style="list-style-type: none"> Valid values when the PricingCycle parameter is set to month: 1 to 9. Valid values when the PricingCycle parameter is set to year: 1 to 3. <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e0f0ff;"> <p>Note This parameter is valid only on the China site (aliyun.com).</p> </div> |

Tags syntax

```
"Tags": [
  {
    "Value": String,
    "Key": String
  }
]
```

Tags properties

| Property | Type | Required or Not | Editable | Description | Constraint |
|----------|--------|-----------------|----------|-------------|--|
| Key | String | No | No | None | None |
| Value | String | Yes | Released | None | <p>Valid values: 1025 to 10000. You cannot use the following commonly used port numbers: 2222, 4500, 4510, 4560, 7505, 9000, 9001, and 9002.</p> |

Response parameters

Fn::GetAtt

- LoadBalancerId: the unique ID of the SLB instance.
- NetworkType: the network type of the SLB instance, which can be vpc or classic.

- **AddressType**: the address type of the SLB instance, which can be intranet or internet.
- **IpAddress**: the IP address of the SLB instance.
- **OrderId**: the ID of the order.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "CreateLoadBalance": {
      "Type": "ALIYUN::SLB::LoadBalancer",
      "Properties": {
        "LoadBalancerName": "createdByHeat",
        "AddressType": "internet",
        "InternetChargeType": "paybybandwidth",
      }
    }
  },
  "Outputs": {
    "LoadBalanceDetails": {
      "Value": {
        "Fn::GetAtt": ["CreateLoadBalance", "LoadBalancerId"]
      }
    }
  }
}
```

5.5.6.8. ALIYUN::SLB::LoadBalancerClone

ALIYUN::SLB::LoadBalancerClone is used to clone an SLB instance.

Syntax

```
{
  "Type": "ALIYUN::SLB::LoadBalancerClone",
  "Properties": {
    "Tags": List,
    "ResourceGroupId": String,
    "VSwitchId": String,
    "LoadBalancerName": String,
    "SourceLoadBalancerId": String,
    "TagsPolicy": String,
    "BackendServersPolicy": String,
    "BackendServers": List
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|-----------------|--------|----------|----------|-------------------------------|------------|
| ResourceGroupId | String | No | No | The ID of the resource group. | None |

| Property | Type | Required | Editable | Description | Constraint |
|----------------------|--------|----------|----------|--|--|
| VSwitchId | String | No | No | The ID of the VSwitch. | The VSwitch must exist in the VPC to which the source SLB instance belongs. If the parameter is not specified, the VSwitch to which the source SLB instance belongs is used. |
| SourceLoadBalancerId | String | Yes | No | The ID of the SLB instance to be cloned. | None |

| Property | Type | Required | Editable | Description | Constraint |
|----------------------|--------|----------|----------|---|---|
| BackendServersPolicy | String | No | No | The clone policy. The ECS instances listened by the new SLB instance and the weight of each ECS instance are specified in the policy. | <p>Valid values:</p> <ul style="list-style-type: none"> clone: The ECS instances listened by the source SLB instance and the ECS instance weights are cloned to the new SLB instance. empty: No ECS instances are attached to the new SLB instance. append: The ECS instances listened by the source SLB instance and the ECS instance weights are cloned to the new SLB instance. New ECS instances with specified weights are also attached to the new SLB instance. replace: New ECS instances with specified weights are attached to the new SLB instance. However, the ECS instances listened by the source SLB instance and the ECS instance weights are not cloned to the new SLB instance. <p>Default value: clone.</p> |

| Property | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|---------------------------------------|--|
| BackendServers | List | No | Yes | The new ECS instances to be listened. | None |
| LoadBalancerName | String | No | No | The name of the new SLB instance. | You can set the name to any string. The name must be 1 to 80 characters in length and can contain letters, digits, hyphens (-), forward slashes (/), periods (.), and underscores (_). |
| Tags | List | No | Yes | The tags of the SLB instance. | Tags must be specified as key-value pairs. A maximum of five tags can be specified. |
| TagsPolicy | String | No | No | The policy of the tags. | Valid values: <ul style="list-style-type: none"> • clone: The tags of the source SLB instance are used. • empty: No tags are configured. • append: The tags of the source SLB instance are reserved while new tags are added. • replace: The tags of the source SLB instance are deleted while new tags are added. Default value: empty. |

BackendServers syntax

```
"BackendServers": [
  {
    "ServerId": String,
    "Weight": Integer
  }
]
```

BackendServers properties

| Property | Type | Required | Editable | Description | Constraint |
|----------|---------|----------|----------|--|---|
| ServerId | String | Yes | Yes | The ID of the ECS instance. | The ECS instances must be in the running state. |
| Weight | Integer | Yes | Yes | The weight of the ECS instance to be attached to the SLB instance. | Valid values: 0 to 100. Default value: 100. |

Response parameters

Fn::GetAtt

LoadBalancerId: the ID of the new SLB instance.

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "CloneLoadBalance": {
      "Type": "ALIYUN::SLB::LoadBalancerClone",
      "Properties": {
        "SourceLoadBalancerId": "150ebed5f06-cn-beijing-btc-***",
        "LoadBalancerName": "rosnew",
        "BackendServersPolicy": "replace",
        "BackendServers": [
          {
            "ServerId": "i-25zsk****",
            "Weight": 20
          }
        ]
      }
    }
  },
  "Outputs": {
    "LoadBalanceDetails": {
      "Value": {"Fn::GetAtt": ["CloneLoadBalance", "LoadBalancerId"]}
    }
  }
}
```

5.5.6.9. ALIYUN::SLB::MasterSlaveServerGroup

ALIYUN::SLB::MasterSlaveServerGroup is used to create a primary/secondary server group.

Notice A primary/secondary server group contains only two ECS instances: a primary server and a secondary server.

Syntax

```
{
  "Type": "ALIYUN::SLB::MasterSlaveServerGroup",
  "Properties": {
    "MasterSlaveServerGroupName": String,
    "MasterSlaveBackendServers": List,
    "LoadBalancerId": String
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|----------------------------|--------|----------|----------|--|--|
| MasterSlaveServerGroupName | String | No | No | The name of the primary/secondary server group. | None |
| MasterSlaveBackendServers | List | Yes | No | The list of backend servers in the primary/secondary server group. | A primary/secondary server group can contain a maximum of two backend servers. If you do not specify this parameter, an empty primary/secondary server group is created. |
| LoadBalancerId | String | Yes | No | The ID of the SLB instance. | None |

MasterSlaveBackendServers syntax

```
"MasterSlaveBackendServers": [
  {
    "ServerId": String,
    "Port": Integer,
    "Weight": Integer,
    "ServerType": String
  }
]
```

MasterSlaveBackendServers properties

| Property | Type | Required | Editable | Description | Constraint |
|----------|------|----------|----------|-------------|------------|
|----------|------|----------|----------|-------------|------------|

| Property | Type | Required | Editable | Description | Constraint |
|------------|---------|----------|----------|--|--|
| ServerId | String | Yes | No | The ID of the ECS instance or Elastic Network Interface (ENI) to be added. | None |
| ServerType | String | No | No | The type of the server. | Default value: Master. Valid values: <ul style="list-style-type: none"> • Master • Slave |
| Port | Integer | Yes | No | The port number used by the backend server. | Valid values: 1 to 65535. |
| Weight | Integer | Yes | No | The weight of the backend server. | Valid values: 0 to 100. |

Response parameters

Fn::GetAtt

MasterSlaveServerGroupId: the ID of the primary/secondary server group.

Examples

```

{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "MasterSlaveServerGroup": {
      "Type": "ALIYUN::SLB::MasterSlaveServerGroup",
      "Properties": {
        "MasterSlaveServerGroupName": "Group1",
        "MasterSlaveBackendServers": [
          {
            "ServerId": "vm****",
            "Port": "80",
            "Weight": "100",
            "ServerType": "Master"
          },
          {
            "ServerId": "vm****",
            "Port": "90",
            "Weight": "100",
            "ServerType": "Slave"
          }
        ],
        "LoadBalancerId": "lb-bp1hv944r69a4j9j*****"
      }
    },
    "Outputs": {
      "MasterSlaveServerGroupId": {
        "Value": {
          "Fn::GetAtt": [
            "MasterSlaveServerGroup",
            "MasterSlaveServerGroupId"
          ]
        }
      }
    }
  }
}

```

5.5.6.10. ALIYUN::SLB::Rule

ALIYUN::SLB::Rule is used to add forwarding rules for a specified HTTP or HTTPS listener.

Statement

```

{
  "Type": "ALIYUN::SLB::Rule",
  "Properties": {
    "ListenerPort": Integer,
    "RuleList": List,
    "LoadBalancerId": String
  }
}

```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|----------------|---------|----------|----------|--|--|
| ListenerPort | Integer | Yes | No | The frontend listener port used by the SLB instance. | Valid values: 1 to 65,535. |
| RuleList | List | Yes | No | The list of forwarding rules to be added. | <p>A maximum of 10 forwarding rules can be added at a time.</p> <p>Each forwarding rule contains the following parameters: RuleName, Domain, Url, and VServerGroupId.</p> <p>You must specify at least one of the following parameters: Domain and URL.</p> <p>The combination of Domain and URL must be unique in a listener.</p> |
| LoadBalancerId | String | No | No | The IDs of SLB instances. | None |

RuleList syntax

```
"RuleList": [
  {
    "Url": String,
    "Domain": String,
    "VServerGroupId": String,
    "RuleName": String
  }
]
```

RuleList properties

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|------|----------|----------|-------------|------------|
|-----------|------|----------|----------|-------------|------------|

| Parameter | Type | Required | Editable | Description | Constraint |
|----------------|--------|----------|----------|---|--|
| Url | String | Yes | Released | The request URL. | The name must be 1 to 80 characters in length and can contain letters, numbers, and special characters. -, ., percent signs (%), question marks (?), #& . |
| Domain | String | Yes | Released | The request domain name associated with the forwarding rule. | None |
| VServerGroupId | String | No | No | The ID of the destination VServer group specified in the forwarding rule. | None |
| RuleName | String | No | No | The name of the forwarding rule. | The name must be 1 to 40 characters in length and can contain letters, numbers, and special characters. -, ., _ . Forwarding rule names must be unique within each listener. |

Response parameters

Fn::GetAtt

Rules: the list of forwarding rules.

Sample request

```

{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "Rule": {
      "Type": "ALIYUN::SLB::Rule",
      "Properties": {
        "ListenerPort": {
          "Ref": "ListenerPort"
        },
        "RuleList": {
          "Fn::Split": [",", {
            "Ref": "RuleList"
          }], {
            "Ref": "RuleList"
          }
        }
      },
      "LoadBalancerId": {
        "Ref": "LoadBalancerId"
      }
    }
  },
  "Parameters": {
    "ListenerPort": {
      "Type": "Number",
      "Description": "The front-end HTTPS listener port of the Server Load Balancer instance. Valid value:\n1-65535",
      "MaxValue": 65535,
      "MinValue": 1
    },
    "RuleList": {
      "MinLength": 1,
      "Type": "CommaDelimitedList",
      "Description": "The forwarding rules to add.",
      "MaxLength": 10
    },
    "LoadBalancerId": {
      "Type": "String",
      "Description": "The ID of Server Load Balancer instance."
    }
  },
  "Outputs": {
    "Rules": {
      "Description": "A list of forwarding rules. Each element of rules contains \"RuleId\".",
      "Value": {
        "Fn::GetAtt": ["Rule", "Rules"]
      }
    }
  }
}

```

5.5.6.11. ALIYUN::SLB::VServerGroup

ALIYUN::SLB::VServerGroup is used to create a VServer group and adds backend servers to the SLB instance.

Syntax

```
{
  "Type" : "ALIYUN::SLB::VServerGroup",
  "Properties" : {
    "VServerGroupName" : String,
    "BackendServers" : List,
    "LoadBalancerId" : String
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|--|--|
| VServerGroupName | String | Yes | No | The name of the VServer group. | None |
| BackendServers | List | Yes | Yes | The list of ECS instances to be added. | A list can contain up to 20 instances. |
| LoadBalancerId | String | Yes | No | The ID of the SLB instance. | None |

BackendServers syntax

```
"BackendServers" : [
  {
    "ServerId" : String,
    "Port" : Integer,
    "Weight" : Integer
  }
]
```

BackendServers properties

| Property | Type | Required | Editable | Description | Constraint |
|----------|---------|----------|----------|--|---------------------------|
| ServerId | String | Yes | Yes | The ID of the ECS instance. | None |
| Port | Integer | Yes | Yes | The backend port used by the SLB instance. | Valid values: 1 to 65535. |
| Weight | Integer | Yes | Yes | The weight of the ECS instance to be attached to the SLB instance. | Valid values: 0 to 100. |

Response parameters

Fn::GetAtt

- VServerGroupId: the ID of the VServer group.
- BackendServers: the list of backend servers added to the SLB instance

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "CreateVServerGroup": {
      "Type": "ALIYUN::SLB::VServerGroup",
      "Properties": {
        "LoadBalancerId": "lb-2zenh4ndwrqg14yt0****",
        "VServerGroupName": "VServerGroup-****",
        "BackendServers": [
          {
            "ServerId": "i-25zsk****",
            "Weight": 20,
            "Port": 8080
          },
          {
            "ServerId": "i-25zsk****",
            "Weight": 100,
            "Port": 8081
          }
        ]
      }
    }
  },
  "Outputs": {
    "VServerGroupId": {
      "Value": {"Fn::GetAttr": ["CreateVServerGroup", "VServerGroupId"]}
    }
  }
}
```

5.5.7. VPC

5.5.7.1. ALIYUN::VPC::EIP

ALIYUN::VPC::EIP is used to apply for an Elastic IP address.

Statement

```
{
  "Type": "ALIYUN::VPC::EIP",
  "Properties": {
    "Isp": String,
    "Period": Number,
    "ResourceGroupId": String,
    "AutoPay": Boolean,
    "InstanceChargeType": String,
    "PricingCycle": String,
    "InternetChargeType": String,
    "Bandwidth": Number
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------|--------|----------|----------|--|---|
| ResourceGroupId | String | Yes | Released | The ID of the resource group to which the RDS instance belongs. | None |
| Bandwidth | Number | Erased | Released | The network bandwidth. Unit: Mbit/s. | If this parameter is not specified, the default value 5Mbps is used. |
| InternetChargeType | String | Yes | Released | The billing method for network usage. Default value: PayByBandwidth. | Valid values: <ul style="list-style-type: none"> PayByBandwidth: pay-by-bandwidth. PayByTraffic Default value: PayByBandwidth. |
| InstanceChargeType | String | Yes | Released | The billing method of the Elastic IP address. Default value: Postpaid. | Valid values: <ul style="list-style-type: none"> Prepaid Postpaid: pay-as-you-go Default value: PostPaid. |
| PricingCycle | String | Yes | Released | The billing cycle of the subscription. Default value: Month. | Valid values: <ul style="list-style-type: none"> Month: paid by month. Year Default value: Month. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p> Note This parameter is required when InstanceChargeType is set to Prepaid.</p> </div> |
| Period | Number | Erased | Released | The subscription period. | Valid values: <ul style="list-style-type: none"> If pay by month is selected, the billing method can be a fee of 1 to 9. If pay-as-you-go is selected, the payment can be in the range of 1 to 3. Default value: 1 <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p> Note This parameter is required when InstanceChargeType is set to Prepaid.</p> </div> |

| Parameter | Type | Required | Editable | Description | Constraint |
|-----------|---------|----------|----------|---|--|
| AutoPay | Boolean | Erased | Released | Specifies whether to enable automatic payment. | Valid values: <ul style="list-style-type: none"> false: Automatic payment is disabled. After an order is generated, you must go to the Order Center to make the payment. true: Automatic payment is enabled. Payments are automatically made. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note This parameter is required when InstanceChargeType is set to Prepaid. </div> |
| isp | String | Yes | Released | The ISP tag used for Finance Cloud. This parameter takes effect only when your region is set to China (Hangzhou). | This parameter is ignored if you are not a Finance Cloud user. |

Response parameters

Fn::GetAtt

- EipAddress: the allocated Elastic IP address.
- AllocationId: the ID of the instance that the Elastic IP address is allocated to.
- OrderId: The order ID that is returned when you set the InstanceChargeType parameter to Prepaid.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "Eip": {
      "Type": "ALIYUN::VPC::EIP",
      "Properties": {
        "InternetChargeType": "PayByTraffic",
        "Bandwidth": 200
      }
    }
  },
  "Outputs": {
    "EipAddress": {
      "Value": {"Fn::GetAtt": ["Eip", "EipAddress"]}
    },
    "AllocationId": {
      "Value": {"Fn::GetAtt": ["Eip", "AllocationId"]}
    },
    "OrderId": {
      "Value": {"Fn::GetAtt": ["Eip", "OrderId"]}
    }
  }
}
```

5.5.7.2. ALIYUN::VPC::EIPAssociation

ALIYUN::VPC::EIPAssociation is used to associate an Elastic IP address with a cloud service instance.

Statement

```
{
  "Type": "ALIYUN::VPC::EIPAssociation",
  "Properties": {
    "AllocationId": String,
    "InstanceId": String,
    "PrivateIpAddress": String,
    "Mode": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------|--------|----------|----------|-----------------------------------|------------|
| AllocationId | String | No | Yes | The ID of the Elastic IP address. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|------------------|--------|----------|----------|--|---|
| InstanceId | String | No | Yes | The ID of the cloud service instance. | The following instance types are supported: <ul style="list-style-type: none"> VPC-connected ECS instances VPC-connected SLB instances NAT gateways HAVIP Elastic network interfaces |
| PrivateIpAddress | String | Yes | True | The private IP address in the CIDR block of the VSwitch. | None |
| Mode | String | Yes | True | The association mode. | Valid values: <ul style="list-style-type: none"> NAT MULTI_BINDED |

Response parameters

Fn::GetAtt

- EipAddress: The allocated Elastic IP address.
- AllocationId: The ID of the instance to which the Elastic IP address is allocated.

Examples

JSON format

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "Eip": {
      "Type": "ALIYUN::VPC::EIP",
      "Properties": {
        "InternetChargeType": "PayByTraffic",
        "Bandwidth": 200
      }
    },
    "EipAssociation": {
      "Type": "ALIYUN::VPC::EIPAssociation",
      "Properties": {
        "InstanceId": "<LoadBalancerId>",
        "InstanceType": "EcsInstance",
        "AllocationId": {
          "Fn::GetAtt": ["Eip", "AllocationId"]
        }
      }
    }
  },
  "Outputs": {
    "EipAddress": {
      "Value": {"Fn::GetAtt": ["EipAssociation", "EipAddress"]}
    },
    "AllocationId": {
      "Value": {"Fn::GetAtt": ["EipAssociation", "AllocationId"]}
    }
  }
}
```

YAML format

```
ROSTemplateFormatVersion: '2015-09-01'
Resources:
  Eip:
    Type: ALIYUN::VPC::EIP
    Properties:
      InternetChargeType: PayByTraffic
      Bandwidth: 200
  EipAssociation:
    Type: ALIYUN::VPC::EIPAssociation
    Properties:
      InstanceId: "<LoadBalancerId>"
      InstanceType: EcsInstance
  AllocationId:
    Fn::GetAtt:
      - Eip
      - AllocationId
Outputs:
  EipAddress:
    Value:
      Fn::GetAtt:
        - EipAssociation
        - EipAddress
  AllocationId:
    Value:
      Fn::GetAtt:
        - EipAssociation
        - AllocationId
```

5.5.7.3. ALIYUN::VPC::PeeringRouterInterfaceBinding

ALIYUN::VPC::PeeringRouterInterfaceBinding is used to associate two router interfaces to be interconnected.

Statement

```
{
  "Type": "ALIYUN::VPC::PeeringRouterInterfaceBinding",
  "Properties": {
    "OppositeRouterId": String,
    "OppositeInterfaceId": String,
    "OppositeInterfaceOwnerId": String,
    "RouterInterfaceId": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|---------------------|--------|----------|----------|--------------------------------------|------------|
| RouterInterfaceId | String | No | No | The ID of the router interface. | None |
| OppositeInterfaceId | String | No | No | The ID of the peer router interface. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|--------------------------|--------|----------|----------|--|------------|
| OppositeRouterId | String | Yes | Released | The ID of the router to which the peer router interface belongs. | None |
| OppositeInterfaceOwnerId | String | Yes | Released | The ID of the owner of the peer router interface. | None |

Response parameters

Fn::GetAtt

RouterInterfaceId: the ID of the vRouter.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "InitiatorRouterInterfaceBinding": {
      "Type": "ALIYUN::VPC::PeeringRouterInterfaceBinding",
      "Properties": {
        "RouterInterfaceId": "ri-2zedgo0ih64g1me29****",
        "OppositeInterfaceId": "ri-2zex1tkyym98pjaor****",
        "OppositeRouterId": "vrt-2zexb35tzorIU0286****"
      }
    }
  }
}
```

5.5.7.4. ALIYUN::VPC::PeeringRouterInterfaceConnection

ALIYUN::VPC::PeeringRouterInterfaceConnection is used to initiate a router interface connection.

Statement

```
{
  "Type": "ALIYUN::VPC::PeeringRouterInterfaceConnection",
  "Properties": {
    "OppositeInterfaceId": String,
    "RouterInterfaceId": String
  }
}
```

Properties

| Parameter | Type | Required | Editable | Description | Constraint |
|---------------------|--------|----------|----------|--|------------|
| OppositeInterfaceId | String | No | No | The ID of the acceptor router interface. | None |

| Parameter | Type | Required | Editable | Description | Constraint |
|-------------------|--------|----------|----------|--|------------|
| RouterInterfaceId | String | No | No | The ID of the router interface to initiate the connection. | None |

Response parameters

Fn::GetAtt

- OppositeInterfaceId: the ID of the acceptor router interface.
- RouterInterfaceId: the ID of the router interface that initiates the connection.

Sample request

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "InitiatorRouterInterfaceBinding": {
      "Type": "ALIYUN::VPC::PeeringRouterInterfaceConnection",
      "Properties": {
        "RouterInterfaceId": "ri-2zedgo0ih64g1me29****",
        "OppositeInterfaceId": "ri-2ze4k5n2aeardu8cy****"
      }
    }
  }
}
```

5.5.7.5. ALIYUN::VPC::RouterInterface

ALIYUN::VPC::RouterInterface is used to create a router interface.

Syntax

```
{
  "Type": "ALIYUN::VPC::RouterInterface",
  "Properties": {
    "OppositeRegionId": String,
    "Description": String,
    "HealthCheckSourceIp": String,
    "RouterType": String,
    "AccessPointId": String,
    "RouterId": String,
    "Role": String,
    "OppositeInterfaceOwnerId": String,
    "OppositeAccessPointId": String,
    "HealthCheckTargetIp": String,
    "OppositeRouterId": String,
    "Spec": String,
    "OppositeRouterType": String,
    "Name": String,
    "PricingCycle": String,
    "Period": Number,
    "AutoPay": Boolean,
    "InstanceChargeType": String
  }
}
```

Properties

| Property | Type | Required | Editable | Description | Constraint |
|---------------|--------|----------|----------|---|--|
| RouterId | String | Yes | No | The ID of the router | None |
| Role | String | Yes | No | The role of the router interface. | <ul style="list-style-type: none"> When RouterType is set to VBR, set the value to InitiatingSide. When OppositeRouterType is set to VBR, set the value to AcceptingSide. |
| RouterType | String | No | No | The type of the router to which the router interface belongs. | Valid values: <ul style="list-style-type: none"> VRouter VBR |
| AccessPointId | String | No | No | The ID of the access point of the router interface. | <ul style="list-style-type: none"> This parameter is required when RouterType is set to VBR. The access point ID cannot be modified after the router interface is created. This parameter is not required when RouterType is set to VRouter. |

| Property | Type | Required | Editable | Description | Constraint |
|--------------------------|--------|----------|----------|--|--|
| Spec | String | No | No | The specifications of the router interface. | <p>The following list includes available specifications and the corresponding bandwidth values:</p> <ul style="list-style-type: none"> • Mini.2: 2 Mbit/s • Mini.5: 5 Mbit/s • Small.1: 10 Mbit/s • Small.2: 20 Mbit/s • Small.5: 50 Mbit/s • Middle.1: 100 Mbit/s • Middle.2: 200 Mbit/s • Middle.5: 500 Mbit/s • Large.1: 1,000 Mbit/s • Large.2: 2,000 Mbit/s • Large.5: 5,000 Mbit/s • Xlarge.1: 10,000 Mbit/s <div style="background-color: #e1f5fe; padding: 10px; border: 1px solid #cfcfcf;"> <p> Note</p> <ul style="list-style-type: none"> • This parameter is required when Role is set to InitiatingSide. • The value Negative is used by default when Role is set to AcceptingSide. </div> |
| OppositeRegionId | String | No | No | The region ID of the peer router interface. | None |
| OppositeInterfaceOwnerId | String | No | No | The ID of the owner of the peer router interface. | The default value is the ID of the current user. |
| OppositeRouterId | String | No | No | The ID of the router to which the peer router interface belongs. | None |

| Property | Type | Required | Editable | Description | Constraint |
|-----------------------|--------|----------|----------|--|--|
| OppositeRouterType | String | No | No | The type of the router to which the peer router interface belongs. | Valid values: <ul style="list-style-type: none"> When RouterType is set to VBR, set the value to VRouter. VBR |
| OppositeAccessPointId | String | No | No | The ID of the access point of the peer router interface. | <ul style="list-style-type: none"> When OppositeRouterType is set to VBR, this parameter is required. The access point ID cannot be modified after the router interface is created. When OppositeRouterType is set to VRouter, this parameter is not required. When OppositeRouterType is set to VBR, the VBR specified by the OppositeRouterId parameter must be in the access point specified by the OppositeAccessPointId parameter. |
| Description | String | No | No | The description of the router interface. | The description must be 2 to 256 characters in length. It cannot start with <code>http://</code> or <code>https://</code> . The parameter is empty by default. |
| Name | String | No | No | The display name of the router interface. | <ul style="list-style-type: none"> The name must be 2 to 128 characters in length and can contain letters, digits, periods(.), underscores(_), and hyphens(-). It must start with a letter but cannot start with <code>http://</code> or <code>https://</code>. |

| Property | Type | Required | Editable | Description | Constraint |
|---------------------|--------|----------|----------|--|---|
| HealthCheckSourceIp | String | No | No | The source IP address of health check packets used in leased line disaster recovery and ECMP scenarios. | <p>This parameter is valid only for VRouter interfaces with a peer router interface on a VBR.</p> <p>It must be an unused IP address in the VPC where the local VRouter is located.</p> <p>The HealthCheckSourceIp and HealthCheckTargetIp parameters must either both be specified or both left unspecified.</p> |
| HealthCheckTargetIp | String | No | No | The destination IP address of health check packets used in leased line disaster recovery and ECMP scenarios. | <p>This parameter is valid only for VRouter interfaces with a peer router interface on a VBR. Typically, you can use the IP address of a customer premises equipment (CPE) on the user side of the leased line, which is the IP address of the peer gateway on the VBR where the peer router interface is located. You can also specify another IP address on the user side of the leased line as the destination IP address.</p> <p>The HealthCheckSourceIp and HealthCheckTargetIp parameters must either both be specified or both left unspecified.</p> |

| Property | Type | Required | Editable | Description | Constraint |
|--------------------|---------|----------|----------|--|---|
| PricingCycle | String | No | No | The billing cycle of the subscription. | Valid values: <ul style="list-style-type: none"> Month Year |
| Period | Number | No | No | The subscription duration. | <ul style="list-style-type: none"> Valid values when the PricingCycle parameter is set to Month: 1 to 9. Valid values when the PricingCycle parameter is set to Year: 1 to 3. |
| AutoPay | Boolean | No | No | Specifies whether to enable automatic payment. | Default value: false. Valid values: <ul style="list-style-type: none"> true false |
| InstanceChargeType | String | No | No | The billing method of the instance. | Valid values: <ul style="list-style-type: none"> Postpaid: pay-as-you-go Prepaid: subscription |

Response parameters

Fn::GetAtt

RouterInterfaceId: the ID of the router interface.

Examples

```
{
  "ROSTemplateFormatVersion": "2015-09-01",
  "Resources": {
    "RouterInterface": {
      "Type": "ALIYUN::VPC::RouterInterface",
      "Properties": {
        "Name": "RouterInterface_1",
        "Description": "VPC initiator RouterInterface",
        "RouterId": "vrt-2ze2i147e5n0bicoe****",
        "Role": "AcceptingSide",
        "OppositeRegionId": "cn-beijing",
        "HealthCheckSourceIp": "10.0.XX.XX",
        "HealthCheckTargetIp": "192.168.XX.XX"
      }
    }
  },
  "Outputs": {
    "RouterInterfaceId": {
      "Value": {"Fn::GetAtt": ["RouterInterface", "RouterInterfaceId"]}
    }
  }
}
```

6.Object Storage Service (OSS)

6.1. What is OSS?

Object Storage Service (OSS) is a secure, cost-effective, and highly reliable cloud storage service provided by Alibaba Cloud. It enables you to store a large amount of data in the cloud.

Compared with user-created server storage, OSS has outstanding advantages in reliability, security, cost-effectiveness, and data processing capabilities. OSS enables you to store and retrieve a variety of unstructured data objects, such as text, images, audios, and videos over the network at any time.

OSS is an object storage service based on key-value pairs. Files uploaded to OSS are stored as objects in buckets. You can obtain the content of an object based on the object key.

In OSS, you can perform the following operations:

- Create a bucket and upload objects to the bucket.
- Obtain an object URL from OSS to share or download the object.
- Modify the attributes or metadata of a bucket or an object. You can also configure the ACL of the bucket or the object.
- Perform basic and advanced operations in the OSS console.
- Perform basic and advanced operations by using OSS SDKs or calling RESTful API operations in your application.

6.2. Usage notes

Before you use OSS, you must understand the following content:

To allow other users to use all or part of OSS features, you must create RAM users and grant permissions to the users by configuring RAM policies.

Before you use OSS, you must also understand the following limits.

| Item | Limit |
|----------------|---|
| Bucket | <ul style="list-style-type: none"> • You can create up to 100 buckets. • After a bucket is created, its name and region cannot be modified. |
| Upload objects | <ul style="list-style-type: none"> • Objects larger than 5 GB cannot be uploaded by using the following modes: console upload, simple upload, form upload, or append upload. To upload an object that is larger than 5 GB, you must use multipart upload. The size of an object uploaded by using multipart upload cannot exceed 48.8 TB. • If you upload an object that has the same name of an existing object in OSS, the new object will overwrite the existing object. |
| Delete objects | <ul style="list-style-type: none"> • Deleted objects cannot be recovered. • You can delete up to 100 objects at a time in the OSS console. To delete more than 100 objects at a time, you must call an API operation or use an SDK. |
| Lifecycle | You can configure up to 1,000 lifecycle rules for each bucket. |

6.3. Quick start

6.3.1. Log on to the OSS console

This topic describes how to log on to the OSS console.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- We recommend that you use Google Chrome.

Procedure

1. In the address bar, enter the URL used to access the Apsara Uni-manager Management Console. Press Enter.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

 **Note** When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login** to go to the Apsara Uni-manager Management Console.
4. In the top navigation bar, choose **Products > Object Storage Service**.

6.3.2. Create buckets

Objects uploaded to OSS are stored in a bucket. You must create a bucket before you upload objects to OSS.

Procedure

1. [Log on to the OSS console](#).
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click **Create Bucket**.
3. On the **Create OSS Bucket** page, configure parameters.

The following table describes the parameters that you can configure.

| Parameter | Description |
|---------------------|--|
| Organization | Select an organization from the drop-down list for the bucket. |
| Resource Set | Select a resource set from the drop-down list for the bucket. |

| Parameter | Description |
|----------------------------------|--|
| Region | <p>Select a region from the drop-down list for the bucket.</p> <p> Note</p> <ul style="list-style-type: none"> ◦ The region of a bucket cannot be changed after the bucket is created. ◦ If you want to access OSS from your ECS instance through the internal network, select the same region where your ECS instance is deployed. |
| Cluster | Select a cluster for the bucket. |
| Bucket Name | <p>Enter the name of the bucket.</p> <p> Note</p> <ul style="list-style-type: none"> ◦ The bucket name must comply with the naming conventions. ◦ The bucket name must be globally unique among all existing buckets in OSS. ◦ The bucket name cannot be changed after the bucket is created. |
| Storage Class | Set the value to Standard . Only Standard is supported. |
| Bucket Capacity | Specify the capacity of the bucket. Valid values: 0 to 2000000. Unit: TB or GB. |
| Access Control List (ACL) | <p>Set the ACL of the bucket. You can select the following options:</p> <ul style="list-style-type: none"> ◦ Private: Only the owner or authorized users of this bucket can read and write objects in the bucket. Other users, including anonymous users cannot access objects in the bucket without authorization. ◦ Public Read: Only the owner or authorized users of this bucket can read and write objects in the bucket. Other users, including anonymous users can only read objects in the bucket. ◦ Public Read/Write: All users, including anonymous users can read and write objects in the bucket. Fees incurred by such operations are paid by the owner of the bucket. Exercise caution when you configure this option. <p> Note You can modify the ACL of a bucket after the bucket is created. For more information, see Modify bucket ACLs.</p> |
| Server-Side Encryption | <p>Configure server-side encryption for the bucket. You can select the following options:</p> <ul style="list-style-type: none"> ◦ None: Server-side encryption is not performed. ◦ AES256: AES256 is used to encrypt each object in the bucket using different data keys. Customer master keys (CMKs) used to encrypt the data keys are rotated regularly. ◦ KMS: CMKs managed by KMS are used to encrypt objects in the bucket. |
| Encryption Algorithm | You can configure this parameter when you select KMS for Server-Side Encryption . |

| Parameter | Description |
|-----------|---|
| Key ID | <p>You can configure this parameter when you select KMS for Server-Side encryption. OSS uses the specified CMK to encrypt objects in the bucket.</p> <div style="background-color: #e0f2f7; padding: 5px;"> <p> Note To select a CMK ID for server-side encryption, you must create the CMK in the KMS console.</p> </div> |

4. Click **Submit**.

6.3.3. Upload objects

After you create a bucket, you can upload objects to it.

Prerequisites

A bucket is created. For more information about how to create a bucket, see [Create buckets](#).

Context

You can upload an object of any format to a bucket. You can use the OSS console to upload an object up to 5 GB in size. To upload an object larger than 5 GB, use OSS SDKs or call an API operation.

Procedure

1. [Log on to the OSS console](#).
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket to which you want to upload objects.
3. On the bucket details page, click the **Files** tab.
4. Click **Upload**.
5. In the **Upload** panel, set the parameters described in the following table.

| Parameter | Description |
|-----------|---|
| Upload To | <p>Set the folder to which objects are uploaded.</p> <ul style="list-style-type: none"> ◦ Current: Objects are uploaded to the current folder. ◦ Specified: Objects are uploaded to the specified folder. You must enter a folder name. If the specified folder does not exist, OSS automatically creates the specified folder and uploads the object to the folder. |
| File ACL | <p>Set the ACL of the object to upload. Default value: Inherited from Bucket.</p> <ul style="list-style-type: none"> ◦ Inherited from Bucket: The ACL of uploaded objects is the same as that of the bucket. ◦ Private: Only the owner or authorized users can read and write objects in the bucket. Other users, including anonymous users cannot access the objects in the bucket without authorization. ◦ Public Read: Only the bucket owner can perform write operations on objects in the bucket. Other users, including anonymous users, can perform only read operations on objects in the bucket. ◦ Public Read/Write: All users, including anonymous users, can read and write objects in the bucket. |

| Parameter | Description |
|-----------|---|
| Upload | <p>Drag one or more objects to upload to this section, or click Upload to select one or more objects to upload.</p> <div style="background-color: #e1f5fe; padding: 10px; border: 1px solid #cfcfcf;"> <p> Notice</p> <ul style="list-style-type: none"> ◦ When the uploaded object has the same name as an existing object in the bucket, the existing object is overwritten. ◦ If you upload a folder, only the files in the folder are uploaded and stored in the same folder in the bucket. ◦ The name of an uploaded object must comply with the following conventions: <ul style="list-style-type: none"> ▪ The name can contain only UTF-8 characters. ▪ The name is case-sensitive. ▪ The name must be 1 to 1,023 bytes in length. ▪ The name cannot start with a forward slash (/) or backslash (\). </div> |

6. In the **Upload Tasks** panel, wait until the upload task is complete.

Do not refresh or close the **Upload Tasks** panel when objects are being uploaded. Otherwise, the upload tasks are interrupted.

6.3.4. Obtain object URLs

You can obtain the URL of an uploaded object and share the URL with other users to preview or download the object.

Prerequisites

An object is uploaded to the bucket. For more information, see [Upload objects](#).

Procedure

1. [Log on to the OSS console](#).
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket in which the object whose URL you want to obtain is stored.
3. Click the **Files** tab.
4. Click the name of the object whose URL you want to obtain. In the **View Details** panel, click **Copy File URL**.

6.4. Buckets

6.4.1. View bucket information

You can view the detailed information about the created buckets in the OSS console.

Prerequisites

A bucket is created. For more information about how to create a bucket, see [Create buckets](#).

Procedure

1. [Log on to the OSS console](#).
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket that you

want to view.

3. On the **Overview** tab, you can view the information about the bucket, including the organization, resource set, endpoints, and basic settings.

6.4.2. Delete a bucket

You can delete a bucket in the OSS console.

Prerequisites

All objects and parts in the bucket are deleted. For more information, see [Delete objects](#) and [Manage parts](#).

 **Warning** Deleted objects, parts, and buckets cannot be recovered. Exercise caution when you delete objects, parts, and buckets.

Procedure

1. [Log on to the OSS console](#).
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket that you want to delete.
3. On the bucket details page that appears, click **Delete Bucket** in the upper right corner. In the message that appears, click **OK**.

6.4.3. Modify bucket ACLs

OSS provides access control list (ACL) to control access to buckets. By default, the ACL of bucket is private. You can modify the ACL of the bucket after it is created.

Prerequisites

A bucket is created. For more information about how to create a bucket, see [Create buckets](#).

Context

You can set the ACL of a bucket to one of the following values:

- **Private:** Only the owner or authorized users of this bucket can read and write objects in the bucket. Other users, including anonymous users cannot access the objects in the bucket without authorization.
- **Public Read:** Only the owner or authorized users of this bucket can write objects in the bucket. Other users, including anonymous users can only read objects in the bucket.
- **Public Read/Write:** Any users, including anonymous users can read and write objects in the bucket.

 **Warning** If you set the ACL of a bucket to Public Read or Public Read/Write, other users can read the data in the bucket without authentication, which results in security risks. To ensure the security of your data, we recommend that you configure the ACL of your bucket to private.

Procedure

1. [Log on to the OSS console](#).
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket for which you want to modify the ACL.
3. Click the **Basic Settings** tab. Find the **Access Control List (ACL)** section.
4. Click **Configure**. Modify the bucket ACL.
5. Click **Save**.

6.4.4. Configure static website hosting

You can configure static website hosting for a bucket in the OSS console so that users can access the website by using the domain name of the bucket.

Prerequisites

A bucket is created. For more information about how to create a bucket, see [Create buckets](#).

Procedure

1. [Log on to the OSS console](#).
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket for which you want to configure static website hosting.
3. Click the **Basic Settings** tab. Find the **Static Pages** section.
4. Click **Configure** and then set the parameters described in the following table.

| Parameter | Description |
|--------------------|---|
| Default Homepage | <p>Specify an index page that functions similar to index.html. Only HTML objects can be set to the index page. Static website hosting is disabled if you do not specify this parameter.</p> <ul style="list-style-type: none"> ◦ If Subfolder Homepage is disabled, you must make sure that the index object exists in the root folder and is readable. ◦ If Subfolder Homepage is enabled, you must make sure that the index object exists in both the root folder and the subfolder and the index object is readable. <div style="background-color: #e6f2ff; padding: 5px; border: 1px solid #d9e1f2;"> <p> Note You can specify only one index object for Default Homepage. If you enable Subfolder Homepage, the index object for the subfolder homepage must have the same name as that for the root folder homepage. However, the content of the index objects can be different.</p> </div> |
| Default 404 Page | <p>Set the default 404 page that is displayed when the requested resource does not exist. Only the HTML, JPG, PNG, BMP, or WebP object in the root folder can be set to the default 404 page. Default 404 Page is disabled if you do not specify this parameter.</p> |
| Subfolder Homepage | <p>Specify whether to enable the subfolder homepage feature.</p> <ul style="list-style-type: none"> ◦ Disable: disables Subfolder Homepage. The default homepage of the root folder is displayed if you access the root domain name of a static website or any URL that ends with a forward slash (/) under this domain name. ◦ Enable: enables Subfolder Homepage. When you access the root domain name of a static website, the default homepage of the root folder is displayed. When you access a URL ending with a forward slash (/), the default homepage of the corresponding folder is displayed. For example, when you access the URL test.oss-cn-hangzhou.aliyuncs.com/subdir/, the default homepage of the subfolder is displayed if the index object exists in the <i>subdir/</i> folder. |

| Parameter | Description |
|--------------------|--|
| Subfolder 404 Rule | <p>This parameter is available if you enable Subfolder Homepage. You can configure this parameter to specify the result to return when you access an object whose name does not end with a forward slash (/) and the object does not exist. For example, if <i>subdir</i> does not exist when you access <code>test.oss-cn-hangzhou.aliyuncs.com/subdir</code>, the following rules apply:</p> <ul style="list-style-type: none"> ◦ Redirect: the default rule that checks whether <i>subdir/index object</i> exists. <ul style="list-style-type: none"> ▪ If the index object exists, HTTP status code 302 is returned with the Location header that specifies <code>test.oss-cn-hangzhou.aliyuncs.com/subdir/</code>. ▪ If the index object does not exist, the default 404 page is returned. If the default 404 page does not exist, HTTP status code 404 is returned. ◦ NoSuckKey: returns the default 404 page. If the default 404 page does not exist, HTTP status code 404 is returned. ◦ Index: the rule that checks whether <i>subdir/index document</i> exists. <ul style="list-style-type: none"> ▪ If the index document exists, the content of the index document is directly returned. ▪ If the index document does not exist, the default 404 page is returned. If the default 404 page does not exist, HTTP status code 404 is returned. |

5. Click **Save**.

6.4.5. Configure hotlink protection

You can configure hot link protection for a bucket in the OSS console to prevent data in your bucket from being accessed by unauthorized domain names.

Prerequisites

A bucket is created. For more information about how to create a bucket, see [Create buckets](#).

Context

The hot link protection feature allows you to configure a Referrer whitelist for a bucket. This way, only requests from domain names included in the Referrer whitelist can access data in the bucket. OSS allows you to configure Referrer whitelists based on the Referrer header field in HTTP and HTTPS requests.

After hot link protection is configured for a bucket, OSS verifies requests to objects in the bucket only when the requests are initiated from signed URLs or anonymous users. Requests that contain the Authorization field in the header are not verified.

Procedure

1. [Log on to the OSS console](#).
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket for which you want to configure hot link protection.
3. Click the **Basic Settings** tab. Find the **Hotlink Protection** section.
4. Click **Configure**. Configure the parameters.
 - Enter domain names or IP addresses in the **Referrer Whitelist** field. Separate multiple Referers by using line feed. You can use asterisks (*) and question marks (?) as wildcards. Example:
 - If you add `www.example.com` to the Referrer whitelist, requests sent from URLs that start with `www.example.com`, such as `www.example.com/123` and `www.example.com.cn` are allowed.
 - If you add `*www.example.com/` to the Referrer whitelist, requests sent from `http://www.example.com/` and `https://www.example.com/` are allowed.

- An asterisk (*) can be used as a wildcard to indicate zero or more characters. For example, if you add *.example.com to the Referer whitelist, requests sent from URLs such as *help.example.com* and *www.example.com* are allowed.
 - A question mark (?) can be used as a wildcard to indicate a single character. For example, if you add example?.com to the Referer whitelist, requests sent from URLs such as *examplea.com* and *exampleb.com* are allowed.
 - You can add domain names or IP addresses that include a port number, such as *www.example.com:8080* and *10.10.10.10:8080*, to the Referer whitelist.
- Select whether to turn on **Allow Empty Referer** to allow requests in which the Referer field is empty.

An HTTP or HTTPS request that contains an empty Referer indicates that the request does not contain the Referer field or the value of the Referer field is empty.

If you do not allow empty Referers fields, only HTTP or HTTPS requests which include an allowed Referer field can access the objects in the bucket.

 **Note** By default, if you use the bucket endpoint to preview an MP4 object, the browser sends a request that contains the Referer field and a request that does not contain the Referer field at the same time. Therefore, you must not only add the bucket endpoint to the Referer whitelist but also allow empty Referer fields. To use the bucket endpoint to preview an object of other formats, you need only to allow empty Referer fields.

- 5. Click **Save**.

6.4.6. Configure logging

When you access OSS, a large number of access logs are generated. You can use the logging feature to store OSS access logs in a specified bucket.

Prerequisites

A bucket is created. For more information about how to create a bucket, see [Create buckets](#).

Procedure

1. [Log on to the OSS console](#).
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket for which you want to configure logging.
3. Click the **Basic Settings** tab. Find the **Logging** section.
4. Click **Configure**. Turn on the **Logging** switch. Select **Destination Bucket** and set **Log Prefix**.
 - **Destination Bucket**: Select the name of the bucket used to store access logs from the drop-down list. You must be the owner of the selected bucket, and the bucket must be in the same region as the bucket for which logging is enabled.
 - **Log Prefix**: Enter the prefix and folder where the access logs are stored. If you specify *log/<TargetPrefix>* as the prefix, access logs are stored in the *log/* directory.
5. Click **Save**.

6.4.7. Configure CORS

You can configure cross-origin resource sharing (CORS) in the OSS console to enable cross-origin access.

Prerequisites

A bucket is created. For more information about how to create a bucket, see [Create buckets](#).

Context

OSS provides CORS over HTML5 to implement cross-origin access. When OSS receives a cross-origin request (or an OPTIONS request) for a bucket, OSS reads the CORS rules of the bucket and checks the relevant permissions. OSS matches the request with the rules one by one. When OSS finds the first match, OSS returns a corresponding header in the response. If no match is found, OSS does not include any CORS header in the response.

Procedure

1. [Log on to the OSS console.](#)
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket for which you want to configure CORS.
3. Click the **Basic Settings** tab. In the **Cross-Origin Resource Sharing (CORS)** section, click **Configure**.
4. Click **Create Rule**. In the **Create Rule** panel, configure the parameters described in the following table.

| Parameter | Required | Description |
|------------------------|----------|--|
| Sources | Yes | <p>Specifies the sources from which you want to allow cross-origin requests. Note the following rules when you configure the sources:</p> <ul style="list-style-type: none"> ◦ You can configure multiple rules for sources. Separate multiple rules with line feeds. ◦ The domain names must include the protocol name, such as HTTP or HTTPS. ◦ Asterisks (*) are supported as wildcards. Each rule can contain up to one asterisk (*). ◦ A domain name must contain the port number if the domain name does not use the default port. Example: https://www.example.com:8080. <p>The following examples show how to configure domain names:</p> <ul style="list-style-type: none"> ◦ To match a specified domain name, enter the full domain name. Example: https://www.example.com:8080. ◦ Use an asterisk (*) as a wildcard in the domain name to match second-level domains. Example: https://*.example.com. ◦ Enter only an asterisk (*) as the wildcard to match all domain names. |
| Allowed Methods | Yes | Specifies the cross-origin request methods that are allowed. |
| Allowed Headers | No | <p>Specifies the response headers for the allowed cross-origin requests. Take note of the following rules when you configure the allowed headers:</p> <ul style="list-style-type: none"> ◦ This parameter is in the key:value format and case-insensitive. Example: content-type:text/plain. ◦ You can configure multiple rules for allowed headers. Separate multiple rules with new lines. ◦ Each rule can contain up to one asterisk (*) as the wildcard. Set this parameter to an asterisk (*) if you do not have special requirements. |

| Parameter | Required | Description |
|-------------------------|----------|---|
| Exposed Headers | No | Specifies the response headers for allowed access requests from applications, such as an XMLHttpRequest object in JavaScript. Exposed headers cannot contain asterisks (*). |
| Cache Timeout (Seconds) | No | Specifies the time the browser can cache the response to a preflight (OPTIONS) request to a specific resource. |

 **Note** You can configure up to 10 rules for each bucket.

5. Click OK.

6.4.8. Configure lifecycle rules

You can configure lifecycle rules for a bucket and manage the rules in the OSS console.

Prerequisites

A bucket is created. For more information, see [Create buckets](#).

Context

Take note of the following items when you configure lifecycle rules for a bucket:

- After a lifecycle rule is configured, it is loaded within 24 hours and takes effect within 24 hours after it is loaded. Check the configurations of a rule before you save the rule.
- Objects that are deleted based on lifecycle rules cannot be recovered. Configure lifecycle rules based on your requirements.
- You can configure up to 100 lifecycle rules in the OSS console. To configure more than 100 rules, use ossutil or OSS SDKs.

Procedure

1. [Log on to the OSS console](#).
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket for which you want to configure lifecycle rules.
3. Click the **Basic Settings** tab. Find the **Lifecycle** section. Click **Configure**.
4. Click **Create Rule**. In the **Create Rule** panel, configure the parameters described in the following table.

| Parameter | Description |
|-----------------------|--|
| Basic Settings | |
| Status | Specify the status of the lifecycle rule. Valid values: Enabled and Disabled . |
| Applied To | <p>Select policies used to match objects with the rule. You can select Files with Specified Prefix or Whole Bucket. Files with Specified Prefix indicates that this rule applies to objects whose names contain a specified prefix. Whole Bucket indicates that this rule applies to all objects in the bucket.</p> <p> Note If you select Files with Specified Prefix, you can configure multiple lifecycle rules for objects whose names contain different prefixes. If you select Whole Bucket, only one lifecycle rule can be configured for the bucket.</p> |

| Parameter | Description |
|-----------------------|---|
| Prefix | If you set Applied To to Files with Specified Prefix , you must specify the prefix of the objects to which the rule applies. For example, if you want the rule applies to objects whose names start with <code>img</code> , enter <code>img</code> . |
| Tagging | Select tagging and configure tags. The rule applies only to objects that have the specified tags. Example: Select Files with Specified Prefix and set Prefix to <code>img</code> , Key to <code>a</code> , and Value to <code>1</code> . The rule applies to all objects that has <code>img</code> in their names and has the tag <code>a=1</code> . For more information about object tagging, see Configure object tagging . |
| Clear Policy | |
| File Lifecycle | Configure rules for objects to specify when objects expire. You can set File Lifecycle to Validity Period (Days) , Expiration Date , or Disabled . If you select Disabled , the configurations of File Lifecycle do not take effect. |
| Delete | Specify when objects expire based on Validity Period (Days) or Expiration Date that you set for File Lifecycle . Expired objects are deleted. <ul style="list-style-type: none"> ◦ Validity Period (Days): Specify the number of days to retain objects after they are last modified. The objects are deleted the next day after they expire. For example, if you set Validity Period (Days) to 30, objects that are last modified on January 1, 2019 are deleted on February 1, 2019. ◦ Expiration Date: Specify the expiration date. Objects that are last modified before this date expire and are deleted. For example, if you set Expiration Date to 2019-1-1, objects that are last modified before January 1, 2019 are deleted. |
| Delete Parts | |
| Part Lifecycle | Specify the operations to perform on expired parts. You can set Part Lifecycle to Validity Period (Days) , Expiration Data , or Disabled . If you select Disabled , the configurations of Part Lifecycle do not take effect. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 10px; margin-top: 10px;"> <p> Notice</p> <ul style="list-style-type: none"> ◦ You must configure at least one of File Lifecycle and Part Lifecycle. ◦ If you select Tagging, Part Lifecycle is unavailable. </div> |
| Delete | Specify when parts expire based on Validity Period (Days) or Expiration Data that you set for Part Lifecycle . Expired parts are deleted. You can configure this parameter in the same way as you configure the Delete parameter in Clear Policy . |

5. Click **OK**.

6.4.9. Configure storage quota

If the capacity of a bucket reaches the specified storage quota, write operations such as `PutObject`, `MultipartUpload`, `CopyObject`, `PostObject`, and `AppendObject` cannot be performed on the bucket. This topic describes how to configure the storage quota of a created bucket.

Prerequisites

A bucket is created. For more information, see [Create buckets](#).

Context

Take note of the following items when you configure the storage quota of a bucket:

- Before you configure the storage quota of a bucket, make sure that the quota does not limit your business because write operations cannot be performed if the bucket capacity reaches the quota.
- In general, it takes about an hour for OSS to determine whether the bucket capacity exceeds the storage quota. In some cases, it can take longer.

Procedure

1. [Log on to the OSS console.](#)
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket for which you want to configure storage quota.
3. Click the **Basic Settings** tab, find the **Storage Quota** section.
4. Click **Configure**. Turn on **Storage Quota** and set **Storage Quota**.

- Units: TB or GB.

- Valid values: -1 to 2000000

The default value is -1, which indicates that the bucket capacity is not limited.

5. Click **Save**.

6.4.10. Configure back-to-origin rules

If you access data in a bucket that has no back-to-origin rules configured and the data does not exist, 404 Not Found is returned. However, if you configure back-to-origin rules that contain the correct origin URL, you can obtain the data based on the back-to-origin rules.

Context

Back-to-origin supports the mirroring and redirection modes. You can configure back-to-origin rules for hot migration and specific request redirection.

- **Mirroring-based back-to-origin**

After mirroring-based back-to-origin rules are configured for a bucket, you can obtain an object in the bucket based on the rules when the requested object is not found. For example, when you perform the GetObject operation on an object and the object is not found, OSS retrieves the object based on the origin URL, returns the object, and then writes the object to OSS. Mirroring-based back-to-origin rules are used to seamlessly migrate data to OSS. This feature allows you to migrate a service that already runs on a user-created origin or in another cloud service to OSS without interrupting services.

- **Redirection-based back-to-origin**

After redirection-based back-to-origin rules are configured for a bucket, you can obtain an object in the bucket based on the rules when the requested object is not found. For example, when you perform the GetObject operation on an object and the object is not found, OSS redirects the request to the origin URL, and then a browser or client returns the content from the origin. You can use this feature to redirect requests for objects and develop various services based on redirection.

You can configure up to 20 back-to-origin rules, which are run in a sequence that they are configured.



Notice Back-to-origin rules and versioning cannot be configured for a bucket at the same time.

Configure a mirroring-based back-to-origin rule

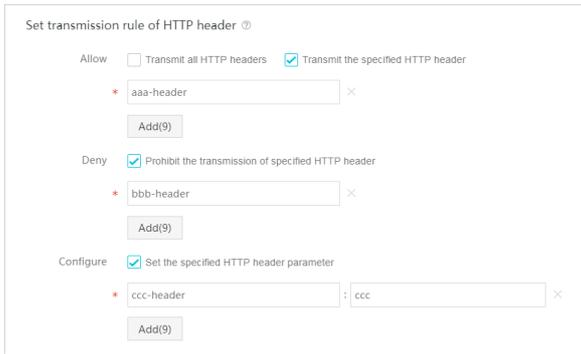
1. [Log on to the OSS console.](#)
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket for which you want to configure the mirroring-based back-to-origin rule.

3. Click the **Basic Settings** tab. In the **Back-to-Origin** section, click **Configure**.
4. Click **Create Rule**. In the **Create Rule** panel, set the parameters described in the following table to create a mirroring-based back-to-origin rule.

| Parameter | Required | Description |
|--------------------------------------|----------|---|
| Mode | Yes | Select Mirroring . In this mode, when a requested object cannot be found in OSS, OSS automatically retrieves the object from the origin, stores the object in OSS, and then returns the content to the requester. |
| Prerequisite | Yes | <p>Configure the conditions that trigger the back-to-origin rule. A rule is triggered only when all conditions are met.</p> <ul style="list-style-type: none"> ◦ HTTP Status Code: The back-to-origin rule is triggered when the specified HTTP status code is returned. The default HTTP status code is 404, which indicates that the rule is triggered when the requested object is not found in OSS and HTTP status code 404 is returned. By default, this option is selected when you select the Mirroring mode. ◦ File Name Prefix: The back-to-origin rule is triggered when the name of the requested object contains the specified prefix. For example, when this parameter is set to <code>abc/</code>, the back-to-origin rule is triggered when you access <code>https://bucketname.endpoint/abc/image.jpg</code>. ◦ File Name Suffix: The back-to-origin rule is triggered when the name of the requested object contains the specified suffix. For example, when this parameter is set to <code>.ida</code>, the back-to-origin rule is triggered when you access <code>https://bucketname.endpoint/image.jpg.ida</code>. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p>Note File Name Prefix and File Name Suffix are optional when you configure only one back-to-origin rule. When you configure multiple back-to-origin rules, you must differentiate the rules by specifying a different prefix or suffix for each of the rules.</p> </div> |
| Replace or Delete File Prefix | No | <p>You can configure this parameter after you select and configure File Name Prefix. When OSS sends a request to the origin, the content of File Name Prefix is replaced with that of Replace or Delete File Prefix.</p> <p>Example: You want to store the object obtained from the origin in the <code>mirror</code> folder under the root folder of the bucket. If the name of the requested object is <code>path/test/photo.jpg</code>, set File Name Prefix to <code>mirror/</code>, Replace or Delete File Prefix to <code>test/</code>, and the third column of Origin URL to <code>path</code>. In this case, when you access <code>https://bucketname.endpoint/mirror/photo.jpg</code>, if <code>photo.jpg</code> does not exist, OSS sends the <code>https://origin URL/path/test/photo.jpg</code> request to obtain this object. If the object is obtained from the origin, OSS stores this object in the <code>mirror/</code> folder and returns the object to you.</p> |

| Parameter | Required | Description |
|--------------------------------------|----------|--|
| Origin URL | Yes | <p>Configure the information about the origin URL.</p> <ul style="list-style-type: none"> First column: Select HTTP or HTTPS based on the protocol used by the origin. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin: 5px 0;"> <p> Note If SNI is enabled on the origin, HTTPS that you select does not take effect.</p> </div> <ul style="list-style-type: none"> Second column: Enter the domain name or IP address of the origin. Internal endpoints and IP addresses are not supported. If you enter an IP address, ensure that the origin corresponding to the IP address can be accessed by using the IP address. Third column: Enter the folder where the requested object is stored. Separate subfolders in a folder with forward slashes (/). Example: <code>abc/123</code>. |
| MD5 Verification | No | <p>If this option is selected, the MD5 hash of the object obtained from the origin is checked. When the response contains the Content-MD5 header value, OSS checks whether the MD5 hash of the object matches the Content-MD5 header value.</p> <ul style="list-style-type: none"> If the MD5 hash of the object matches the Content-MD5 header value, the client obtains the object, and OSS saves the object by using mirroring-based back-to-origin. If the MD5 hash of the object does not match the Content-MD5 header value, OSS calculates the Content-MD5 header value of the object based on the data integrity and does not save the object. However, the client can obtain the object because the object is returned to the client. |
| Keep Forward Slash in Origin URL | No | <p>OSS does not support object names that start with a forward slash (/). Therefore, when the name of the requested object starts with a forward slash (/), you must select this option to obtain the requested object based on the back-to-origin rules.</p> <p>For example, the origin is <code>https://www.example.com</code>, the name of the requested object is <code>/object.txt</code>, and the name of the bucket is <code>examplebucket</code>. If this option is selected, the default public endpoint to access the requested object is <code>https://examplebucket.endpoint/object.txt</code>, and the origin URL is <code>https://www.example.com/object.txt</code>. The object is obtained from OSS, returned to the client, and then saved in the bucket as <code>object.txt</code>.</p> |
| Other Parameter | No | <p>If this option is selected, the query string included in the request to OSS is transferred to the origin.</p> |
| 3xx Response | No | <p>If this option is selected, OSS follows the origin to direct the request to obtain the resource and stores the resource in buckets. If this option is not selected, OSS passes the 3xx response without obtaining the resource.</p> |
| Set Transmission Rule of HTTP Header | No | <p>You can configure the transmission rule for HTTP headers to customize transmission actions such as passthrough, filtering, or modification. For more information, see the following examples of transmission rule configurations for HTTP headers.</p> |

The following figure provides a sample transmission rule for HTTP headers.



You send a request that contains the following HTTP headers to OSS.

```
GET /object
host : bucket.oss-cn-hangzhou.aliyuncs.com
aaa-header : aaa
bbb-header : bbb
ccc-header : 111
```

After the mirroring-based back-to-origin rule is triggered, OSS sends the following request to the origin:

```
GET /object
host : source.com
aaa-header : aaa
ccc-header : ccc
```

Transmission rules do not support the following HTTP headers:

- o Headers that contain the following prefixes:
 - X-OSS-
 - OSS-
 - x-drs-
- o All standard HTTP headers. Examples:
 - content-length
 - authorization2
 - authorization
 - range
 - date

Configure a redirection-based back-to-origin rule

1. [Log on to the OSS console.](#)
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket for which you want to configure a redirection-based back-to-origin rule.
3. Click the **Basic Settings** tab. In the **Back-to-Origin** section, click **Configure**.
4. Click **Create Rule**. In the **Create Rule** panel, configure the parameters described in the following table to create a redirection-based back-to-origin rule.

| Parameter | Required | Description |
|-----------|----------|-------------|
|-----------|----------|-------------|

| Parameter | Required | Description |
|--------------------------------------|----------|--|
| Mode | Yes | Select Redirection . In this mode, OSS redirects requests that meet the prerequisites to the origin URL over HTTP, and then a browser or client returns the content from the origin to the requester. |
| Prerequisite | Yes | <p>Configure the conditions that trigger the back-to-origin rule. A rule is triggered only when all conditions are met.</p> <ul style="list-style-type: none"> ◦ HTTP Status Code: The back-to-origin rule is triggered when the specified HTTP status code is returned. The default HTTP status code is 404, which indicates that the rule is triggered when the requested object is not found in OSS and HTTP status code 404 is returned. ◦ File Name Prefix: The back-to-origin rule is triggered when the name of the requested object contains the specified prefix. For example, when this parameter is set to <code>abc/</code>, the back-to-origin rule is triggered when you access <code>https://bucketname.endpoint/abc/image.jpg</code>. ◦ File Name Suffix: The back-to-origin rule is triggered when the name of the requested object contains the specified suffix. For example, when this parameter is set to <code>.jpg</code>, the back-to-origin rule is triggered when you access <code>https://bucketname.endpoint/image.jpg</code>. <p>Note File Name Prefix and File Name Suffix are optional when you configure only one back-to-origin rule. When you configure multiple back-to-origin rules, you must differentiate the rules by specifying a different prefix or suffix for each of the rules.</p> |
| Replace or Delete File Prefix | No | <p>You can configure this parameter after you select and configure File Name Prefix. When OSS sends a request to the origin, the content of File Name Prefix is replaced with that of Replace File Name Prefix.</p> <p>Note If you select this option, you can only set Replace File Name Prefix for Origin URL.</p> |

| Parameter | Required | Description |
|------------------|----------|--|
| Origin URL | Yes | <p>Configure the information about the origin URL. You must configure the back-to-origin rule based on the origin URL.</p> <p>Configuring the origin URL in accordance with the following rules:</p> <ul style="list-style-type: none"> ◦ First column: Select HTTP or HTTPS based on the protocol used by the origin. ◦ Second column: Enter the domain name or IP address of the origin. <div style="background-color: #e6f2ff; padding: 5px; margin: 10px 0;"> <p> Note You can enter IP addresses only when the origin can be directly accessed by using IP addresses.</p> </div> <ul style="list-style-type: none"> ◦ Third column: Configure the redirection-based back-to-origin rule based on the prefix and suffix configurations. <ul style="list-style-type: none"> ▪ Add Prefix or Suffix: Add a prefix or suffix to the redirected URL. The prefix is configured in the third column. The suffix is configured in the fourth column. <p>You can select this option to configure information for the redirected URL when the URL of the requested data excludes a prefix or suffix. For example, the prefix is set to <code>123/</code>, and the suffix is set to <code>.img</code>. When you access <code>https://bucketname.endpoint/image</code>, the request is redirected to <code>https://Origin URL/123/image.jpg</code>.</p> ▪ Redirect to Fixed URL: Access to the requested object is redirected to a specified object. The object address is specified in the third column. Example: <code>abc/myphoto.jpg</code>. <div style="background-color: #e6f2ff; padding: 5px; margin: 10px 0;"> <p> Note You can also set the value of this parameter to the URL of a website. For example, if you set the URL to <code>https://www.aliyun.com/index.html</code>, access to the bucket is redirected to the Alibaba Cloud homepage.</p> </div> <ul style="list-style-type: none"> ▪ Replace File Name Prefix: If you set this parameter and File Name Prefix for Prerequisite, the prefix in the name of the object to which the access is redirected is replaced with that of the third column. If you set this parameter without setting File Name Prefix for Prerequisite, the specified prefix is added to the name of the object to which the access is redirected. <p>Separate subfolders in a folder with forward slashes (/). The folder name must end with a forward slash (/). The folder name cannot contain asterisks (*).</p> |
| Other Parameter | No | If this option is selected, the query string included in the request to OSS is transferred to the origin. |
| Redirection Code | Yes | You can select the redirection code from the drop-down list. Select Source from Alibaba Cloud CDN if the redirect request is from Alibaba Cloud CDN. |

5. Click OK.

6.4.11. Configure server-side encryption

OSS supports server-side encryption. When you upload an object to a bucket for which server-side encryption is enabled, OSS encrypts the object and stores the encrypted object. When you download the encrypted object from OSS, OSS automatically decrypts the object and returns the decrypted object to you. A header is added in the response to indicate that the object is encrypted on the OSS server.

Context

OSS supports the following encryption methods:

- Server-side encryption by using KMS (SSE-KMS)

OSS uses the default customer master key (CMK) managed by KMS or a specified CMK to encrypt objects. The CMK is managed by KMS to ensure confidentiality, integrity, and availability at minimal costs.

- Server-side encryption by using OSS-managed keys (SSE-OSS)

OSS uses data keys to encrypt objects and manages the data keys. In addition, OSS uses master keys that are regularly rotated to encrypt data keys.

You can enable server-side encryption in the OSS console by using one of the following methods:

- [Method 1: Enable server-side encryption when you create a bucket](#)
- [Method 2: Enable server-side encryption on the Basic Settings tab](#)

Method 1: Enable server-side encryption when you create a bucket

1. [Log on to the OSS console.](#)
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click **Create Bucket**.
3. On the **Create OSS Bucket** page, set parameters.

You can set the following parameters to configure server-side encryption for the bucket.

- **Server-Side Encryption:** Specify the encryption methods.
 - **None:** Server-side encryption is not performed.
 - **AES256:** AES256 is used to encrypt each object in the bucket by using different data keys. The CMKs used to encrypt the data keys are rotated regularly.
 - **SM4:** SM4 is used to encrypt each object in the bucket by using different data keys. The CMKs used to encrypt the data keys are rotated regularly.
 - **KMS:** CMKs managed by KMS are used to encrypt objects in the bucket.
- **Encryption Algorithm:** This parameter can be configured when you select **KMS** for **Server-Side Encryption**. You can select **SM4** or **AES256**.
- **Key ID:** This parameter can be configured when you select **KMS** for **Server-Side Encryption**. OSS uses the specified CMK to encrypt objects in the bucket.

 **Note** To select a CMK ID for server-side encryption, you must create the CMK in the KMS console.

4. Click **Submit**.

Method 2: Enable server-side encryption on the Basic Settings tab

1. [Log on to the OSS console.](#)
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket for which you want to configure server-side encryption.
3. Click the **Basic Settings** tab. Find the **Server-side Encryption** section.
4. Click **Configure** and set the following parameters:

- **Encryption Method:** Specify the encryption method.
 - **None:** Server-side encryption is not performed.
 - **OSS-Managed:** Keys managed by OSS are used to encrypt your data.
 - **KMS:** CMKs managed by KMS are used to encrypt objects in the bucket.
- **Encryption Algorithm:** You can select **SM4** or **AES256**.
 - **AES256:** AES256 is used to encrypt each object in the bucket by using different data keys. CMKs used to encrypt the data keys are rotated regularly.
 - **SM4:** SM4 is used to encrypt each object in the bucket by using different data keys. CMKs used to encrypt the data keys are rotated regularly.
- **CMK:** This parameter can be configured when you select **KMS** for **Encryption Method**. OSS uses the specified CMK to encrypt objects in the bucket.

 **Note** To select a CMK ID for server-side encryption, you must create the CMK in the KMS console.

5. Click **Save**.

 **Notice** The configurations of the default encryption method for a bucket do not affect the encryption configurations of existing objects within the bucket.

6.4.12. Bind a bucket to a VPC network

You can bind your bucket to a specified virtual private cloud (VPC) network to allow only requests from IP addresses within the VPC network to access your bucket.

Prerequisites

A VPC network is created. For more information, see the "Create a VPC" chapter of *VPC User Guide*.

Procedure

1. [Log on to the OSS console](#).
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket that you want to bind to the VPC network.
3. Click the **Overview** tab. Click **Bind VPC** in the **VPC Info** section.
4. On the **Bind VPC** page, select the VPC network that you create.
You can also click **Create VPC** to create a new VPC network.
5. Click **Submit**.

6.4.13. Configure CRR

Cross-region replication (CRR) allows you to perform automatic and asynchronous (near real-time) replication on objects across buckets in different regions. If you enable CRR, operations such as the creation, overwriting, and deletion of objects can be synchronized from the source bucket to the destination bucket.

Prerequisites

The source bucket and destination bucket are created. For more information, see [Create buckets](#).

Context

This feature meets the requirements of geo-disaster recovery or data replication. Objects in the destination bucket are extra replicas of objects in the source bucket. They have the same names, content, and metadata, such as the

created time, owner, user metadata, and object ACL.

Procedure

1. [Log on to the OSS console.](#)
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket for which you want to configure CRR.
3. On the bucket details page, click the **Basic Settings** tab. Find the **Cross-Region Replication** section.
4. Click **Enable**. In the **Cross-Region Replication** panel, configure the parameters described in the following table.

| Parameter | Description |
|----------------------------------|---|
| Source Region | The region where the current bucket is located. |
| Source Bucket | The name of the current bucket. |
| Destination Region | Select the region where the destination bucket is located. The source and destination buckets for CRR must be located in different regions. Data cannot be synchronized between buckets located within the same region. |
| Destination Bucket | Select the destination bucket to which data is synchronized. The source bucket and destination bucket specified in a CRR rule are not allowed to synchronize data with other buckets. For example, if you configure a CRR rule to synchronize data from Bucket A to Bucket B, Bucket A and Bucket B are not allowed to synchronize data with other buckets. |
| Applied To | Select the source data that you want to synchronize. <ul style="list-style-type: none"> ◦ All Files in Source Bucket: All objects within the source bucket are synchronized to the destination bucket. ◦ Files with Specified Prefix: Only objects whose names contain one of the specified prefixes are synchronized to the destination bucket. For example, if you have a folder named <i>management/</i> in the root folder of a bucket and want to synchronize objects in a subfolder named <i>abc/</i> in <i>management/</i>, you can enter the prefix <i>management/abc/</i>. You can specify up to 10 prefixes. |
| Operations | Select the synchronization policy. <ul style="list-style-type: none"> ◦ Add/Change: Only newly added and changed data is synchronized from the source bucket to the destination bucket. ◦ Add/Delete/Change: All changes to data including creation, modification, and deletion of objects are synchronized from the source bucket to the destination bucket. |
| Replicate Historical Data | Specify whether to synchronize historical data that is generated before you enable cross-cloud replication. <ul style="list-style-type: none"> ◦ Yes: Historical data is synchronized to the destination bucket. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin: 5px 0;"> <p> Notice When historical data is synchronized, objects in the destination bucket may be overwritten by historical data from the source bucket with the same name. Before you select this option, make sure that the data is consistent.</p> </div> <ul style="list-style-type: none"> ◦ No: Only objects that are uploaded or updated after CRR is enabled are synchronized to the destination bucket. |

5. Click **OK**.

 **Note**

- It takes about 3 to 5 minutes to take effect after CRR is configured. Synchronization information is displayed after the source bucket is synchronized.
- In CRR, data is asynchronously replicated in near-real time. Depending on how much data needs to be replicated, it can take a few minutes to several hours to replicate the data to the destination bucket.

6.5. Objects

6.5.1. Search for objects

You can search for objects whose names contain specific prefixes in buckets or folders in the OSS console.

Prerequisites

Objects are uploaded to the bucket. For more information, see [Upload objects](#).

Context

When you search for objects based on a prefix, search strings are case-sensitive and cannot contain forward slashes (/). You can search for objects only in the root folder of the current bucket or in the current folder. Subfolders and objects stored in subfolders cannot be searched.

Procedure

1. [Log on to the OSS console](#).
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket in which the objects that you want to search for are stored.
3. Click the **Files** tab.
4. Search for objects.
 - Search for objects or folders within the root folder of the bucket

In the upper-right corner, enter the prefix to search in the search box and press Enter or click the  icon to search for related objects. Objects and subfolders whose names contain the specified prefix within the root folder of the bucket are displayed.

- Search for objects or subfolders within a specified folder
Click the folder in which the objects or subfolders that you want to search for are stored. In the upper-right corner, enter the prefix to search in the search box and press Enter or click the  icon to search for related objects. Objects and subfolders whose names contain the specified prefix within the current folder are displayed.

6.5.2. Configure object ACLs

You can configure the ACL of an object in the OSS console to control access to the object.

Prerequisites

An object is uploaded to the bucket. For more information, see [Upload objects](#).

Procedure

1. [Log on to the OSS console](#).

- In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket that contains the object whose ACL you want to configure.
- In the left-side navigation pane, click **Files**.
- Click the name of the object whose ACL you want to configure. In the **View Details** panel, click **Set ACL** on the right side of **File ACL**.
You can also choose **More > Set ACL** in the Actions column corresponding to the object whose ACL you want to configure.
- In the **Set ACL** panel, configure the ACL of the object.
You can set the ACL of the object to one of the following values:
 - Inherited from Bucket**: The ACL of the object is the same as that of the bucket.
 - Private**: Only the owner or authorized users of this bucket can read and write objects in the bucket. Other users, including anonymous users cannot access the objects in the bucket without authorization.
 - Public Read**: Only the owner or authorized users of this bucket can write objects in the bucket. Other users, including anonymous users can only read objects in the bucket.
 - Public Read/Write**: All users, including anonymous users can read and write objects in the bucket. Fees incurred by such operations are charged to the owner of the bucket. Exercise caution when you set the object ACL to this value.
- Click **OK**.

6.5.3. Create folders

You can use the OSS console to create and simulate basic features of folders in Windows. This topic describes how to create a folder by using the OSS console.

Prerequisites

A bucket is created. For more information, see [Create buckets](#).

Context

OSS does not use a hierarchical structure for objects, but instead uses a flat structure. All elements are stored in buckets as objects. To facilitate object grouping and to simplify management, the OSS console displays objects whose names end with a forward slash (/) as folders. These objects can be uploaded and downloaded. You can use OSS folders in the OSS console in the same manner as you use folders in Windows.

 **Note** The OSS console displays objects whose names end with a forward slash (/) as folders, regardless of whether these objects contain data. The objects can only be downloaded by calling an API operation or by using OSS SDKs.

Procedure

- [Log on to the OSS console](#).
- In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket in which you want to create folders.
- Click the **Files** tab. On the page that appears, click **Create Folder**.
- In the **Create Folder** panel, enter the folder name.
The folder name must comply with the following conventions:
 - The name can contain only UTF-8 characters and cannot contain emojis.
 - The name cannot start with a forward slash (/) or backslash (\). The name cannot contain consecutive forward slashes (/). You can use forward slashes (/) in a folder name to quickly create a subfolder. For example, when you create a folder named *example/test/*, the folder named *example/* is created in the root

folder of the bucket and the subfolder named *test/* is created in the *example/* folder.

- The name cannot be two consecutive periods (..).
 - The folder name must be 1 to 254 characters in length.
5. Click **OK**.

6.5.4. Delete objects

You can delete uploaded objects in the OSS console when they are no longer needed.

Context

You can delete a single object or batch delete multiple objects. You can batch delete up to 100 objects. To delete specific objects or batch delete more than 100 objects, we recommend that you use API operations or OSS SDKs.

 **Notice** Deleted objects cannot be recovered. Exercise caution when you delete objects.

Procedure

1. [Log on to the OSS console](#).
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket in which the objects you want to delete are stored.
3. In the left-side navigation pane, click **Files**.
4. Select one or more objects that you want to delete in the object list, and then choose **Batch Operation > Delete**.
You can also choose **More > Completely Delete** in the Actions column corresponding to the object you want to delete.
5. In the dialog box that appears, click **OK**.

6.5.5. Manage parts

When you use multipart upload to upload an object, the object is split into several parts. After all of the parts are uploaded to the OSS server, you can call CompleteMultipartUpload to combine the parts into a complete object.

Context

You can also configure lifecycle rules to clear parts that are not needed on a regular basis. For more information, see [Manage lifecycle rules](#).

Procedure

1. [Log on to the OSS console](#).
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket in which the parts you want to delete are stored.
3. Click the **Files** tab. On the page that appears, click **Parts**.
4. In the **Parts** panel, delete the parts.
 - To delete all parts in the bucket, select all parts and then click **Delete All**.
 - To delete specific parts in the bucket, select these parts and then click **Delete**.
5. In the dialog box that appears, click **OK**.

6.5.6. Configure object tagging

You can configure object tagging to classify objects. Object tagging uses key-value pairs to identify objects. You can perform operations on multiple objects that have the same tag. For example, you can configure lifecycle rules for objects that have the same tag.

Context

Object tagging uses key-value pairs to identify objects. You can manage multiple objects that have the same tag. For example, you can configure lifecycle rules for objects that have the same tag or authorize Resource Access Management (RAM) users to access objects that have the same tag.

When you configure object tagging, take note of the following items:

- You can add up to 10 tags to an object. The tags added to an object must have unique tag keys.
- A tag key can be up to 128 bytes in length. A tag value can be up to 256 bytes in length.
- Tag keys and tag values are case-sensitive.
- The key and value of a tag can contain letters, digits, spaces, and the following special characters:
+ - = . _ : /
- Only the bucket owner and authorized users have read and write permissions on object tags. These permissions are independent of object access control lists (ACLs).
- In cross-region replication (CRR), object tags are also replicated to the destination bucket.

Procedure

1. [Log on to the OSS console.](#)
2. In the left-side navigation pane, click **Buckets**. On the Buckets page, click the name of the bucket in which the object that you want to configure tagging is stored.
3. On the bucket details page, click the **Files** tab.
4. Choose **More > Tagging** in the Actions column corresponding to the object to which you want to add tags.
5. In the **Tagging** panel, configure the **Key** and **Value** of the tag.
You can click **Add** to add up to more 10 tags to the object.
6. Click **OK**.

6.6. Create single tunnels

You can create single tunnels between OSS and a virtual private cloud (VPC) to access OSS resources from the VPC.

Prerequisites

A VPC and a vSwitch are created.

For more information, see the *Create a VPC* and *Create a vSwitch* topics in *VPC User Guide*.

Procedure

1. [Log on to the OSS console.](#)
2. In the left-side navigation panel, click **Create Single Tunnel**.
3. Click **Create**.
4. On the **Create Single Tunnel** page, configure the parameters described in the following table.

| Parameter | Required | Description |
|---------------------|----------|---|
| Organization | Yes | Select the organization of the VPC from which you want to access OSS resources. |

| Parameter | Required | Description |
|--------------|----------|--|
| Resource Set | Yes | After you select an organization, the resource set is automatically selected based on the organization. |
| Region | Yes | After you select an organization, a region is automatically selected based on the organization. |
| Description | No | Enter the description of the single tunnel you want to create. The description cannot exceed 180 characters in length. |
| VPC | Yes | Select the VPC that you created. You can also click Create VPC to create a VPC. |
| vSwitch | Yes | Select the vSwitch that you created. You can also click Create vSwitch to create a vSwitch. |

5. Click **Submit**.

6.7. Add OSS paths

You can add the paths of OSS resources in the console for quicker access.

Prerequisites

A bucket is created. For more information about how to create a bucket, see [Create buckets](#).

Procedure

1. [Log on to the OSS console](#).
2. Click the + icon on the right side of **My OSS Paths**.
3. In the **Add Authorized OSS Path** panel, add a path.

You can configure the following parameters to add a path.

- **Region:** Select the region of the bucket in the path that you want to add.
- **File Path:** Add the path of the resource that you want to access. The path is in the `bucket/object-prefix` format. For example, if the OSS resource that you want to access is the root folder of a bucket named *example*, set File Path to *example*. If the OSS resource that you want to access is the *test* folder in the root folder of the bucket named *example*, set File Path to *example/test/*.

7. Tablestore

7.1. What is Tablestore?

Tablestore is a NoSQL database service independently developed by Alibaba Cloud. Tablestore is a proprietary software program that is certified by the relevant authorities in China. Tablestore is built on the Apsara system of Alibaba Cloud, and can store large amounts of structured data and allow real-time access to the data.

Tablestore provides the following features:

- Offers schema-free data storage. You do not need to define attribute columns before you use them, or perform table-level changes to add or delete attribute columns. You can set the time to live (TTL) parameter for a table to manage the lifecycle of data. Expired data is deleted from the table.
- Uses the multi-node cluster architecture. Each management node on the platform implements a high availability mechanism. Therefore, the failure of a service node within a cluster does not affect the overall operating of businesses.
- Adopts the triplicate technology to keep three copies of data across different servers in different racks. A cluster can support single storage type instances (SSD only) or mixed storage type instances (SSD and SATA) to meet different budget and performance requirements.
- Adopts a fully redundant architecture to prevent single points of failure (SPOFs). The support for smooth online upgrades, hot cluster upgrades, and automatic data migration enables you to dynamically add or remove nodes for maintenance without incurring service interruptions. The concurrent read/write throughput and storage capacity can be linearly scaled. Each cluster can have no less than 500 hosts.
- Supports highly concurrent read/write operations. Concurrent read/write capabilities can be scaled out as the number of hosts increases. The read/write performance is indirectly related to the amount of data in a single table.
- Supports identity authentication and multi-tenancy. Comprehensive access control and isolation mechanisms are provided to safeguard your data. VPC and access over HTTPS are supported. Provides multiple authentication and authorization mechanisms so that you can define access permissions on individual tables and operations.

7.2. Precautions

Before you use Tablestore, you need to take note of the following precautions and limits.

The following table describes the limits for Tablestore. A part of the limits indicate the maximum allowable values rather than the suggested values. To ensure better performance, set the table scheme and data size in a single row based on actual conditions, and adjust the following configurations.

| Item | Limit | Description |
|--|----------------|---|
| The number of instances under an Apsara Stack tenant account | 1024 | To raise the limit, contact the technical support personnel. |
| The number of tables in an instance | 1024 | To raise the limit, contact the technical support personnel. |
| The length of an instance name | 3 to 16 bytes | The instance name can contain uppercase and lowercase letters, digits, and hyphens (-). It must start with a letter and cannot end with a hyphen (-). |
| The length of a table name | 1 to 255 bytes | The table name can contain uppercase and lowercase letters, digits, and underscores (_). It must start with a letter or underscore (_). |

| Item | Limit | Description |
|---|----------------|--|
| The length of a column name | 1 to 255 bytes | The column name can contain uppercase and lowercase letters, digits, and underscores (_). It must start with a letter or underscore (_). |
| The number of columns in a primary key | 1 to 4 | A primary key can contain one to four primary key columns. |
| The size of the value in a string type primary key column | 1 KB | The size of the value in a STRING primary key column cannot exceed 1 KB. |
| The size of the value in a STRING attribute column | 2 MB | The size of the value in a STRING attribute column cannot exceed 2 MB. |
| The size of the value in a BINARY primary key column | 1 KB | The size of the value in a BINARY primary key column cannot exceed 1 KB. |
| The size of the value in a BINARY attribute column | 2 MB | The size of the value in a BINARY attribute column cannot exceed 2 MB. |
| The number of attribute columns in a single row | Unlimited | A single row can contain an unlimited number of attribute columns. |
| The number of attribute columns written by one request | 1,024 | During a PutRow, UpdateRow, or BatchWriteRow operation, the number of attribute columns written in a row cannot exceed 1,024. |
| The data size of a row | Unlimited | The total size of all column names and column values for a row is unlimited. |

7.3. Quick start

7.3.1. Log on to the Tablestore console

This topic describes how to log on to the Tablestore console.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

Note When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top navigation bar, choose **Products > Tablestore**.

7.3.2. Create an instance

An instance is a logical entity in Tablestore and is used to manage tables. An instance is the basic unit of the resource management system of Tablestore. Tablestore controls application access and implements resource measurement at the instance level. This topic describes how to create an instance.

Procedure

1. [Log on to the Tablestore console](#).
2. On the **Overview** tab, click **Create Instance**.

Note You can create different instances to manage the associated tables for different business, or create different instances for development, testing, and production environments of the same business. By default, Tablestore allows you to create up to 1,024 instances and up to 1,024 tables in each instance that belongs to an Apsara Stack tenant account.

3. On the **Create Instance** page, configure parameters.

| Parameter | Description |
|----------------------|--|
| Region | Select a region from the drop-down list for the instance. |
| Organization | Select an organization from the drop-down list for the instance. |
| Resource Set | Select a resource set from the drop-down list for the instance. |
| Instance Name | Enter a name for the instance. Instance naming conventions: The name must be 3 to 16 characters in length and can contain only letters, digits, and hyphens (-). It must start with a letter and cannot start with case-insensitive string <code>ali</code> or <code>ots</code> . |
| Description | Enter a description for the instance. |
| Instance Type | Select an instance type from the drop-down list for the instance. Tablestore provides high-performance instances and capacity instances. The instance types vary based on the type of cluster you deploy. |

4. Click **Submit**.
5. In the **Submitted** dialog box, click **Back to Console**.
On the **Overview** tab, you can view the created instance.
After the instance is created, you can perform the following operations on the instance:

- Click the name of the instance or click **Manage Instance** in the **Actions** column that corresponds to the instance. On the **Instance Management** page, click each tab to perform various operations.
 - On the **Instance Details** tab, you can view the Instance Access URL, Basic Information, and Tables sections.
 - On the **Instance Monitoring** tab, you can view monitoring data by using time ranges, metric categories, and operations.
 - On the **Network Management** tab, you can bind or unbind virtual private clouds (VPCs) and view the list of VPCs.
- Click **Release** in the **Actions** column to release an instance.

 **Notice**

- Before you release an instance, ensure that all tables are deleted, and VPCs are unbound from instances.
- To create an instance when you release an existing instance, ensure that the name of the instance you want to create is different from that of the existing instance to avoid conflicts.

7.3.3. Create tables

This topic describes how to create a table in the Tablestore console.

Procedure

1. [Log on to the Tablestore console](#).
2. On the **Overview** page, click the name of the required instance or click **Manage Instance** in the **Actions** column corresponding to the instance.
3. On the **Instance Details** tab, click **Create Table**.

 **Note** You can create a maximum of 1,024 tables in each instance.

4. In the **Create Table** dialog box, set **Table Name** and **Primary Key**.

The following table describes the parameters you can configure.

| Parameter | Description |
|------------|--|
| Table Name | The name of the table. This name is used to uniquely identify a table in an instance. The name must be 1 to 255 bytes in length and can contain letters, digits, and underscores (_). The name must start with a letter or an underscore (_). |

| Parameter | Description |
|-------------|---|
| Primary Key | <p>One or more primary key columns in the table that uniquely identify each record in the table.</p> <p>Enter a primary key name and select a data type. Click Add a Primary Key to add a primary key column.</p> <p>You can add one to four primary key columns. By default, the first primary key column is the partition key. The configurations and order of primary key columns cannot be modified after the table is created.</p> <div style="background-color: #e0f2f1; padding: 10px; border: 1px solid #ccc;"> <p> Note</p> <ul style="list-style-type: none"> ◦ In Tablestore, only a primary key column can be used as an auto-increment primary key column. Partition keys cannot be used as auto-increment primary key columns. ◦ After a primary key column is set to an auto-increment primary key column, Tablestore automatically generates a value for the auto-increment primary key column when you write a row of data. You do not need to specify a value for the auto-increment primary key column. The values of auto-increment primary key columns are incremental and unique within the rows that share the same partition key. </div> <ul style="list-style-type: none"> ◦ Naming conventions of primary key columns: The name must be 1 to 255 bytes in length and can contain letters, digits, and underscores (_). The name must start with a letter or underscore (_) ◦ Data types supported by primary key columns are String, Integer, and Binary. |

5. Optional. Configure advanced parameters.

If you need to configure parameters such as Time to Live and Max Versions, perform this operation.

- i. Turn on **Advanced Settings**.

ii. Configure advanced parameters.

The following table describes the advanced parameters you can configure.

| Parameter | Description |
|---------------------------|--|
| Time to Live | The period for which data in the table can be retained. When the retention period exceeds the Time to Live (TTL) value, the system deletes the expired data. The minimum TTL value is 86,400 seconds (one day). A value of -1 indicates that data never expires. |
| Max Versions | The maximum number of versions of data that can be retained for an attribute column. When the versions of data in an attribute column exceed the Max Versions value, the system deletes the earliest versions of data to keep the maximum number of versions equal to the Max Versions value. Valid values: 1 to 10. |
| Max Version Offset | The difference between the version number and the data written time must be within the value of Max Version Offset. Otherwise, an error occurs when the data is written. Unit: seconds. The valid version range for attribute columns is calculated based on the following formula: Valid version range = [Data written time - Max version offset value, Data written time + Max version offset value). |
| Reserved Read Throughput | You can set this parameter only for high-performance instances. The read and write throughput that is allocated and reserved for the table. |
| Reserved Write Throughput | Valid values: integers from 0 to 5000. When the specified reserved read and write throughput is 0, Tablestore does not reserve related resources for the table. |

6. Optional. Create secondary indexes.

If you need to create secondary indexes, perform this operation.

i. Turn on **Create Secondary Index**.

ii. Click the **+ Add** button in the Pre-defined Column section. Enter the name of the pre-defined column and select a data type from the drop-down list.

- This operation is performed to create a predefined column for the base table. Tablestore uses a schema-free model. You can write an unlimited number of columns to a row and do not need to specify a fixed number of predefined columns in a schema. When you create a table, you can also predefine columns and specify their data types.
- You can add up to 14 predefined columns. To delete the predefined column you add, click the  icon on the left of the corresponding predefined column.
- The name of a predefined column must be 1 to 255 bytes in length and can contain letters, digits, and underscores (_). The name must start with a letter or underscore (_).
- The data types of predefined columns include STRING, INTEGER, BINARY, FLOAT, and BOOLEAN.

- iii. Click **Add Secondary Index**. Enter Index Name and set Primary Key and Pre-defined Column for the index table.
 - The name of an index table must be 1 to 255 bytes in length and can contain letters, digits, and underscores (_). The name must start with a letter or underscore (_).
 - You can set the primary key of the index table to the primary key or predefined columns of the base table.
 - Pre-defined Column is optional. You can set the predefined columns of the index table to only the predefined columns of the base table.
7. Click **OK**.

After a table is created, you can view the table in the **Table List** section. If the created table is not displayed in the list of tables, click the  icon to refresh the list of tables.

After a table is created, you can perform the following operations on the table:

- Click the name of the table or click **Details** in the Actions column. On the **Manage Table** page, you can perform the following operations:
 - On the **Details** tab, you can view the description of the table and the primary key columns list, and modify the attributes of the table.
 - On the **Data Editor** tab, you can insert or update data, query data, view data details, and delete multiple data at a time.
- Click the  icon in the Actions column corresponding to a table and choose **Delete** from the shortcut menu. Click **OK** in the Delete Table dialog box. The table is deleted.

 **Notice** If you delete a table, the table and the data in the table are permanently deleted from Tablestore and cannot be recovered. Exercise caution when you perform this operation.

7.3.4. Read and write data in the console

After a table is created, you can read data from and write data to the table in the console.

Add data

1. [Log on to the Tablestore console](#).
2. On the **Overview** page, click the name of the required instance or click **Manage Instance** in the Actions column corresponding to the instance.
3. In the **Table List** section of the **Instance Details** tab, click the name of the target table and click the **Data Editor** tab. You can also click **Data Editor** in the Actions column.
4. On the **Data Editor** tab, click **Insert**.
5. In the **Insert** dialog box that appears, set **Primary Key Value**. Click **Add Column**. Set **Name**, **Type**, **Value**, and **Version**.

By default, **System Time** is selected, indicating that the current system time is used as the version number of the data. You can also clear **System Time** and enter the version number of the data.

6. Click **OK**.

Rows that contain the written data are displayed on the **Data Editor** tab.

Update data

You can update data in the attribute columns of a row.

1. [Log on to the Tablestore console](#).

2. On the **Overview** page, click the name of the required instance or click **Manage Instance** in the Actions column corresponding to the instance.
3. In the **Table List** section of the **Instance Details** tab, click the name of the target table and click the **Data Editor** tab. You can also click **Data Editor** in the Actions column.
4. On the **Data Editor** tab, select the row of data to update. Click **Update**.
5. In the **Update** dialog box that appears, modify the type and value for the primary key, add or remove attribute columns, and update or delete data in attribute columns.
 - You can click **+Add Column** to add an attribute column. You can also click the  icon to delete an attribute column.
 - If you select **Update**, you can modify data in attribute columns. If you select **Delete**, select the number of version to delete. If you select **Delete All**, all versions of the data are deleted.
6. Click **OK**.

Query data

In the Tablestore console, you can query data in a single row (GetRow) or query data within a specified range (RangeQuery).

To query data in a single row, perform the following operations:

1. [Log on to the Tablestore console](#).
2. On the **Overview** page, click the name of the required instance or click **Manage Instance** in the Actions column corresponding to the instance.
3. In the **Table List** section of the **Instance Details** tab, click the name of the target table and click the **Data Editor** tab. You can also click **Data Editor** in the Actions column.
4. On the **Data Editor** tab, click **Search**.
5. Set filter conditions.
 - i. In the **Search** dialog box, Set Modes to **Get Row**.
 - ii. By default, the system returns all columns. To return specified attribute columns, turn off **All Columns**. Enter the names of the attribute columns to return. Separate the names of the attribute columns with commas (,).
 - iii. Set **Primary Key Value**.

The integrity and accuracy of the primary key value affect the query results.
 - iv. Set **Count of Versions** to specify the maximum number of versions to return.
6. Click **OK**.

Data that meets the filter conditions is displayed on the **Data Editor** tab.

To perform range query, perform the following steps:

1. [Log on to the Tablestore console](#).
2. On the **Overview** page, click the name of the required instance or click **Manage Instance** in the Actions column corresponding to the instance.
3. In the **Table List** section of the **Instance Details** tab, click the name of the target table and click the **Data Editor** tab. You can also click **Data Editor** in the Actions column.
4. On the **Data Editor** tab, click **Search**.
5. Set filter conditions.
 - i. In the **Search** dialog box, Set Modes to **Range Search**.
 - ii. By default, the system returns all columns. To return specified attribute columns, turn off **All Columns**. Enter the names of the attribute columns to return. Separate the names of the attribute columns with commas (,).

- iii. Set **Start Primary Key Column** and **End Primary Key Column**.

You can set **Start Primary Key Column** to **Min Value** or **Custom** and **End Primary Key Column** to **Max Value** or **Custom**. If you select **Custom**, enter a custom value.

 **Note**

- The value in the first primary key column takes priority when the range query mode is used. When the minimum and maximum values for the first primary key column are the same, the system uses the value in the second primary key column to perform the query. The query rules for the subsequent primary keys are the same as those for the first two primary keys.
- The **Custom** range is a left-open and right-closed interval.

- iv. Set **Count of Versions** to specify the maximum number of versions to return.
 - v. Set **Sequence** to **Forward Search** or **Backward Search**.
6. Click **OK**.

Data that meets the filter conditions is displayed based on the specified order on the **Data Editor** tab.

Delete data

You can delete data you no longer need.

1. [Log on to the Tablestore console](#).
2. On the **Overview** page, click the name of the required instance or click **Manage Instance** in the Actions column corresponding to the instance.
3. In the **Table List** section of the **Instance Details** tab, click the name of the target table and click the **Data Editor** tab. You can also click **Data Editor** in the Actions column.
4. On the **Data Editor** tab, select the row of data you want to delete. Click **Delete**.
5. In the **Delete** message that appears, click **OK**.

7.3.5. Bind a VPC to a Tablestore instance

After you bind a VPC to a Tablestore instance, you can access the Tablestore instance from the ECS instances in the VPC in the same region.

Prerequisites

- A VPC that is within the same region as the Tablestore instance is created.
- After the VPC is created, create an ECS instance in the VPC.

Procedure

1. [Log on to the Tablestore console](#).
2. On the **Overview** page, click the name of the required instance or click **Manage Instance** in the Actions column corresponding to the instance.
3. Click the **Network Management** tab.
4. On the **Network Management** tab, click **Bind VPC**.
5. In the **Bind VPC** dialog box, select a VPC and switch, enter **Instance VPC Name**.

The name of a VPC can contain only letters and digits and must start with a letter. The name must be 3 to 16 bytes in length.

6. Click **OK**.

After the VPC is bound to the instance, you can view the information of the VPC in the **VPC List** on the **Network Management** tab. You can use the VPC address to access the Tablestore instance from the ECS instances in the VPC.

After you bind a VPC, you can perform the following operations:

- Click **VPC Instance List** in the Actions column to view the VPC instances list, which contains the instance name, instance VPC name, and VPC domain name.
- Click **Unbind** in the Actions column to unbind the VPC from the instance. After the VPC is unbound, the ECS instance in the VPC can no longer access the Tablestore instance by using the VPC address. To access the Tablestore instance from the ECS instance, you must bind the VPC to the Tablestore instance again.

8. ApsaraDB RDS for MySQL

8.1. What is ApsaraDB RDS?

ApsaraDB Relational Database Service (RDS) is a stable, reliable, and scalable online database service. Based on the distributed file system and high-performance storage, ApsaraDB RDS provides a set of solutions for disaster recovery, backup, restoration, monitoring, and migration.

ApsaraDB RDS supports four storage engines, which are MySQL, SQL Server, PolarDB, and PostgreSQL. You can create database instances based on these storage engines to meet your business requirements.

RDS MySQL

Originally based on a branch of MySQL, ApsaraDB RDS for MySQL provides excellent performance. It is a tried and tested solution that handled the high-volume concurrent traffic during Double 11. ApsaraDB RDS for MySQL provides basic features, such as whitelist configuration, backup and restoration, Transparent Data Encryption (TDE), data migration, and management for instances, accounts, and databases. ApsaraDB RDS for MySQL also provides the following advanced features:

- **Read-only instance:** In scenarios where ApsaraDB RDS handles a small number of write requests but a large number of read requests, you can create read-only instances to scale up the reading capability and increase the application throughput.
- **Read/write splitting:** The read/write splitting feature provides a read/write splitting endpoint. This endpoint enables an automatic link for the primary instance and all its read-only instances. An application can connect to the read/write splitting endpoint to read and write data. Write requests are distributed to the primary instance and read requests are distributed to read-only instances based on their weights. To scale up the reading capability of the system, you can add more read-only instances.

8.2. Log on to the ApsaraDB RDS console

This topic describes how to log on to the ApsaraDB RDS console.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

 **Note** When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.

4. In the top navigation bar, choose **Products > Database Services > ApsaraDB RDS**.

8.3. Quick start

8.3.1. Limits

To ensure instance stability and security, ApsaraDB RDS for MySQL has some service limits, as listed in the following table.

| Operation | Description |
|---------------------------------------|---|
| Instance parameters | Instance parameters can be modified by using the ApsaraDB RDS console or API operations. Due to security and stability considerations, only specific parameters can be modified. |
| Root permissions of databases | The root or system administrator permissions are not provided. |
| Database backup | <ul style="list-style-type: none"> Logical backup can be performed by using the command line interface (CLI) or graphical user interface (GUI). Physical backup can be performed only by using the ApsaraDB RDS console or API operations. |
| Database restoration | <ul style="list-style-type: none"> Logical restoration can be performed by using the CLI or GUI. Physical restoration can be performed only by using the ApsaraDB RDS console or API operations. |
| ApsaraDB RDS for MySQL storage engine | <p>Only InnoDB is supported.</p> <ul style="list-style-type: none"> To ensure performance and security, we recommend that you use the InnoDB storage engine. The TokuDB engine is not supported. Percona no longer provides support for TokuDB, which leads to bugs that cannot be fixed and business losses in extreme cases. The MyISAM engine is not supported. Due to the inherent shortcomings of the MyISAM engine, some data may be lost. Only some existing instances use the MyISAM engine. MyISAM engine tables in newly created instances are automatically converted to InnoDB engine tables. The Memory engine is not supported. Newly created Memory tables are automatically converted into InnoDB tables. |
| Database replication | ApsaraDB RDS for MySQL provides dual-node clusters based on a primary/secondary replication architecture. The secondary instances in this replication architecture are hidden and cannot be accessed directly. |
| Instance restart | Instances must be restarted by using the ApsaraDB RDS console or API operations. |
| Account and database management | ApsaraDB RDS for MySQL manages accounts and databases by using the ApsaraDB RDS console. ApsaraDB RDS for MySQL also allows you to create a privileged account to manage users, passwords, and databases. |
| Standard account | <ul style="list-style-type: none"> Authorization is not allowed. The ApsaraDB RDS console allows you to manage accounts and databases. Instances that support standard accounts also support privileged accounts. |

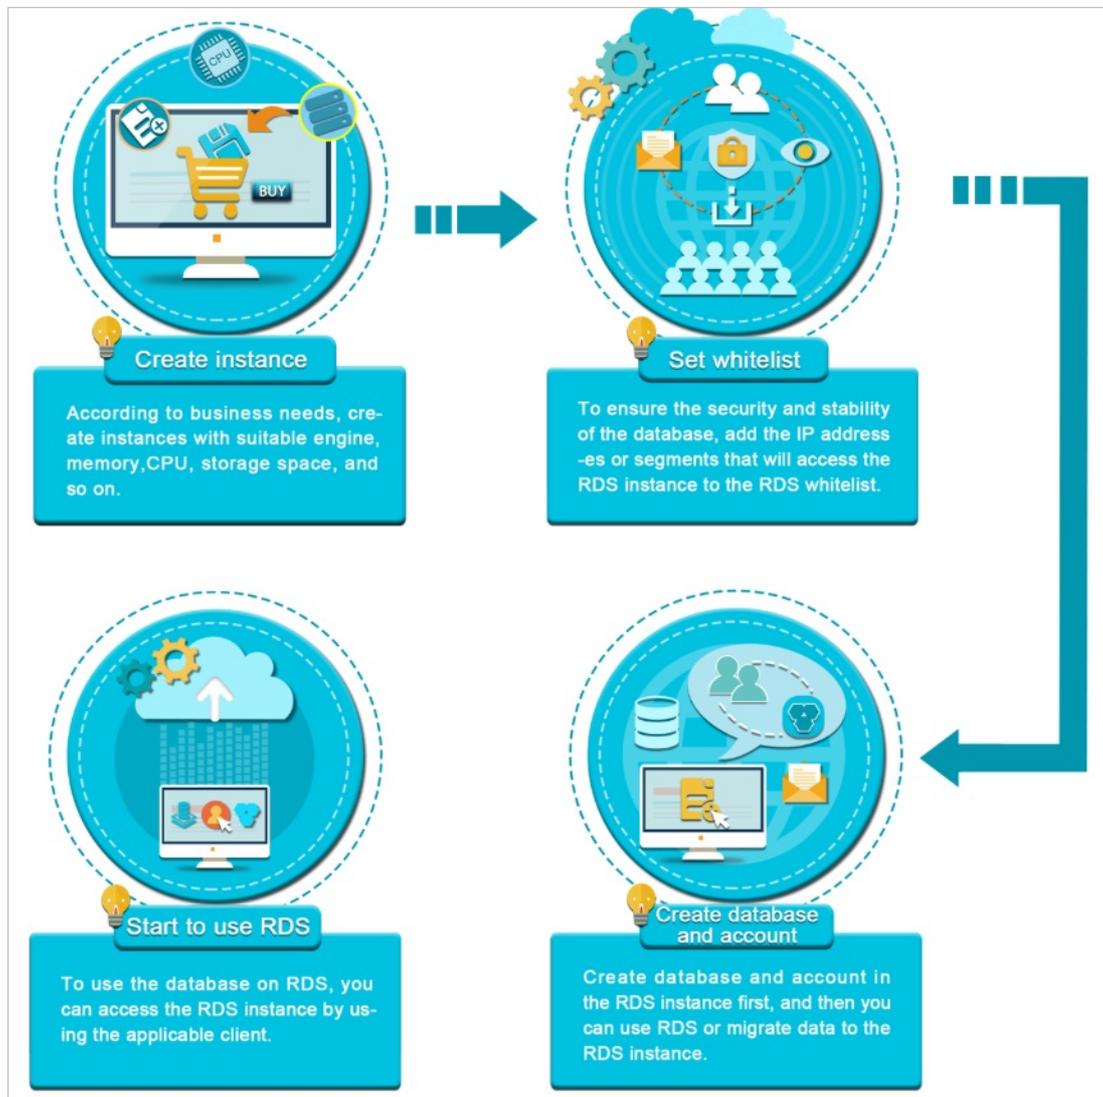
| Operation | Description |
|--------------------|--|
| Privileged account | <ul style="list-style-type: none"> • Authorization is allowed to standard accounts. • The ApsaraDB RDS console does not provide interfaces to manage accounts or databases. These operations can be performed only by using code or DMS. • The privileged account cannot be reverted to a standard account. |

8.3.2. Procedure

ApsaraDB RDS quick start covers the following operations: creating an instance, configuring a whitelist, creating a database, creating an account, and connecting to the instance. This topic describes how to use ApsaraDB RDS and provides all the necessary information to create an ApsaraDB RDS instance. ApsaraDB RDS for MySQL is used in the example.

Typically, after an instance is created, you must perform several operations to make the instance ready for use, as shown in [Quick start flowchart](#).

Quick start flowchart



- **Create an instance**

An instance is a virtual database server on which you can create and manage multiple databases.

- [Configure a whitelist](#)

After you create an ApsaraDB RDS instance, you must configure its whitelist to allow access from external devices.

Whitelists make your ApsaraDB RDS instance more secure. We recommend that you maintain whitelists on a regular basis. The whitelist configuration process does not affect the normal operations of the ApsaraDB RDS instance.

- [Create a database](#) and [Create an account](#)

Before you use a database, you must first create the database and an account in the ApsaraDB RDS instance.

- [Connect to an ApsaraDB RDS for MySQL instance](#)

After you create an ApsaraDB RDS instance, configure a whitelist, and create a database and an account, you can connect to the instance by using a database client.

8.3.3. Create an instance

This topic describes how to create one or more instances in the ApsaraDB RDS console.

Prerequisites

An Apsara Stack tenant account is created.

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, click **Create Instance** in the upper-right corner.
3. Configure the following parameters.

| Section | Parameter | Description |
|----------------|------------------------|---|
| Basic Settings | Organization | The organization to which the instance belongs. |
| | Resource Set | The resource set to which the instance belongs. |
| Region | Region | The region in which you want to create the instance. Services in different regions cannot communicate over an internal network. After you create an instance, its region cannot be changed. |
| | Zone of Primary Node | The zone where the primary instance is deployed. |
| | Deployment Method | Specifies whether to deploy the primary and secondary instances in separate zones. ApsaraDB RDS supports Multi-zone Deployment and Single-zone Deployment . If you select Multi-zone Deployment , you must configure Zone of Secondary Node . |
| | Zone of Secondary Node | The zone where the secondary instance is deployed. This parameter is available only when Deployment Method is set to Multi-zone Deployment . <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 5px;"> ? Note If you select the same zone for both the primary and secondary instances, the deployment is equivalent to single-zone deployment. </div> |
| | Quantity | The number of ApsaraDB RDS instances that you want to create. Default value: 1. |

| Section | Parameter | Description |
|----------------|------------------|---|
| Specifications | Instance Name | <p>The name of the instance.</p> <ul style="list-style-type: none"> The name must be 2 to 64 characters in length The name must start with a letter. The name can contain letters, digits, and the following special characters: _ - : The name cannot start with http:// or https://. |
| | Connection Type | <p>The connection type of the instance. ApsaraDB RDS instances support the following connection types:</p> <ul style="list-style-type: none"> Internet: ApsaraDB RDS instances of this connection type can be connected over the Internet. Internal Network: ApsaraDB RDS instances of this connection type can be connected over an internal network. <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e0f0ff;"> <p> Note The value of this parameter cannot be changed after the instance is created. Proceed with caution.</p> </div> |
| | Database Engine | The database engine of the instance. Set the value to MySQL . |
| | Engine Version | <p>The version of the database engine. Valid values:</p> <ul style="list-style-type: none"> 8.0 5.7 5.6 |
| | Edition | The edition of the instance. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| | Storage Type | The storage type of the instance. Select local SSD. |
| | Instance Type | The instance type of the instance. Memory size determines the maximum number of connections and the input/output operations per second (IOPS). The actual values are displayed in the console. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| | Storage Capacity | The storage capacity of the instance, which includes the space to store data, system files, binlog files, and transaction files. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |

| Section | Parameter | Description |
|---------|----------------------|---|
| Network | Network Type | The network type of the instance. ApsaraDB RDS instances support the following network types: <ul style="list-style-type: none"> Classic Network: Cloud services in the classic network are not isolated. Unauthorized access to a cloud service is blocked only by the security group or whitelist policy of the service. <div style="background-color: #e0f2f1; padding: 5px; margin: 5px 0;"> ? Note Instances that use standard SSDs cannot be deployed in the classic network. </div> <ul style="list-style-type: none"> VPC: A virtual private cloud (VPC) helps you build an isolated network environment on Apsara Stack. You can customize route tables, CIDR blocks, and gateways in a VPC. We recommend that you select VPC for improved security. |
| | VPC | Select a VPC. <div style="background-color: #e0f2f1; padding: 5px; margin: 5px 0;"> ? Note When Network Type is set to VPC, you must specify this parameter. </div> |
| | vSwitch | Select a vSwitch. <div style="background-color: #e0f2f1; padding: 5px; margin: 5px 0;"> ? Note When Network Type is set to VPC, you must specify this parameter. </div> |
| | IP Address Whitelist | The IP addresses that are allowed to connect to the instance. |

4. Click **Submit**.

8.3.4. Initialization settings

8.3.4.1. Configure a whitelist

To ensure database security and reliability, you must modify the whitelist of an ApsaraDB RDS instance before you enable the instance. You must add the IP addresses or CIDR blocks that are used for database access to the whitelist.

Context

The whitelist improves the access security of your ApsaraDB RDS instance. We recommend that you maintain the whitelist on a regular basis. The whitelist configuration process does not affect the normal operations of the ApsaraDB RDS instance.

To configure a whitelist, perform the following operations:

- Configure a whitelist: Add IP addresses to allow them to connect to the ApsaraDB RDS instance.
- Configure an ECS security group: Add an ECS security group for the ApsaraDB RDS instance to allow ECS instances in the group to connect to the ApsaraDB RDS instance.

Precautions

- The default whitelist can be modified or cleared, but cannot be deleted.
- You can add up to 1,000 IP addresses or CIDR blocks to a whitelist. If you want to add a large number of IP addresses, we recommend that you merge them into CIDR blocks, such as 192.168.1.0/24.

Configure a standard IP address whitelist

1. For more information, see [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find an instance. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
3. In the left-side navigation pane, click **Data Security**.
4. On the **Whitelist Settings** tab, click **Edit** corresponding to the **default** whitelist.

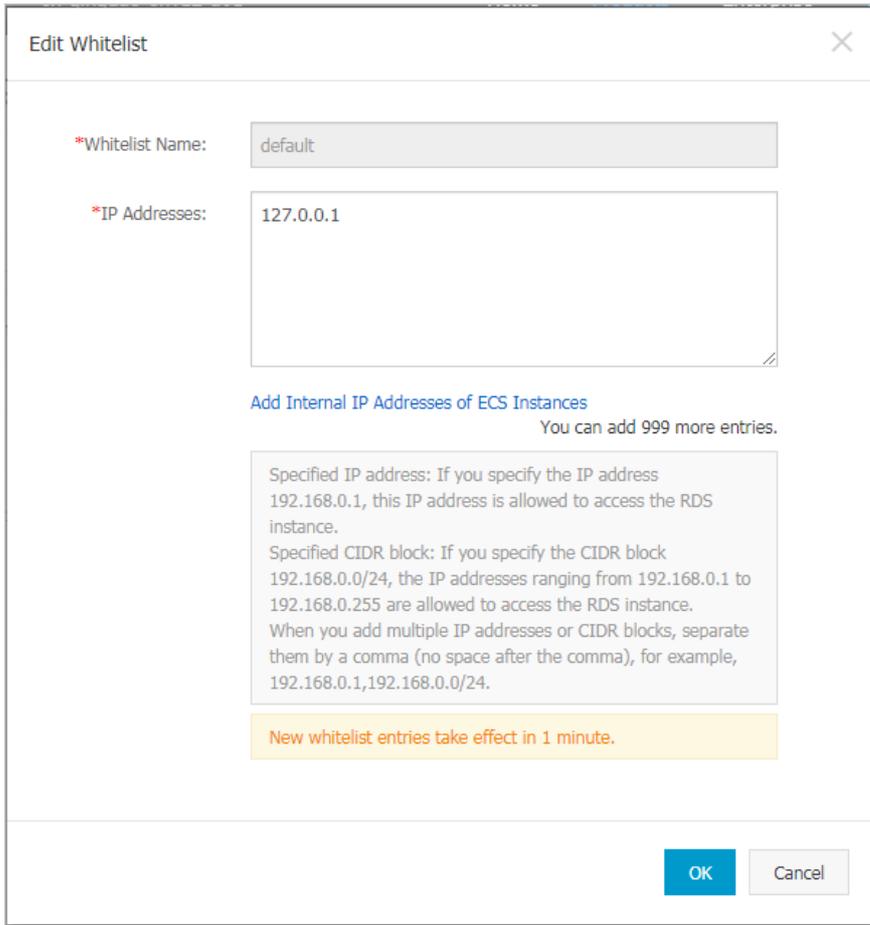


Note

- If you want to connect an ECS instance to an ApsaraDB RDS instance by using an internal endpoint, you must make sure that the two instances are in the same region and have the same network type. Otherwise, the connection fails.
- You can also click **Create Whitelist** to create a new whitelist.

5. In the **Edit Whitelist** dialog box, enter the IP addresses or CIDR blocks that are allowed to access your ApsaraDB RDS instance, and then click **OK**.
 - If you add the CIDR block 10.10.10.0/24, all IP addresses in the 10.10.10.X format are allowed to access the ApsaraDB RDS instance.
 - If you enter more than one IP address or CIDR block, you must separate them with commas (.). Do not add spaces before or after the commas. Example: 192.168.0.1,172.16.213.9.
 - If you click **Add Internal IP Addresses of ECS Instances**, the IP addresses of all of the ECS instances that are created in your Alibaba Cloud account appear. Then, you can select the required IP addresses and add them to the whitelist.

Note If you add a new IP address or CIDR block to the **default** whitelist, the default address 127.0.0.1 is deleted.



8.3.4.2. Create an account

After you create an ApsaraDB RDS instance and configure its IP address whitelist, you must create a database and an account on the instance. This topic describes how to create privileged and standard accounts.

Context

ApsaraDB RDS for MySQL supports two types of database accounts: privileged and standard. You can manage all your accounts and databases in the ApsaraDB RDS console. For more information about permissions that can be granted to each type of account, see [Account permissions](#).

| Account type | Description |
|---------------------------|---|
| Privileged account | <ul style="list-style-type: none"> You can create and manage privileged accounts by using the ApsaraDB RDS console or API operations. You can create only one privileged account on each ApsaraDB RDS instance. The privileged account can be used to manage all standard accounts and databases on the instance. A privileged account allows you to manage permissions to a fine-grained level. For example, you can grant each standard account the permissions to query specific tables. A privileged account has the permissions to disconnect all standard accounts on the instance. |

| Account type | Description |
|-------------------------|--|
| Standard account | <ul style="list-style-type: none"> You can create and manage standard accounts by using the ApsaraDB RDS console, API operations, or SQL statements. You can create up to 500 standard accounts on an instance. You must manually grant standard accounts the specific database permissions. You cannot use a standard account to create, manage, or disconnect other accounts from databases. |

| Account type | Maximum number of databases | Maximum number of tables | Maximum number of accounts |
|--------------------|-----------------------------|--------------------------|--|
| Privileged account | Unlimited | < 200,000 | Varies based on the kernel parameter settings of the instance. |
| Standard account | 500 | < 200,000 | Varies based on the kernel parameter settings of the instance. |

Create a privileged account

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Accounts**.
5. On the **Accounts** page, click the **Accounts** tab.
6. Click **Create Account**.



7. On the **Create Account** page, configure the following parameters.

| Parameter | Description |
|-------------------------|--|
| Database Account | Enter the name of the account. The account name must meet the following requirements: <ul style="list-style-type: none"> The name is 1 to 16 characters in length. The name starts with a lowercase letter and ends with a lowercase letter or digit. The name contains lowercase letters, digits, and underscores (_). |
| Account Type | Select Privileged Account. |
| Password | Enter the password of the account. The password must meet the following requirements: <ul style="list-style-type: none"> The password is 8 to 32 characters in length. The password contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. Special characters include ! @ # \$ % ^ & * () _ + - = |

| Parameter | Description |
|--------------------------|---|
| Re-enter Password | Enter the password of the account again. |
| Description | Optional. Enter information about the account to facilitate subsequent management. The description can be up to 256 characters in length. |

8. Click **Create**.

Reset the permissions of a privileged account

If an issue occurs on the privileged account, you can enter the password of the privileged account to reset permissions. For example, you can reset the permissions if the permissions are unexpectedly revoked.

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Accounts**.
5. On the **Accounts** page, click the **Accounts** tab.
6. On the **Accounts** tab, find the privileged account and click **Reset Permissions** in the **Actions** column.
7. On the **Initialize Account** page, enter the password of the privileged account and click **OK**.

Create a standard account

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Accounts**.
5. On the **Accounts** page, click the **Accounts** tab.
6. Click **Create Account**.



7. On the **Create Account** page, configure the following parameters.

| Parameter | Description |
|-------------------------|--|
| Database Account | Enter the name of the account. The account name must meet the following requirements: <ul style="list-style-type: none"> ◦ The name is 1 to 16 characters in length. ◦ The name starts with a lowercase letter and ends with a lowercase letter or digit. ◦ The name contains lowercase letters, digits, and underscores (_). |
| Account Type | Select Standard Account. |

| Parameter | Description |
|-----------------------------|---|
| Authorized Databases | <p>Select one or more databases on which you want to grant permissions to the account. You can also leave this parameter empty at this time and authorize databases after the account is created.</p> <ol style="list-style-type: none"> i. Select one or more databases from the Unauthorized Databases section and click Add to add them to the Authorized Databases section. ii. In the Authorized Databases section, select the Read/Write, Read-only, DDL Only, or DML Only permissions on each authorized database. <p>If you want to grant the same permissions on multiple databases to the account, click the button in the upper-right corner of the section. The button may appear as Set All to Read/Write.</p> |
| Password | <p>Enter the password of the account. The password must meet the following requirements:</p> <ul style="list-style-type: none"> ◦ The password is 8 to 32 characters in length. ◦ The password contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. ◦ Special characters include ! @ # \$ % ^ & * () _ + - = |
| Re-enter Password | Enter the password of the account again. |
| Description | Optional. Enter information about the account to facilitate subsequent management. The description can be up to 256 characters in length. |

8. Click **Create**.

Account permissions

| Account type | Authorization type | Permission | | | | |
|--------------------|-------------------------|--------------------|---|-------------|----------------|-------------------|
| | | Privileged account | - | SELECT | INSERT | UPDATE |
| DROP | RELOAD | | | PROCESS | REFERENCES | INDEX |
| ALTER | CREATE TEMPORARY TABLES | | | LOCK TABLES | EXECUTE | REPLICATION SLAVE |
| REPLICATION CLIENT | CREATE VIEW | | | SHOW VIEW | CREATE ROUTINE | ALTER ROUTINE |
| CREATE USER | EVENT | | | TRIGGER | - | - |

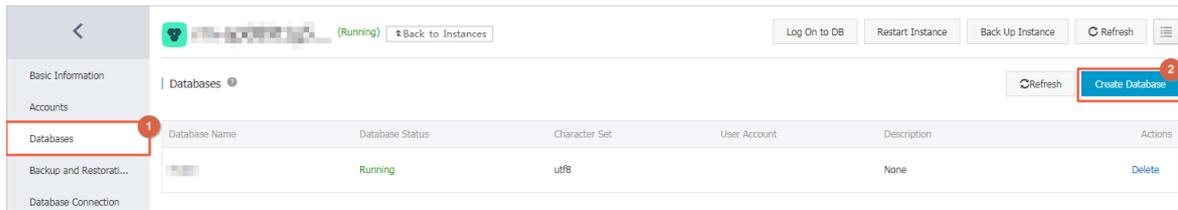
| Account type | Authorization type | Permission | | | | |
|------------------|--------------------|--------------------|-------------------|--------------------|----------------|-------------------------|
| Standard account | Read-only | SELECT | LOCK TABLES | SHOW VIEW | PROCESS | REPLICATION SLAVE |
| | | REPLICATION CLIENT | - | - | - | - |
| | Read/write | SELECT | INSERT | UPDATE | DELETE | CREATE |
| | | DROP | REFERENCES | INDEX | ALTER | CREATE TEMPORARY TABLES |
| | | LOCK TABLES | EXECUTE | CREATE VIEW | SHOW VIEW | CREATE ROUTINE |
| | | ALTER ROUTINE | EVENT | TRIGGER | PROCESS | REPLICATION SLAVE |
| | | REPLICATION CLIENT | - | - | - | - |
| | DDL-only | CREATE | DROP | INDEX | ALTER | CREATE TEMPORARY TABLES |
| | | LOCK TABLES | CREATE VIEW | SHOW VIEW | CREATE ROUTINE | ALTER ROUTINE |
| | | PROCESS | REPLICATION SLAVE | REPLICATION CLIENT | - | - |
| | DML-only | SELECT | INSERT | UPDATE | DELETE | CREATE TEMPORARY TABLES |
| | | LOCK TABLES | EXECUTE | SHOW VIEW | EVENT | TRIGGER |
| | | PROCESS | REPLICATION SLAVE | REPLICATION CLIENT | - | - |

8.3.4.3. Create a database

After you create an ApsaraDB RDS instance and configure its whitelist, you must create a database and an account in the instance.

Procedure

1. Log on to the ApsaraDB for RDS console.
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Databases**.
5. Click **Create Database**.



6. Configure the following parameters.

| Parameter | Description |
|---------------------------------|--|
| Database Name | <ul style="list-style-type: none"> The name must be 1 to 64 characters in length. The name must start with a letter and end with a letter or digit. The name can contain lowercase letters, digits, underscores (_), and hyphens (-). The name must be unique within the instance. |
| Supported Character Sets | Select utf8, gbk, latin1, utf8mb4, or all. If you want to use other character sets, select all, and then select the required character set from the list. |
| Description | Optional. Enter information about the database to facilitate subsequent management. The description can be up to 256 characters in length. |

7. Click **Create**.

8.3.5. Connect to an ApsaraDB RDS for MySQL instance

After you complete the initial configuration of your ApsaraDB RDS for MySQL instance, you can connect to it from an Elastic Compute Service (ECS) instance or an on-premises client.

Context

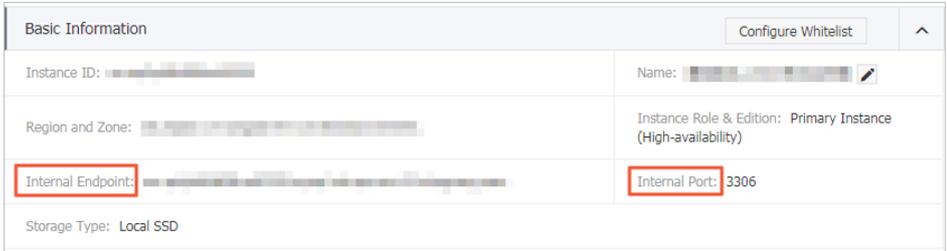
After you perform operations such as [Create an instance](#), [Configure a whitelist](#), and [Create an account](#), you can use a general database client or configure the endpoint, port number, and account information in an application to connect to the MySQL instance.

If you need to connect an ECS instance to an ApsaraDB RDS instance, you must make sure that both instances are in classic networks or in the same VPC, and the IP address of the ECS instance is correctly configured in the ApsaraDB RDS whitelist.

Connect to an instance from a client

ApsaraDB RDS for MySQL is fully compatible with open source MySQL. You can connect to an ApsaraDB RDS instance from a database client by using a method similar to the method that you use to connect to an open source MySQL database. In the following example, the [HeidiSQL](#) client is used.

1. Start the HeidiSQL client.
2. In the lower-left corner of the Session manager dialog box, click **New**.
3. Enter information about the ApsaraDB RDS instance that you want to connect. The following table describes the required parameters.

| Parameter | Description |
|-----------------------|--|
| Network type | Select the network type of the ApsaraDB RDS instance that you want to connect. For this example, select MariaDB or MySQL (TCP/IP) . |
| Host name / IP | <p>Enter the internal or public endpoint of the ApsaraDB RDS instance.</p> <ul style="list-style-type: none"> ◦ If your client is deployed on an ECS instance that is in the same region and has the same network type as the ApsaraDB RDS instance, use the internal endpoint. For example, if your ECS and ApsaraDB RDS instances are both in a VPC located in the China (Hangzhou) region, you can use the internal endpoint of the ApsaraDB RDS instance to create a secure connection. ◦ In other scenarios, use the public endpoint. <p>To view the internal and public endpoints and port numbers of the ApsaraDB RDS instance, perform the following operations:</p> <ol style="list-style-type: none"> Log on to the ApsaraDB for RDS console. Find the ApsaraDB RDS instance to which you want to connect and click its ID. In the Basic Information section, view the internal endpoint and internal port number of the instance.  |
| User | Enter the username of the account that you use to connect to the ApsaraDB RDS instance. |
| Password | Enter the password of the account. |
| Port | If you connect to the instance over an internal network, enter the internal port number of the instance. If you connect to the instance over the Internet, enter the public port number of the instance. |

4. Click **Open**. If the connection information is correct, you can connect to the instance.

8.4. Instances

8.4.1. Create an instance

This topic describes how to create one or more instances in the ApsaraDB RDS console.

Prerequisites

An Apsara Stack tenant account is created.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, click **Create Instance** in the upper-right corner.
3. Configure the following parameters.

| Section | Parameter | Description |
|----------------|------------------------|--|
| Basic Settings | Organization | The organization to which the instance belongs. |
| | Resource Set | The resource set to which the instance belongs. |
| Region | Region | The region in which you want to create the instance. Services in different regions cannot communicate over an internal network. After you create an instance, its region cannot be changed. |
| | Zone of Primary Node | The zone where the primary instance is deployed. |
| | Deployment Method | Specifies whether to deploy the primary and secondary instances in separate zones. ApsaraDB RDS supports Multi-zone Deployment and Single-zone Deployment . If you select Multi-zone Deployment , you must configure Zone of Secondary Node . |
| | Zone of Secondary Node | The zone where the secondary instance is deployed. This parameter is available only when Deployment Method is set to Multi-zone Deployment . <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e6f2ff;"> ? Note If you select the same zone for both the primary and secondary instances, the deployment is equivalent to single-zone deployment. </div> |
| Specifications | Quantity | The number of ApsaraDB RDS instances that you want to create. Default value: 1. |
| | Instance Name | The name of the instance. <ul style="list-style-type: none"> ◦ The name must be 2 to 64 characters in length ◦ The name must start with a letter. ◦ The name can contain letters, digits, and the following special characters: _ - : ◦ The name cannot start with http:// or https://. |
| | Connection Type | The connection type of the instance. ApsaraDB RDS instances support the following connection types: <ul style="list-style-type: none"> ◦ Internet: ApsaraDB RDS instances of this connection type can be connected over the Internet. ◦ Internal Network: ApsaraDB RDS instances of this connection type can be connected over an internal network. <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e6f2ff;"> ? Note The value of this parameter cannot be changed after the instance is created. Proceed with caution. </div> |
| | Database Engine | The database engine of the instance. Set the value to MySQL . |

| Section | Parameter | Description |
|----------------|-----------------------------|---|
| | Engine Version | The version of the database engine. Valid values: <ul style="list-style-type: none"> ◦ 8.0 ◦ 5.7 ◦ 5.6 |
| | Edition | The edition of the instance. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| | Storage Type | The storage type of the instance. Select local SSD. |
| | Instance Type | The instance type of the instance. Memory size determines the maximum number of connections and the input/output operations per second (IOPS). The actual values are displayed in the console. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| | Storage Capacity | The storage capacity of the instance, which includes the space to store data, system files, binlog files, and transaction files. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| Network | Network Type | The network type of the instance. ApsaraDB RDS instances support the following network types: <ul style="list-style-type: none"> ◦ Classic Network: Cloud services in the classic network are not isolated. Unauthorized access to a cloud service is blocked only by the security group or whitelist policy of the service. <div style="background-color: #e0f2f1; padding: 5px; margin: 5px 0;"> ? Note Instances that use standard SSDs cannot be deployed in the classic network. </div> <ul style="list-style-type: none"> ◦ VPC: A virtual private cloud (VPC) helps you build an isolated network environment on Apsara Stack. You can customize route tables, CIDR blocks, and gateways in a VPC. We recommend that you select VPC for improved security. |
| | VPC | Select a VPC. <div style="background-color: #e0f2f1; padding: 5px; margin: 5px 0;"> ? Note When Network Type is set to VPC, you must specify this parameter. </div> |
| | vSwitch | Select a vSwitch. <div style="background-color: #e0f2f1; padding: 5px; margin: 5px 0;"> ? Note When Network Type is set to VPC, you must specify this parameter. </div> |
| | IP Address Whitelist | The IP addresses that are allowed to connect to the instance. |

4. Click **Submit**.

8.4.2. Create an ApsaraDB RDS for MySQL instance with standard SSDs

Cloud disks are block-level data storage products provided by Alibaba Cloud for Elastic Compute Service (ECS). They provide low latency, high performance, durability, and reliability. This topic describes how to create one or more instances that use standard SSDs in the ApsaraDB RDS console.

Prerequisites

An instance that runs MySQL 5.7 on RDS High-availability Edition can be created.

Context

An ApsaraDB RDS instance with standard SSDs uses a distributed triplicate mechanism to ensure high data reliability. If service disruptions occur within a zone due to hardware faults, data in that zone is copied to an unaffected disk in another zone to ensure data availability.

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, click **Create Instance** in the upper-right corner.
3. Configure the following parameters.

| Section | Parameter | Description |
|----------------|------------------------|--|
| Basic Settings | Organization | The organization to which the instance belongs. |
| | Resource Set | The resource set to which the instance belongs. |
| Region | Region | The region in which you want to create the instance. Services in different regions cannot communicate over an internal network. After you create an instance, its region cannot be changed. |
| | Zone of Primary Node | The zone where the primary instance is deployed. |
| | Deployment Method | Specifies whether to deploy the primary and secondary instances in separate zones. ApsaraDB RDS supports Multi-zone Deployment and Single-zone Deployment . If you select Multi-zone Deployment , you must configure Zone of Secondary Node . |
| | Zone of Secondary Node | The zone where the secondary instance is deployed. This parameter is available only when Deployment Method is set to Multi-zone Deployment . <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> ? Note If you select the same zone for both the primary and secondary instances, the deployment is equivalent to single-zone deployment. </div> |
| | Quantity | The number of ApsaraDB RDS instances that you want to create. Default value: 1. |
| | Instance Name | The name of the instance. <ul style="list-style-type: none"> ◦ The name must be 2 to 64 characters in length. ◦ The name must start with a letter. ◦ The name can contain letters, digits, and the following special characters: <code>_ - :</code> ◦ The name cannot start with <code>http://</code> or <code>https://</code>. |

| Section | Parameter | Description |
|------------------|---|--|
| Specifications | Connection Type | <p>The connection type of the instance. ApsaraDB RDS instances support the following connection types:</p> <ul style="list-style-type: none"> ◦ Internet: ApsaraDB RDS instances of this connection type can be connected over the Internet. ◦ Internal Network: ApsaraDB RDS instances of this connection type can be connected over an internal network. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p> Note The value of this parameter cannot be changed after the instance is created. Proceed with caution.</p> </div> |
| | Database Engine | The database engine of the instance. Select MySQL . |
| | Engine Version | The version of the database engine. Select 5.7 . |
| | Edition | The edition of the instance. Select High-availability . |
| | Storage Type | The storage type of the instance. Select cloud ssd . |
| | Encrypted | Specifies whether to encrypt the standard SSD. This parameter is available only when Storage Type is set to cloud ssd . If you select Encrypted, you must specify the Encryption Key parameter. If you do not have a key, you must first create one in Key Management Service (KMS). For more information, see <i>Create a CMK in Key Management Service User Guide</i> . |
| | Encryption Key | The key that is used to encrypt the standard SSD. This parameter is available only when you select Encrypted . |
| | Instance Type | The instance type of the instance. Memory size determines the maximum number of connections and the input/output operations per second (IOPS). The actual values are displayed in the console. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| Storage Capacity | The storage capacity of the instance, which includes the space to store data, system files, binlog files, and transaction files. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . | |
| Network | Network Type | <p>The network type of the instance. ApsaraDB RDS instances support the following network types:</p> <ul style="list-style-type: none"> ◦ Classic Network: Cloud services in the classic network are not isolated. Unauthorized access to a cloud service is blocked only by the security group or whitelist policy of the service. ◦ VPC: A virtual private cloud (VPC) helps you build an isolated network environment on Apsara Stack. You can customize route tables, CIDR blocks, and gateways in a VPC. We recommend that you select VPC for improved security. |
| | VPC | <p>Select a VPC.</p> <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p> Note When Network Type is set to VPC, you must specify this parameter.</p> </div> |

| Section | Parameter | Description |
|---------|----------------------|--|
| | vSwitch | Select a vSwitch.  Note When Network Type is set to VPC, you must specify this parameter. |
| | IP Address Whitelist | The IP addresses that are allowed to connect to the instance. |

4. Click **Submit**.

8.4.3. View basic information of an instance

This topic describes how to view the details of an ApsaraDB RDS instance, such as its basic information, internal network connection information, status, and configurations.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. Use one of the following methods to go to the **Basic Information** page of an instance:
 - On the **Instances** page, click the ID of an instance to go to the **Basic Information** page.
 - On the **Instances** page, click **Manage** in the **Actions** column corresponding to an instance to go to the **Basic Information** page.

8.4.4. Restart an instance

This topic describes how to manually restart an ApsaraDB RDS for MySQL instance. This applies if the number of connections exceeds a specific threshold or if an instance has performance issues.

Prerequisites

The instance is in the **Running** state.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. Click **Restart Instance** in the upper-right corner.

 **Note** When you restart an instance, applications are disconnected from the instance. We recommend that you make appropriate service arrangements before you restart an instance. Proceed with caution.

5. In the Restart Instance message, click **Confirm**.

8.4.5. Change the specifications of an instance

This topic describes how to change specifications of your instance, such as the instance type and storage capacity, if the specifications do not meet the requirements of your application.

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. Find your ApsaraDB RDS instance and click its ID.
3. In the **Configure Information** section of the **Basic Information** page, click **Change Specifications**.
4. On the **Change Specifications** page, set **Edition**, **Storage Type**, **Instance Type**, and **Storage Capacity**.
5. Click **Submit**.

8.4.6. Set a maintenance window

This topic describes how to set a maintenance window for an ApsaraDB RDS instance.

Context

To ensure the stability of ApsaraDB RDS instances, the backend system performs maintenance of the instances at irregular intervals. The default maintenance window is from 02:00 to 06:00. You can set the maintenance window to the off-peak period of your business to avoid impact on business.

Precautions

- To ensure the stability of the maintenance process, the instance changes to the **Maintaining Instance** state before the maintenance window. When the instance is in this state, access to data in the database and query operations such as performance monitoring are not affected. However, apart from account and database management and IP address whitelist configuration, modification operations such as upgrade, downgrade, and restart are temporarily unavailable.
- During the maintenance window, one or two transient connections may occur. Make sure that you configure automatic reconnection policies for your applications to avoid service disruptions.

Procedure

1. Connect to your ApsaraDB RDS instance. For more information, see [Log on to the ApsaraDB RDS console](#).
2. Click the ID of an instance or click **Manage** in the **Actions** column corresponding to the instance.
3. In the **Configuration Information** section, click **Configure** to the right of **Maintenance Window**.
4. Select a maintenance window and click **Save**.

 **Note** The maintenance window is in UTC+8.

8.4.7. Change the data replication mode

You can set the data replication mode between primary and secondary ApsaraDB RDS instances to improve database availability.

Prerequisites

The ApsaraDB RDS instance uses local SSDs.

Context

- Semi-sync

After an application-initiated update is complete on the primary instance of a cluster, logs are synchronized to all secondary instances. This transaction is considered committed after at least one secondary instance has received the logs, regardless of whether the secondary instance finishes executing the updates specified in the logs.

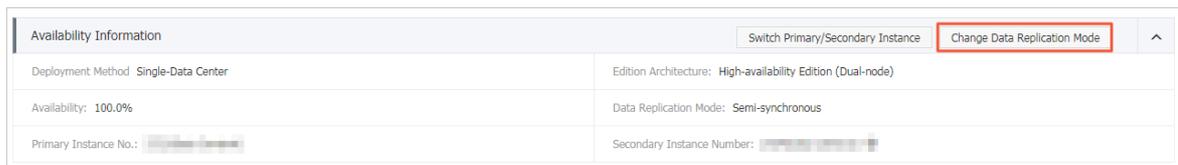
If the secondary instances are unavailable or a network exception occurs between the primary and secondary instances, semi-synchronous replication degrades to the Asynchronous mode.

- Asynchronous

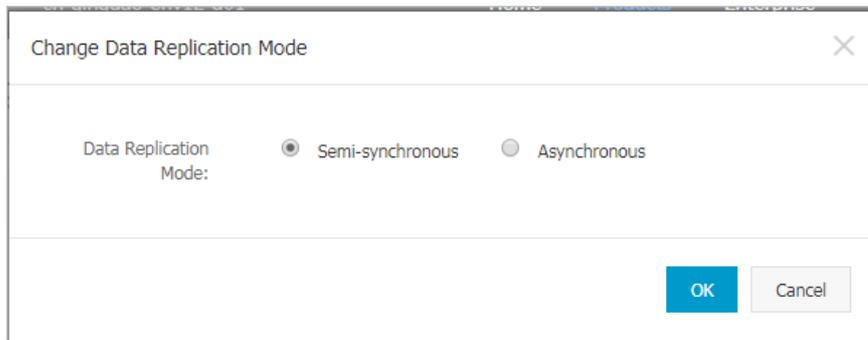
When your application initiates a request to add, delete, or modify data, the primary instance responds to your application immediately after it completes the operation. At the same time, the primary instance starts to asynchronously replicate data to its secondary instances. During asynchronous data replication, operations on the primary instance are not affected if the secondary instances are unavailable. If the primary instance is unavailable, data remains consistent.

Procedure

1. For more information, see [Log on to the ApsaraDB RDS console](#).
2. Find your ApsaraDB RDS instance and click its ID.
3. In the left-side navigation pane, click **Service Availability**.
4. Click **Change Data Replication Mode**.



5. In the Change Data Replication Mode dialog box, select a data replication mode and click **OK**.



8.4.8. Release an instance

This topic describes how to manually release an instance.

Precautions

- Only instances in the running state can be released.
- After an instance is released, the instance data is immediately deleted. We recommend that you back up instance data before you release an instance.

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. In the Actions column corresponding to the instance you want to release, choose **More > Release Instance**.
3. In the **Release Instance** message, click **Confirm**.

8.4.9. Upgrade the minor version of an instance

ApsaraDB RDS for MySQL supports automatic and manual updates of the minor version. These updates increase performance, unveil new features, and fix known issues.

Overview

ApsaraDB RDS for MySQL automatically upgrades the minor version by default. You can log on to the ApsaraDB RDS console, go to the **Basic Information** page of your ApsaraDB RDS instance, and then view the current **Minor Version Upgrade Mode** in the Configuration Information section.

- **Automatic Upgrade:** When a new minor version is released, the system automatically upgrades the minor version of your instance during the specified maintenance window. For more information, see [Set a maintenance window](#).
- **Manual Upgrade:** You can manually upgrade the minor version on the **Basic Information** page. For more information, see [Manually upgrade the minor version](#).

Precautions

- When you upgrade the minor version of your ApsaraDB RDS instance, a 30-second network interruption may occur. We recommend that you upgrade the minor version during off-peak hours or make sure that your applications are configured with automatic reconnection policies.
- The minor version of your ApsaraDB RDS instance cannot be downgraded after it is upgraded.
- After you upgrade the specifications of your ApsaraDB RDS instance, the minor version of the instance is upgraded.

Configure the minor version upgrade mode

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the **Configuration Information** section of the Basic Information page, click **Configure** to the right of **Minor Version Upgrade Mode**.
5. In the Set Minor Version Upgrade Mode dialog box, select **Auto** or **Manual**, and click **OK**.

Manually upgrade the minor version

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the **Configuration Information** section of the page, click **Upgrade Minor Version**.

 **Note** The **Upgrade Minor Version** button is displayed only when a new minor version is available.

5. In the dialog box that appears, specify the upgrade time and click **OK**.

FAQ

- **Q:** After I upgraded the minor version of my ApsaraDB RDS instance, why does the `SELECT @@version` statement still return the source minor version that I used before the upgrade?
A: The `SELECT @@version` statement returns the minor version of Alibaba Cloud, not the minor version of the ApsaraDB RDS for MySQL instance. You need to execute the `show variables like '%rds_release_date%'` statement to view the minor version of your instance.
- **Q:** When an upgrade takes effect, is my instance upgraded only to the next minor version?
A: No, when an upgrade takes effect, your instance is upgraded to the latest minor version.

8.4.10. Modify parameters of an instance

This topic describes how to view and modify the values of some parameters and query parameter modification records in the console.

Precautions

- To ensure instance stability, you can select specific parameters to modify in the ApsaraDB RDS console.
- When you modify parameters on the **Editable Parameters** tab, refer to the **Value Range** column corresponding to each parameter.
- After some parameters are modified, you must restart your ApsaraDB RDS instance for the changes to take effect. For more information, see the **Force Restart** column on the **Editable Parameters** tab. We recommend that you modify the parameters of an instance during off-peak hours and make sure that your applications are configured with automatic reconnection policies.

Modify parameters

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Parameters**.
5. You can perform the following operations:

Export the parameter settings of the ApsaraDB RDS instance to your computer.

On the **Editable Parameters** tab, click **Export Parameters**. The parameter settings of the ApsaraDB RDS instance are exported as a TXT file to your computer.

Modify and import the parameter settings.

- i. After you have modified parameters in the exported parameter file, click **Import Parameters** and copy the parameter settings to the field.
- ii. Click **OK**.
- iii. In the upper-right corner of the page, click **Apply Changes**.

Note

- If the new parameter value takes effect only after an instance restarts, the system prompts you to restart the ApsaraDB RDS instance. We recommend that you restart the ApsaraDB RDS instance during off-peak hours and make sure that your applications are configured with automatic reconnection policies.
- Before the new parameter values are applied, you can click **Cancel Changes** to cancel them.

Modify a single parameter.

- i. On the **Editable Parameters** tab, find the parameter that you want to reconfigure, and click the  icon in the **Actual Value** column.
- ii. Enter a new value based on the prompted value range.
- iii. Click **Confirm**.

- iv. In the upper-right corner of the page, click **Apply Changes**.

 **Note**

- If the new parameter value takes effect only after an instance restarts, the system prompts you to restart the ApsaraDB RDS instance. We recommend that you restart the ApsaraDB RDS instance during off-peak hours and make sure that your applications are configured with automatic reconnection policies.
- Before the new parameter value is applied, you can click **Cancel Changes** to cancel it.

View the parameter modification history

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Parameters**.
5. On the Parameters page, click the **Edit History** tab.
6. Select a time range and then click **Search**.

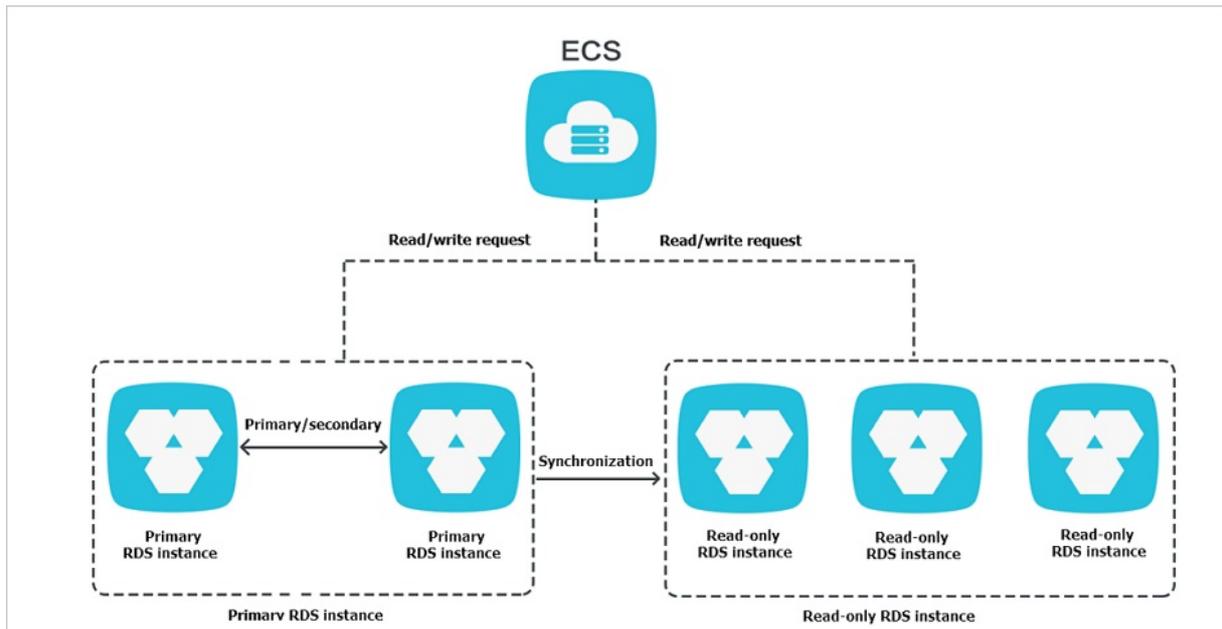
8.4.11. Read-only instances

8.4.11.1. Overview of read-only instances

ApsaraDB RDS for MySQL allows you to create read-only instances. In scenarios where an instance has a small number of write requests but a large number of read requests, you can create read-only instances to distribute database access loads away from the primary instance. This topic describes the features and limits of read-only instances.

To scale up the reading capability and distribute database access loads, you can create one or more read-only instances in a region. Read-only instances can increase the application throughput when a large amount of data is being read.

A read-only instance with a single physical node and no backup node uses the native replication capability of MySQL to synchronize changes from the primary instance to all its read-only instances. Read-only instances must be in the same region as the primary instance but do not have to be in the same zone as the primary instance. The following figure shows the topology of read-only instances.



Read-only instances have the following features:

- Specifications of a read-only instance can be different from those of the primary instance and can be changed at any time. This facilitates elastic scaling.
- Read-only instances do not require account or database maintenance. Account and database information is synchronized from the primary instance.
- The whitelists of read-only instances can be independently configured.
- System performance monitoring is provided.

ApsaraDB RDS provides up to 20 system performance monitoring views, including those for disk capacity, IOPS, connections, CPU utilization, and network traffic. You can view the load of instances.

- ApsaraDB RDS provides a variety of optimization recommendations, such as storage engine check, primary key check, large table check, and check for excessive indexes and missing indexes. You can optimize your databases based on the optimization recommendations and specific applications.

8.4.11.2. Create a read-only instance

You can create read-only instances of different specifications based on your business requirements.

Precautions

- A maximum of five read-only instances can be created for a primary instance.
- Backup settings and temporary backup are not supported.
- Instance restoration is not supported.
- Data migration to read-only instances is not supported.
- Database creation and deletion are not supported.
- Account creation, deletion, authorization, and password changes are not supported.
- After a read-only instance is created, you cannot restore data by directly overwriting the primary instance with a backup set.

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic**

Information page.

- In the **Distributed by Instance Role** section on the right side of the **Basic Information** page, click **Create Read-only Instance**.
- On the **Create Read-only Instance** page, configure the following parameters.

| Section | Parameter | Description |
|-----------------------|-------------------------|--|
| Region | Region | The region in which you want to create the read-only instance. |
| Specifications | Database Engine | The database engine of the read-only instance, which is the same as that of the primary instance and cannot be changed. |
| | Engine Version | The version of the database engine, which is the same as that of the primary instance and cannot be changed. |
| | Edition | Set the value to Read-only . |
| | Instance Type | The instance type of the read-only instance. The instance type of the read-only instance can be different from that of the primary instance, and can be changed at any time to facilitate flexible upgrade and downgrade. |
| | Storage Capacity | The storage capacity of the read-only instance. To ensure sufficient I/O throughput for data synchronization, we recommend that you select at least the same instance type and storage capacity as the primary instance for the read-only instance. Valid values: 20 to 6000. Unit: GB. The value is in 1 GB increments. |
| Network Type | Network Type | The network type of the read-only instance, which is the same as that of the primary instance and cannot be changed. |
| | VPC | The VPC in which you want to create the read-only instance. |
| | vSwitch | The vSwitch in the VPC. |

- After you configure the preceding parameters, click **Submit**.

8.4.11.3. View details of read-only instances

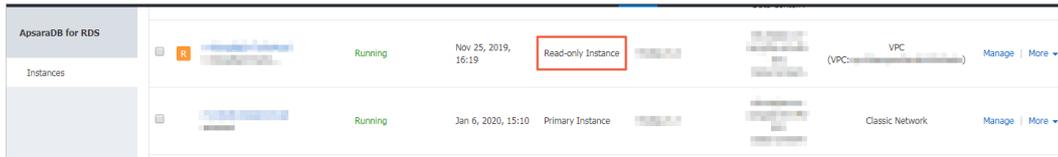
This topic describes how to view details of read-only instances. You can go to the Basic Information page of a read-only instance from the Instances page or from the read-only instance list of the primary instance. Read-only instances are managed in the same way as primary instances. The read-only instance management page shows the management operations that can be performed.

View details of a read-only instance from the Instances page

- [Log on to the ApsaraDB for RDS console](#).
- On the **Instances** page, click the ID of a read-only instance. The **Basic Information** page appears.

In the instance list, Instance Role of read-only instances is displayed as Read-only Instance, as shown in [View a read-only instance](#).

View a read-only instance



View details of a read-only instance from the Basic Information page of the primary instance

1. Log on to the [ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. On the **Basic Information** page, move the pointer over the number below **Read-only Instance** in the **Distributed by Instance Role** section. The ID of the read-only instance is displayed.
5. Click the ID of the read-only instance to go to the read-only instance management page.

8.5. Accounts

8.5.1. Create an account

After you create an ApsaraDB RDS instance and configure its IP address whitelist, you must create a database and an account on the instance. This topic describes how to create privileged and standard accounts.

Context

ApsaraDB RDS for MySQL supports two types of database accounts: privileged and standard. You can manage all your accounts and databases in the ApsaraDB RDS console. For more information about permissions that can be granted to each type of account, see [Account permissions](#).

| Account type | Description |
|---------------------------|---|
| Privileged account | <ul style="list-style-type: none"> You can create and manage privileged accounts by using the ApsaraDB RDS console or API operations. You can create only one privileged account on each ApsaraDB RDS instance. The privileged account can be used to manage all standard accounts and databases on the instance. A privileged account allows you to manage permissions to a fine-grained level. For example, you can grant each standard account the permissions to query specific tables. A privileged account has the permissions to disconnect all standard accounts on the instance. |
| Standard account | <ul style="list-style-type: none"> You can create and manage standard accounts by using the ApsaraDB RDS console, API operations, or SQL statements. You can create up to 500 standard accounts on an instance. You must manually grant standard accounts the specific database permissions. You cannot use a standard account to create, manage, or disconnect other accounts from databases. |

| Account type | Maximum number of databases | Maximum number of tables | Maximum number of accounts |
|--------------------|-----------------------------|--------------------------|--|
| Privileged account | Unlimited | < 200,000 | Varies based on the kernel parameter settings of the instance. |

| Account type | Maximum number of databases | Maximum number of tables | Maximum number of accounts |
|------------------|-----------------------------|--------------------------|--|
| Standard account | 500 | < 200,000 | Varies based on the kernel parameter settings of the instance. |

Create a privileged account

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Accounts**.
5. On the **Accounts** page, click the **Accounts** tab.
6. Click **Create Account**.



7. On the **Create Account** page, configure the following parameters.

| Parameter | Description |
|--------------------------|--|
| Database Account | Enter the name of the account. The account name must meet the following requirements: <ul style="list-style-type: none"> ◦ The name is 1 to 16 characters in length. ◦ The name starts with a lowercase letter and ends with a lowercase letter or digit. ◦ The name contains lowercase letters, digits, and underscores (_). |
| Account Type | Select Privileged Account. |
| Password | Enter the password of the account. The password must meet the following requirements: <ul style="list-style-type: none"> ◦ The password is 8 to 32 characters in length. ◦ The password contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. ◦ Special characters include ! @ # \$ % ^ & * () _ + - = |
| Re-enter Password | Enter the password of the account again. |
| Description | Optional. Enter information about the account to facilitate subsequent management. The description can be up to 256 characters in length. |

8. Click **Create**.

Reset the permissions of a privileged account

If an issue occurs on the privileged account, you can enter the password of the privileged account to reset permissions. For example, you can reset the permissions if the permissions are unexpectedly revoked.

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.

3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Accounts**.
5. On the **Accounts** page, click the **Accounts** tab.
6. On the **Accounts** tab, find the privileged account and click **Reset Permissions** in the **Actions** column.
7. On the **Initialize Account** page, enter the password of the privileged account and click **OK**.

Create a standard account

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Accounts**.
5. On the **Accounts** page, click the **Accounts** tab.
6. Click **Create Account**.



7. On the **Create Account** page, configure the following parameters.

| Parameter | Description |
|-----------------------------|--|
| Database Account | Enter the name of the account. The account name must meet the following requirements: <ul style="list-style-type: none"> ◦ The name is 1 to 16 characters in length. ◦ The name starts with a lowercase letter and ends with a lowercase letter or digit. ◦ The name contains lowercase letters, digits, and underscores (_). |
| Account Type | Select Standard Account. |
| Authorized Databases | Select one or more databases on which you want to grant permissions to the account. You can also leave this parameter empty at this time and authorize databases after the account is created. <ol style="list-style-type: none"> i. Select one or more databases from the Unauthorized Databases section and click Add to add them to the Authorized Databases section. ii. In the Authorized Databases section, select the Read/Write, Read-only, DDL Only, or DML Only permissions on each authorized database. If you want to grant the same permissions on multiple databases to the account, click the button in the upper-right corner of the section. The button may appear as Set All to Read/Write . |
| Password | Enter the password of the account. The password must meet the following requirements: <ul style="list-style-type: none"> ◦ The password is 8 to 32 characters in length. ◦ The password contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. ◦ Special characters include ! @ # \$ % ^ & * () _ + - = |

| Parameter | Description |
|--------------------------|---|
| Re-enter Password | Enter the password of the account again. |
| Description | Optional. Enter information about the account to facilitate subsequent management. The description can be up to 256 characters in length. |

8. Click **Create**.

Account permissions

| Account type | Authorization type | Permission | | | | |
|--------------------|--------------------|--------------------|-------------------------|--------------------|----------------|-------------------------|
| Privileged account | - | SELECT | INSERT | UPDATE | DELETE | CREATE |
| | | DROP | RELOAD | PROCESS | REFERENCES | INDEX |
| | | ALTER | CREATE TEMPORARY TABLES | LOCK TABLES | EXECUTE | REPLICATION SLAVE |
| | | REPLICATION CLIENT | CREATE VIEW | SHOW VIEW | CREATE ROUTINE | ALTER ROUTINE |
| | | CREATE USER | EVENT | TRIGGER | - | - |
| Standard account | Read-only | SELECT | LOCK TABLES | SHOW VIEW | PROCESS | REPLICATION SLAVE |
| | | REPLICATION CLIENT | - | - | - | - |
| | Read/write | SELECT | INSERT | UPDATE | DELETE | CREATE |
| | | DROP | REFERENCES | INDEX | ALTER | CREATE TEMPORARY TABLES |
| | | LOCK TABLES | EXECUTE | CREATE VIEW | SHOW VIEW | CREATE ROUTINE |
| | | ALTER ROUTINE | EVENT | TRIGGER | PROCESS | REPLICATION SLAVE |
| | | REPLICATION CLIENT | - | - | - | - |
| | DDL-only | CREATE | DROP | INDEX | ALTER | CREATE TEMPORARY TABLES |
| | | LOCK TABLES | CREATE VIEW | SHOW VIEW | CREATE ROUTINE | ALTER ROUTINE |
| | | PROCESS | REPLICATION SLAVE | REPLICATION CLIENT | - | - |

| Account type | Authorization type | Permission | | | | |
|--------------|--------------------|-------------|-------------------|--------------------|--------|-------------------------|
| DML-only | | SELECT | INSERT | UPDATE | DELETE | CREATE TEMPORARY TABLES |
| | | LOCK TABLES | EXECUTE | SHOW VIEW | EVENT | TRIGGER |
| | | PROCESS | REPLICATION SLAVE | REPLICATION CLIENT | - | - |

8.5.2. Reset the password

You can use the ApsaraDB RDS console to reset the password of your database account.

Prerequisites

The instance is in the **Running** state.

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Accounts**.
5. Find an account and click **Reset Password** in the **Actions** column.
6. In the dialog box that appears, enter and confirm the new password, and then click **OK**.

-  **Note** The password must meet the following requirements:
- The password must be 8 to 32 characters in length.
 - The password must contain at least three of the following characters: uppercase letters, lowercase letters, digits, and special characters.
 - Special characters include ! @ # \$ % ^ & * () _ + - =

8.5.3. Modify account permissions

You can modify the account permissions of your ApsaraDB RDS instance at any time.

Prerequisites

You can modify the permissions of a standard account. The permissions of privileged accounts can only be reset to the default settings and cannot be changed to a specific set of permissions.

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic**

Information page.

4. In the left-side navigation pane, click **Accounts**.
5. Find an account and click **Edit Permissions** in the Actions column.
6. Configure the following parameters

| Parameter | Description |
|-----------------------------|--|
| Authorized Databases | In the Unauthorized Databases section, select a database and click Add to authorize the database. In the Authorized Databases section, select a database and click Remove to remove the permissions from the database. |
| Authorized Databases | You can set permissions on each database in the Authorized Database section. You can also click the button such as Set All to Read/Write in the upper-right corner to set the permissions of the account on all authorized databases. <ul style="list-style-type: none">◦ Read-only: grants the database read-only permissions to the account.◦ Read/Write: grants the database read/write permissions to the account.◦ DDL Only: grants the database permissions of DDL operations to the account.◦ DML Only: grants the database permissions of DML operations to the account. |

7. Click **OK**.

8.5.4. Delete an account

You can delete a database account in the ApsaraDB RDS console.

Prerequisites

You can use the console to delete privileged and standard accounts that are no longer used.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Accounts**.
5. Find the account you want to delete and click **Delete** in the Actions column.
6. In the message that appears, click **Confirm**.

 **Note** Accounts in the **Processing** state cannot be deleted.

8.6. Databases

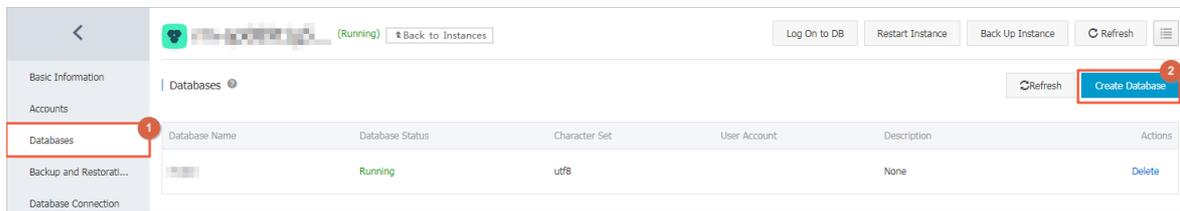
8.6.1. Create a database

After you create an ApsaraDB RDS instance and configure its whitelist, you must create a database and an account in the instance.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.

3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Databases**.
5. Click **Create Database**.



6. Configure the following parameters.

| Parameter | Description |
|---------------------------------|--|
| Database Name | <ul style="list-style-type: none"> ◦ The name must be 1 to 64 characters in length. ◦ The name must start with a letter and end with a letter or digit. ◦ The name can contain lowercase letters, digits, underscores (_), and hyphens (-). ◦ The name must be unique within the instance. |
| Supported Character Sets | Select utf8, gbk, latin1, utf8mb4, or all. If you want to use other character sets, select all , and then select the required character set from the list. |
| Description | Optional. Enter information about the database to facilitate subsequent management. The description can be up to 256 characters in length. |

7. Click **Create**.

8.6.2. Delete a database

You can delete databases that are no longer used in the ApsaraDB RDS console.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Databases**.
5. Find the database you want to delete and click **Delete** in the **Actions** column.
6. In the Delete Database message, click **Confirm**.

8.7. Database connection

8.7.1. Change the endpoint and port number of an instance

This topic describes how to view and change the endpoint and port number of an instance.

View the endpoint and port number

1. Log on to the [ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Connection**.
You can view **Internal Endpoint** and **Internal Port** of the instance on the **Database Connection** page. If you apply for a public endpoint, you can also view **Public Endpoint** and **Public Port**.

Change the endpoint and port number

1. Log on to the [ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Connection**.
5. Click **Change Endpoint**.
6. In the **Change Endpoint** dialog box, select a connection type from the **Connection Type** drop-down list.
7. Set **Endpoint** and **Port**, and click **OK**.

Change Endpoint

Connection Type: Internal Endpoint

Endpoint: [redacted].mysql.rds.intra.env17e.shuguang.com
The endpoint must be 8 to 64 characters and can contain letters, digits, and hyphen (-). It must start with a lowercase letter.

Port: 3306
Port Range: 1000 to 65534

OK Cancel

Note

- The prefix of an endpoint must be 8 to 64 characters in length and can contain only letters, digits, and hyphens (-). It must start with a lowercase letter.
- The port number must be a value within the range of 1000 to 65534.

8.7.2. Log on to an ApsaraDB RDS instance by using DMS

This topic describes how to log on to an ApsaraDB RDS instance by using Data Management (DMS).

Prerequisites

The IP address whitelist is configured. For more information about how to configure an IP address whitelist, see

Configure a whitelist.

Procedure

1. Log on to the ApsaraDB for RDS console.
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the upper-right corner of the page, click **Log On to DB** to go to the Database Logon page.
5. In the **Login instance** dialog box of the **DMS** console, check the **database type**, **instance area**, and **connection string address**. If the preceding information is correct, enter the **database account** and **database password**, as shown in the following figure.

| Parameter | Description |
|----------------------------------|---|
| Database type | The database engine of the instance. By default, this parameter is set to the database engine of the instance that you want to access. |
| Instance Area | The region where the instance resides. By default, this parameter is set to the region where the current instance resides. |
| Connection string address | The endpoint and port number that are used to connect to the instance. By default, this parameter is set to the endpoint and port number of the current instance. |
| Database account | The account that is used to connect to the database. |
| Database Password | The password of the account that is used to connect to the database. |

6. Click **Login**.

Note If you want the browser to remember the password, select **Remember password** before you click **Login**.

8.7.3. Hybrid access from both the classic network and VPCs

This topic describes how to use the hybrid access solution of ApsaraDB RDS to change the network type of an instance from classic network to Virtual Private Network (VPC) without network interruptions.

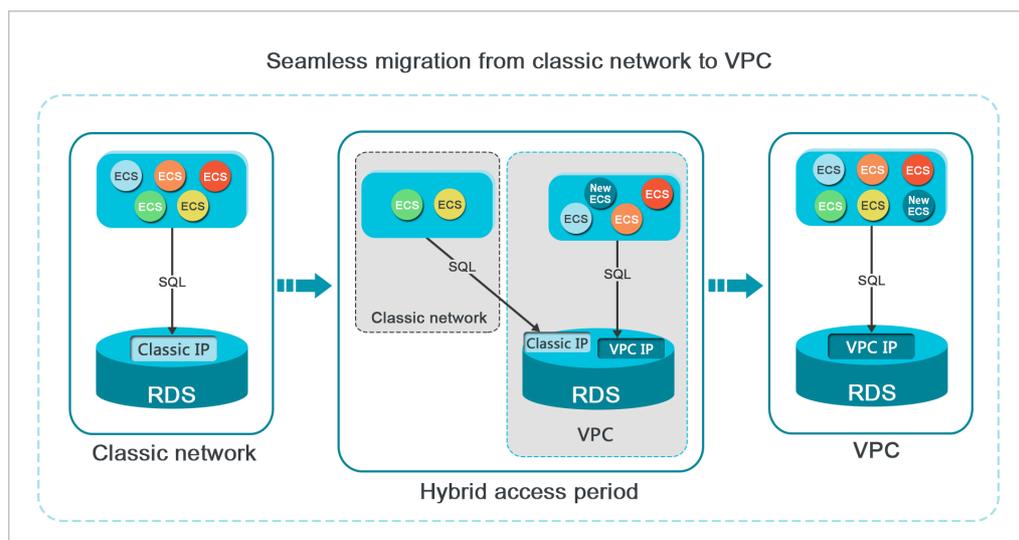
Background

In the past, when you change the network type of an ApsaraDB RDS instance from classic network to VPC, the internal endpoint of the instance would remain the same but the IP address bound to the endpoint would change to the corresponding IP address in the VPC. This change would cause a 30-second network interruption, and ECS instances within the classic network would not be able to access the ApsaraDB RDS instance by using the internal endpoint within this period. To smoothly change the network type, ApsaraDB RDS provides the hybrid access solution.

Hybrid access refers to the ability of an ApsaraDB RDS instance to be accessed by ECS instances in both the classic network and VPCs. During the hybrid access period, the ApsaraDB RDS instance reserves the original internal endpoint of the classic network and adds the internal endpoint of VPCs. This prevents network interruptions during the network type switchover.

For security and performance purposes, we recommend that you use only the internal VPC endpoint. Therefore, ApsaraDB RDS allows the configured hybrid access solution to remain valid only for a specific period. When the hybrid access period elapses, ApsaraDB RDS releases the internal classic network endpoint. In this case, your applications cannot connect to your ApsaraDB RDS instance by using the internal classic network endpoint. You must add the internal VPC endpoint to all your applications during the hybrid access period. This ensures a smooth network migration and avoids interruptions to your workloads.

For example, your company wants to use the hybrid access solution to change the network type from classic network to VPC. During the hybrid access period, some applications can access the database by using the internal endpoint of VPCs, and the other applications can access the database by using the original internal endpoint of the classic network. When all the applications access the database by using the internal endpoint of VPCs, the internal endpoint of the classic network can be released. The following figure illustrates the scenario.



Limits

During the hybrid access period, the instance has the following limits:

- The network type of the instance cannot be changed to classic network.
- The instance cannot be migrated to another zone.

Prerequisites

- The network type of the instance is classic network.
- Available VPCs and vSwitches exist in the zone where the instance resides.

Change the network type from classic network to VPC

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Connection**.
5. On the Instance Connection tab, click **Switch to VPC**.
6. In the Switch to VPC dialog box, select a VPC and a vSwitch and specify whether to retain the endpoint used in the classic network.

Determine whether to select **Reserve Original Classic Endpoint** based on the details described in the following table.

| Operation | Description |
|---|--|
| Clear the Reserve Original Classic Network Endpoint option | <p>The endpoint used in the classic network is replaced with an endpoint in the VPC.</p> <p>When you change the network type, a network interruption of about 30 seconds occurs, and the connection between ECS instances in the classic network and the ApsaraDB RDS instance are interrupted.</p> |
| Select the Reserve Original Classic Network Endpoint option | <p>The endpoint used in the classic network is retained, and a new endpoint to be used in the VPC is generated. In such cases, the ApsaraDB RDS instance runs in hybrid access mode. ECS instances in both the classic network and a VPC can connect to the ApsaraDB RDS instance over the internal network.</p> <p>When you change the network type, no network interruption occur. Connections between ECS instances in the classic network and the ApsaraDB RDS instance are available till the endpoint used in the classic network expires.</p> <p>Specify the expiration time of the classic network endpoint. You must add the new VPC endpoint to the ECS instance before the endpoint in the classic network expires. This ensures smooth network switchover.</p> |

7. Add the internal IP addresses of ECS instances in the selected VPC to a VPC whitelist. This allows the ECS instances to access the ApsaraDB RDS instance over the internal network. If no VPC whitelists are available, create a whitelist. For more information, see [Configure a whitelist](#).

Change the expiration time for the original internal endpoint of the classic network

During the period in which your instance can be connected over the classic network or VPCs, you can specify the expiration time for the endpoint of the classic network. The setting takes effect immediately. For example, if the endpoint of the classic network expires on August 18, 2017 and you change the expiration time to 14 days later on August 15, 2017, the endpoint of the classic network is released on August 29, 2017.

To change the expiration time, perform the following operations:

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Connection**.
5. Click **Change Expiration Time**.
6. In the **Change Expiration Time** dialog box, select an expiration time and click **OK**.

8.7.4. Change the network type of an instance

This topic describes how to change the network type of an ApsaraDB RDS instance between classic network and VPC.

Context

- **Classic network:** ApsaraDB RDS instances in the classic network are not isolated. Unauthorized access to these instances can be blocked only by whitelists.
- **VPC:** Each VPC is an isolated network. We recommend that you select the VPC network type because it is more secure than the classic network.

You can configure route tables, CIDR blocks, and gateways in a VPC. To smoothly migrate applications to the cloud, you can use leased lines or VPNs to connect on-premises data center to a VPC to create a virtual data center.

Change the network type from VPC to classic network

Precautions

- After you change the network type from VPC to classic network, the internal endpoint of your ApsaraDB RDS instance remains unchanged. However, the IP address that is associated with the internal endpoint changes.
- After you change the network type from VPC to classic network, ECS instances in the same VPC as the ApsaraDB RDS instance can no longer connect to the ApsaraDB RDS instance by using the internal endpoint. You must update the endpoint for the applications deployed on the ECS instances.
- When you change the network type, a 30-second network interruption may occur. To avoid business interruption, change the network type during off-peak hours or make sure that your applications are configured with automatic reconnection policies.

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Connection**.
5. In the upper-right corner of the Database Connection section, click **Switch to Classic Network**.
6. In the message that appears, click **OK**.

Change the network type from classic network to VPC

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Connection**.
5. In the upper-right corner of the Database Connection section, click **Switch to VPC**.
6. In the **Switch to VPC** dialog box, select a VPC and vSwitch and specify whether to **Reserve Original Classic**

Network Endpoint. Click OK. For more information about **Reserve Original Classic Network Endpoint**, see [Hybrid access from both the classic network and VPCs](#).

8.7.5. Switch an ApsaraDB RDS for MySQL instance to a new VPC or vSwitch

This topic describes how to switch an ApsaraDB RDS for MySQL instance to a new VPC or vSwitch.

Prerequisites

The instance is deployed in a VPC.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Connection**.
5. In the upper-right corner of the Database Connection section, click **Switch vSwitch**.
6. Select a VPC and a vSwitch, and then click **OK**.
7. In the message that appears, click **Switch**.

Note

- A 30-second network interruption occurs when you switch the VPC and vSwitch of an ApsaraDB RDS for MySQL instance. Make sure that your application is configured to automatically reconnect to the ApsaraDB RDS for MySQL instance.
- We recommend that you clear the cache immediately after the instance is switched to a new VPC and vSwitch. Otherwise, data can only be read but cannot be written.

8.8. Database proxy

8.8.1. Dedicated proxy

This topic describes the dedicated proxy feature provided by ApsaraDB RDS for MySQL. The dedicated proxy feature provides advanced features such as read/write splitting, connection pooling, and transaction splitting.

Prerequisites

Your ApsaraDB RDS instance runs one of the following MySQL versions and RDS editions:

- MySQL 5.6 on RDS High-availability Edition (with local SSDs)
- MySQL 5.7 on RDS Enterprise Edition
- MySQL 5.7 on RDS High-availability Edition (with local SSDs)
- MySQL 8.0 on RDS High-availability Edition (with local SSDs)

Context

The dedicated proxy feature uses dedicated computing resources. This feature has the following benefits:

- A unified proxy endpoint is provided to connect to all the dedicated proxies that are enabled on your ApsaraDB RDS instance. This reduces maintenance costs by relieving the need to update the endpoints on your application. The proxy endpoint remains valid unless you release the dedicated proxies. For example, you may enable

read/write splitting during peak hours, and then release read-only instances and disable read/write splitting after peak hours. In these cases, you do not need to update the endpoints on your application because the proxy endpoint remains connected.

- Dedicated proxies exclusively serve your ApsaraDB RDS instance and its read-only instances. You do not need to compete with other users for resources. This ensures service stability.
- Dedicated proxies are scalable. You can add dedicated proxies based on your business requirements to handle more workloads.

Limits

- Dedicated proxies do not support Secure Sockets Layer (SSL) encryption.
- Dedicated proxies do not support compression protocols.

Precautions

- When you change the specifications of your ApsaraDB RDS instance or its read-only instances, a network interruption may occur.
- If you connect your application to the proxy endpoint, all the requests that are encapsulated in transactions are routed to your ApsaraDB RDS instance. This applies if you do not enable the transaction splitting feature.
- If you use the proxy endpoint to implement read/write splitting, the read consistency of requests that are not encapsulated in transactions cannot be guaranteed. If you require read consistency, you must encapsulate these requests in transactions.
- If you connect your application to the proxy endpoint, the `SHOW PROCESSLIST` statement returns a combination of results from the primary ApsaraDB RDS instance and all of its read-only instances.
- If you execute `multi-statements` or stored procedures, the read/write splitting feature is disabled and all the subsequent requests over the current connection are routed to the primary ApsaraDB RDS instance. To enable the read/write splitting feature again, you must close the current connection and establish a new connection.
- The dedicated proxy feature supports the `/*FORCE_MASTER*/` and `/*FORCE_SLAVE*/` hints. However, requests that contain hints have the highest route priorities and are not constrained by consistency or transaction limits. Before you use these hints, you must check whether these hints are suitable for your workloads. In addition, these hints cannot contain statements such as `/*FORCE_SLAVE*/ set names utf8;`. Such statements can change environment variables. If you include such statements in these hints, errors may occur when you process your subsequent workloads.
- After you enable the dedicated proxy feature, each connection is replicated to the primary ApsaraDB RDS instance and all of its read-only instances in compliance with the 1:N connection model. We recommend that you specify the same connection specifications for these instances. If these instances have different connection specifications, the number of connections allowed varies based on the lowest connection specifications among these instances.
- If you create or restart a read-only instance after you enable the dedicated proxy feature, only the requests sent over new connections are routed to the new or restarted read-only instance.
- The `max_prepared_stmt_count` parameter must be set to the same value for the primary ApsaraDB RDS instance and all of its read-only instances.

Enable the dedicated proxy feature

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Proxy**.
5. Click **Enable now**.

Overview of the Proxy Service tab

After the dedicated proxy feature is enabled, you can use the generated proxy endpoint to implement features such as read/write splitting, connection pooling, and transaction splitting.

| Section | Parameter | Description |
|----------------|-------------------------------------|--|
| Proxy Endpoint | Instance ID | The ID of the ApsaraDB RDS instance. |
| | Enabled Proxies | The number of enabled dedicated proxies. You can process more requests by enabling more dedicated proxies. After public preview ends, you must pay for the proxies that you enable. |
| | Read/Write Splitting | Specifies whether to enable the read/write splitting feature for the proxy endpoint. For more information, see Read/write splitting . |
| | Short-Lived Connection Optimization | The type of connection pool for the proxy endpoint. This feature is suitable for scenarios where PHP short-lived connections are established. For more information, see Short-lived connection optimization . Note You can click Enable or Disable to the right of Short-Lived Connection Optimization to enable or disable this feature. |
| | Transaction Splitting | Specifies whether to enable the transaction splitting feature for the proxy endpoint. For more information, see Transaction splitting . Note You can click Enable or Disable to the right of Transaction Splitting to enable or disable this feature. |
| | Endpoint | The proxy endpoint that is generated after the dedicated proxy feature is enabled. This endpoint connects to all the dedicated proxies that are enabled on the ApsaraDB RDS instance. The read/write splitting feature is also bound to this endpoint. Note You can click Copy Address to the right of Endpoint to copy the endpoint. |
| | Port | The port that is used to connect to the proxy endpoint. |
| | Endpoint Type | The network type of the proxy endpoint. |
| | Proxy Type | The type of proxy that is enabled. Only Dedicated Proxy is supported. |

| Section | Parameter | Description |
|---------|-----------------|---|
| Proxy | CPU and Memory | The CPU and memory of the dedicated proxies. Only 2 Cores, 4 GB is supported. |
| | Enabled Proxies | <p>The number of dedicated proxies that are enabled on the ApsaraDB RDS instance. Up to 60 dedicated proxies are supported.</p> <p>Note We recommend that you specify the number of dedicated proxies as the total number of CPU cores of your ApsaraDB RDS instance and its read-only instances divided by 8 and rounded up to the nearest integer.</p> <p>For example, if your ApsaraDB RDS instance has 8 CPU cores and its read-only instances have 4 CPU cores, the recommended number of dedicated proxies is 2, as calculated in the following formula: $\lceil (8 + 4) / 8 \rceil = 2$.</p> |

Adjust the number of dedicated proxies

Note When you adjust the number of dedicated proxies, your network connection is interrupted. Make sure that your applications are configured with automatic reconnection policies.

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Proxy**.
5. In the **Proxy** section of the Proxy Service tab, change the number in the **Adjusted Proxies** column and click **Apply** in the **Adjustment Plan** column.
6. In the Configure Proxy Resources dialog box, select **Migrate Immediately** to apply the change. You can also select **Next Maintenance Period** to set a maintenance window for the change to take effect. Click **OK**.

View the monitoring data of dedicated proxies

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Proxy**.
5. Click the **Monitoring Data** tab.
6. Select a time range to view the **CPU Utilization** metric within that time range.

Disable the dedicated proxy feature

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Proxy**.
5. In the upper-right corner of the page, click **Disable Proxy Service**.

6. In the message that appears, click **OK**.

8.8.2. Short-lived connection optimization

This topic describes the short-lived connection optimization feature provided by ApsaraDB RDS for MySQL in its dedicated proxy feature.

Prerequisites

- Your ApsaraDB RDS instance runs one of the following MySQL versions and RDS editions:
 - MySQL 5.6 on RDS High-availability Edition (with local SSDs)
 - MySQL 5.7 on RDS Enterprise Edition
 - MySQL 5.7 on RDS High-availability Edition (with local SSDs)
 - MySQL 8.0 on RDS High-availability Edition (with local SSDs)
- The database proxy feature is enabled for the instance. For more information, see [Dedicated proxy](#).

Context

The short-lived connection optimization feature is used to reduce workloads on the ApsaraDB RDS instance caused by frequent short-lived connections. When a client is disconnected, the system checks whether the closed connection is idle. If the connection is considered idle, the dedicated proxy retains the connection in the connection pool for a short period of time. When the client initiates a request for access to your instance again, the dedicated proxy searches the connection pool for an idle connection that matches the request. The connection pool is matched based on the values of the user, clientip, and dbname fields in the request. If the dedicated proxy finds an idle connection that matches the request, it reuses the matched idle connection. If no idle connection can be matched, a new connection is established with your instance to reduce database connection overheads.

 **Note** The short-lived connection optimization feature does not reduce concurrent connections with the instance. It decreases the frequency in which connections are established between an application and your instance to reduce overheads of the primary MySQL thread and improve efficiency to process business requests. However, idle connections in the connection pool still occupy the database threads for a short period of time.

Precautions

You cannot configure different permissions for different source IP addresses by using the same account. Otherwise, errors may occur when connections in the connection pool are reused. For example, if a user account has permissions on database_a when its source IP address is 192.168.1.1 but does not have permissions on database_a when its source IP address is 192.168.1.2, the short-lived connection optimization feature may encounter permission errors.

Enable short-lived connection optimization

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Proxy**.
5. Click the **Proxy Service** tab.
6. In the **Proxy Endpoint** section, click **Enable** to the right of **Short-Lived Connection Optimization**.

8.8.3. Transaction splitting

This topic describes the transaction splitting feature provided by the database proxy of ApsaraDB RDS. This feature identifies and distributes read requests initiated before write requests within a transaction to read-only instances. This reduces workloads on the primary instance.

Prerequisites

- Your ApsaraDB RDS instance runs one of the following MySQL versions and RDS editions:
 - MySQL 5.6 on RDS High-availability Edition (with local SSDs)
 - MySQL 5.7 on RDS Enterprise Edition
 - MySQL 5.7 on RDS High-availability Edition (with local SSDs)
 - MySQL 8.0 on RDS High-availability Edition (with local SSDs)
- The database proxy feature is enabled for the instance. For more information, see [Dedicated proxy](#).

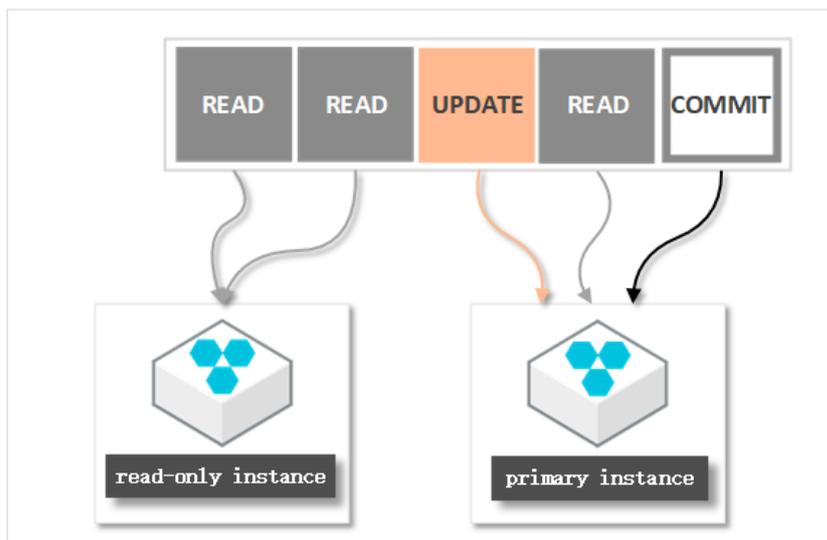
Context

By default, the dedicated proxy sends all requests in transactions to the primary instance to ensure the correctness of the transactions. If the framework encapsulates all requests in transactions, the primary instance becomes heavily loaded. In this case, you can enable the transaction splitting feature.

When transaction splitting is enabled and the default isolation level READ COMMITTED is used, the ApsaraDB RDS instance starts a transaction only for write requests when autocommit is disabled (set autocommit=0). Read requests that arrive before the transaction is started are distributed to read-only instances by the load balancer.

Note

- Explicit transactions do not support splitting, such as transactions started by using the BEGIN or START statement.
- After you enable the transaction splitting feature, global consistency cannot be ensured. Before you enable this feature, we recommend that you evaluate whether this feature is suitable for your workloads.



Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Proxy**.

5. On the **Proxy Service** tab, click **Enable** to the right of **Transaction Splitting**.

 **Note**

- When you no longer need transaction splitting, you can click **Disable** to the right of **Transaction Splitting**.
- The operation to enable or disable transaction splitting takes effect only on new connections.

8.8.4. Read/write splitting

8.8.4.1. Enable read/write splitting

This topic describes the read/write splitting feature of ApsaraDB RDS for MySQL in its dedicated proxy feature and how to enable this feature.

Prerequisites

- Your ApsaraDB RDS instance runs one of the following MySQL versions and RDS editions:
 - MySQL 5.6 on RDS High-availability Edition (with local SSDs)
 - MySQL 5.7 on RDS Enterprise Edition
 - MySQL 5.7 on RDS High-availability Edition (with local SSDs)
 - MySQL 8.0 on RDS High-availability Edition (with local SSDs)
- The database proxy or dedicated database proxy feature is enabled. For more information, see [Enable the dedicated proxy feature](#).
- At least one read-only instance is created. For more information about how to create a read-only instance, see [Create a read-only instance](#).

Context

If your primary instance needs to process a large number of read requests but only a small number of write requests, you can create one or more read-only instances to offload read requests from your primary instance. This ensures service stability. For more information, see [Create a read-only instance](#).

After you create read-only instances, you can enable read/write splitting. In this case, a read/write splitting endpoint is provided. After you add the endpoint to your application, write requests are routed to the primary instance and read requests are routed to the read-only instances.

□

Differences between the read/write splitting endpoint and the internal and public endpoints

- After you enable read/write splitting and add the read/write splitting endpoint to your application, all requests are first routed to this endpoint, and then to the primary and read-only instances based on the request types and the read weights of these instances.
- If the internal or public endpoint of the primary instance is added to your application, all requests are routed to the primary instance. In this case, you must add the endpoints and read weights of the primary and read-only instances to your application to implement read/write splitting.

Logic to route requests

- The following requests are routed only to the primary instance:
 - All data manipulation language (DML) operations, such as INSERT, UPDATE, DELETE, and SELECT FOR UPDATE.
 - All data definition language (DDL) statements used to perform operations such as creating databases or tables, deleting databases or tables, and changing schemas or permissions.
 - All requests that are encapsulated in transactions.

- Requests for user-defined functions.
 - Requests for stored procedures.
 - Requests for EXECUTE statements.
 - Multi-statements. For more information, see [Multi-Statement](#).
 - Requests that involve temporary tables.
 - Requests for SELECT last_insert_id() statements.
 - All requests to query or modify user environment variables.
 - Requests for SHOW PROCESSLIST statements.
 - All requests for KILL statements in SQL (not KILL commands in Linux).
- The following requests are routed to the primary instance or its read-only instances:
 - Read requests that are not encapsulated in transactions.
 - Requests for COM_STMT_EXECUTE statements.
 - The following requests are routed to all the primary and read-only instances:
 - All requests to modify system environment variables.
 - Requests for USE statements.
 - Requests for COM_STMT_PREPARE statements.
 - Requests for COM_CHANGE_USER, COM_QUIT, and COM_SET_OPTION statements.

Benefits

- Easier maintenance by using a unified endpoint

If you do not enable the read/write splitting feature, you must add the endpoints of the primary and read-only instances to your application. After you add the endpoints, your database system routes write requests to the primary instance and read requests to the read-only instances.

If you enable the read/write splitting feature, you can use a dedicated proxy endpoint to implement read/write splitting. After your application is connected to this endpoint, your database system routes read and write requests to the primary and read-only instances based on the read weights of these instances. This reduces maintenance costs.

You can also improve the read capability of your database system by creating read-only instances. You do not need to modify the configuration data on your application.

- Higher performance and lower maintenance costs by using a native link

You can build your own proxy layer on the cloud to implement read/write splitting. In this case, data needs to be parsed and forwarded by multiple components before the data reaches your database system. As a result, response latencies increase. The read/write splitting feature is built in the ApsaraDB RDS ecosystem and can efficiently reduce response latencies, increase processing speeds, and reduce maintenance costs.

- Ideal in various use scenarios based on configurable read weights and thresholds

You can specify the read weights of the primary and read-only instances. You can also specify the latency threshold for data replication to the read-only instances.

- High availability based on instance-level health checks

The read/write splitting feature enables ApsaraDB RDS to actively check the health status of the primary and read-only instances. If a read-only instance unexpectedly exits or its data replication latency exceeds the specified threshold, ApsaraDB RDS stops routing read requests to the instance. ApsaraDB RDS redirects the read requests that are destined for the faulty read-only instance to other healthy instances in your database system. This ensures service availability in the event of faults on individual read-only instances. After the faulty read-only instance is recovered, ApsaraDB RDS resumes routing read requests to the instance.

 **Note** We recommend that you create at least two read-only instances to mitigate the impacts of single points of failure (SPOFs).

Precautions

- When you change the specifications of your ApsaraDB RDS instance or its read-only instances, a network interruption may occur.
- After you create a read-only instance, only the requests over new connections can be routed to the read-only instance.
- The dedicated proxy endpoint does not support SSL encryption.
- The dedicated proxy endpoint does not support compression.
- If the endpoint of the dedicated proxy is used for connection, all the requests encapsulated in transactions are routed to the primary instance.
- If you use the dedicated proxy endpoint to implement read/write splitting, the read consistency of the requests that are not encapsulated in transactions cannot be guaranteed. If you require read consistency, you must encapsulate these requests in transactions.
- If you connect your application to the dedicated proxy endpoint, the `SHOW PROCESSLIST` statement returns a combination of results from the primary ApsaraDB RDS instance and all of its read-only instances.
- If the short-lived connection optimization feature is enabled, the `SHOW PROCESSLIST` statement may return idle connections.
- If you execute **multi-statements** or stored procedures, the read/write splitting feature is disabled and all the subsequent requests over the current connection are routed to the primary ApsaraDB RDS instance. To enable the read/write splitting feature again, you must close the current connection and establish a new connection.
- The dedicated proxy feature supports the `/*FORCE_MASTER*/` and `/*FORCE_SLAVE*/` hints. However, requests that contain hints have the highest route priorities and are not constrained by consistency or transaction limits. Before you use these hints, you must check whether these hints are suitable for your workloads. In addition, these hints cannot contain statements such as `/*FORCE_SLAVE*/ set names utf8;`. Such statements can change environment variables. If you include such statements in these hints, errors may occur when you process your subsequent workloads.

Prerequisites

A read-only instance is created for the primary instance. For more information, see [Create a read-only instance](#).

Enable read/write splitting

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Proxy**.
5. On the **Read/Write Splitting** tab, click **Enable now**.
6. Configure the following parameters.

| Parameter | Description |
|-----------|-------------|
| | |

| Parameter | Description |
|---------------------------------|--|
| Latency Threshold | <p>The maximum latency that is allowed for data replication from the primary instance to its read-only instances. If the latency of data replication to a read-only instance exceeds the specified threshold, ApsaraDB RDS stops routing read requests to the instance. This applies even if the instance has a high read weight.</p> <p>Valid values: 0 to 3600. Unit: seconds. The read-only instances may replicate data from the primary instance at a specific latency due to SQL statement execution limits. We recommend that you set this parameter to a value that is greater than or equal to 30.</p> |
| Read Weight Distribution | <p>The read weight of each instance in your database system. A higher read weight indicates more read requests to process. For example, the primary instance is attached with three read-only instances, and the read weights of the primary and read-only instances are 0, 100, 200, and 200. In this case, the primary instance processes only write requests, and the three read-only instances process all of the read requests at the 1:2:2 ratio.</p> <ul style="list-style-type: none"> ◦ Automatic Distribution: Your database system assigns a read weight to each instance based on the instance specifications. After you create a read-only instance, ApsaraDB RDS automatically assigns a read weight to the instance and adds the instance to the read/write splitting link. ◦ Customized Distribution: You must manually specify the read weight of each instance. Valid values: 0 to 10000. After you create a read-only instance, ApsaraDB RDS sets the read weight of the read-only instance to 0. You must manually modify the read weight of the created read-only instance. |

7. Click OK.

8.8.4.2. Configure read/write splitting

This topic describes how to configure the latency threshold and specify read weights for an ApsaraDB RDS instance in the ApsaraDB RDS console.

Prerequisites

- Your ApsaraDB RDS instance runs one of the following MySQL versions and RDS editions:
 - MySQL 5.6 on RDS High-availability Edition (with local SSDs)
 - MySQL 5.7 on RDS Enterprise Edition
 - MySQL 5.7 on RDS High-availability Edition (with local SSDs)
 - MySQL 8.0 on RDS High-availability Edition (with local SSDs)
- Read/write splitting is enabled. For more information, see [Enable read/write splitting](#).

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Proxy**.
5. On the **Read/Write Splitting** tab, click **Configure Read/Write Splitting**.
6. Configure the following parameters.

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|---------------------------------|--|
| Latency Threshold | <p>The maximum latency that is allowed for data replication from the primary instance to its read-only instances. If the latency of data replication to a read-only instance exceeds the specified threshold, ApsaraDB RDS stops routing read requests to the instance. This applies even if the instance has a high read weight.</p> <p>Valid values: 0 to 3600. Unit: seconds. The read-only instances may replicate data from the primary instance at a specific latency due to SQL statement execution limits. We recommend that you set this parameter to a value that is greater than or equal to 30.</p> |
| Read Weight Distribution | <p>The read weight of each instance in your database system. A higher read weight indicates more read requests to process. For example, the primary instance is attached with three read-only instances, and the read weights of the primary and read-only instances are 0, 100, 200, and 200. In this case, the primary instance processes only write requests, and the three read-only instances process all of the read requests at the 1:2:2 ratio.</p> <ul style="list-style-type: none"> ◦ Automatic Distribution: Your database system assigns a read weight to each instance based on the instance specifications. After you create a read-only instance, ApsaraDB RDS automatically assigns a read weight to the instance and adds the instance to the read/write splitting link. ◦ Customized Distribution: You must manually specify the read weight of each instance. Valid values: 0 to 10000. After you create a read-only instance, ApsaraDB RDS sets the read weight of the read-only instance to 0. You must manually modify the read weight of the created read-only instance. |

7. Click OK.

8.8.4.3. Disable read/write splitting

This topic describes how to disable the read/write splitting feature of an ApsaraDB RDS instance in the ApsaraDB RDS console.

Prerequisites

- Your ApsaraDB RDS instance runs one of the following MySQL versions and RDS editions:
 - MySQL 5.6
 - MySQL 5.7 on RDS Enterprise Edition
 - MySQL 5.7 on RDS High-availability Edition (with local SSDs)
 - MySQL 8.0 on RDS High-availability Edition (with local SSDs)
- Read/write splitting is enabled. For more information, see [Enable read/write splitting](#).

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Proxy**.
5. On the **Read/Write Splitting** tab, click **Disable Read/Write Splitting**.
6. Click **Confirm**.

8.8.4.4. Upgrade an ApsaraDB RDS for MySQL instance from shared proxy to dedicated proxy

To prevent stability, scalability, and performance issues with the shared proxy feature, ApsaraDB RDS for MySQL provides the dedicated proxy feature. If your instance has the shared proxy feature enabled and needs to use the read/write splitting feature, you can upgrade your instance from shared proxy to dedicated proxy. This topic describes how to upgrade an ApsaraDB RDS for MySQL instance from shared proxy to dedicated proxy.

Context

The read/write splitting feature of ApsaraDB RDS for MySQL is implemented based on database proxies. For some existing ApsaraDB RDS instances that run MySQL 5.6 or 5.7, the read/write splitting feature is implemented based on shared proxies. However, shared proxies cannot ensure service stability. We recommend that you upgrade these instances from shared proxy to dedicated proxy. Dedicated proxies have the following advantages over shared proxies:

- Dedicated proxies provide better stability and isolation.
- Dedicated proxies provide higher performance.
- Dedicated proxies support scaling. You can increase the number of dedicated proxies to handle more workloads.
- Dedicated proxies support performance monitoring. You can adjust the number of dedicated proxies based on the monitoring data and your business plan. The adjustment takes effect immediately after it is applied.
- A unified dedicated proxy endpoint is provided. This reduces maintenance costs by relieving the need to modify the endpoint configuration on your application. In addition, the dedicated proxy endpoint remains valid unless you release the dedicated proxies. For example, you may enable read/write splitting during peak hours, and then release read-only instances and disable read/write splitting after peak hours. In these cases, you do not need to update the endpoints on your application because the proxy endpoint remains connected.
- A unified dedicated proxy endpoint is used to implement features such as read/write splitting, short-lived connection optimization, and transaction splitting.

Prerequisites

- The shared proxy feature is enabled for your ApsaraDB RDS instance.
- Your instance runs MySQL 5.7 on RDS High-availability Edition (in minor engine version 20190925 or later) or MySQL 5.6 on RDS High-availability Edition (in minor engine version 20200229 or later).

Note

- If your ApsaraDB RDS instance runs MySQL 5.6 on RDS Enterprise Edition, you cannot upgrade the instance from shared proxy to dedicated proxy.
- If a sample error message of **current db not support db proxy** is reported, upgrade the minor engine version first. For more information, see [Upgrade the minor version of an instance](#).

- If the read/write splitting feature is not enabled for your instance, upgrade your instance to the Linux Virtual Server (LVS) link instead of upgrading from shared proxy to dedicated proxy.

Precautions

- During the upgrade, the endpoints of the primary and read-only ApsaraDB RDS instances may be disconnected for 30 seconds, and the read/write splitting endpoint may be unavailable for 30 seconds. We recommend that you perform upgrade operations during off-peak hours.
- The number of allowed connections is determined by the lowest connection specifications among the primary and read-only ApsaraDB RDS instances. We recommend that you specify the same connection specifications for these instances.
- If you create a read-only ApsaraDB RDS instance after you upgrade to the dedicated proxy feature, only the requests over new connections are routed to the new read-only instance. This also applies if you restart a read-only instance.
- The `max_prepared_stmt_count` parameter must be set to the same value for the primary and read-only ApsaraDB RDS instances.
- For more information, see [Dedicated proxy](#).

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Proxy**.
5. On the **Database Proxy** page, click the **Read/Write Splitting** tab.
6. In the **Basic Information of Read/Write Splitting** section, click **Upgrade to Dedicated Proxy**.
7. In the message that appears, click **OK**.

8.9. Monitoring and alerts

8.9.1. View resource and engine monitoring data

The ApsaraDB RDS console provides a variety of performance metrics to monitor the status of your instances.

Prerequisites

The ApsaraDB RDS instance runs one of the following MySQL versions and RDS editions:

- MySQL 5.6
- MySQL 5.7 on RDS High-availability Edition with local SSDs
- MySQL 5.7 on RDS High-availability Edition with standard SSDs

Procedure

1. For more information, see [Log on to the ApsaraDB RDS console.](#)
2. Find an instance and click the instance ID.
3. In the left-side navigation pane, click **Monitoring and Alerts**.
4. On the **Monitoring and Alerts** page, select **Resource Monitoring** or **Engine Monitoring**, and select a time range to view the corresponding monitoring data. The following table describes the metrics.

| Category | Metric | Description |
|---------------------|---|--|
| Resource Monitoring | Disk Space (MB) | <p>The disk space usage of the instance. It consists of the following items:</p> <ul style="list-style-type: none"> ◦ Instance size ◦ Data usage ◦ Log size ◦ Temporary file size ◦ Other system file size <p>Unit: MB.</p> |
| | IOPS (Input/Output Operations per Second) | The number of input/output operations per second (IOPS) of the instance. |
| | Total Connections | The number of active connections to the instance and the total number of connections to the instance. |

| Category | Metric | Description |
|-------------------|---|--|
| | CPU Utilization and Memory Usage (%) | The CPU utilization and memory usage of the instance. These metrics do not include the CPU utilization and memory usage for the operating system. |
| | Network Traffic (KB) | The inbound and outbound traffic of the instance per second. Unit: KB. |
| Engine Monitoring | Transactions per Second (TPS)/Queries per Second (QPS) | The average number of transactions per second and the average number of SQL statements executed per second. |
| | InnoDB Buffer Pool Read Hit Ratio, Usage Ratio, and Dirty Block Ratio (%) | The read hit ratio, usage ratio, and dirty block ratio of the InnoDB buffer pool. |
| | InnoDB Read/Write Volume (KB) | The amount of data that InnoDB reads and writes per second. Unit: KB. |
| | InnoDB Buffer Pool Read/Write Frequency | The number of read and write operations that InnoDB performs per second. |
| | InnoDB Log Read/Write/fsync | The average frequency of physical writes to log files per second by InnoDB, the log write request frequency, and the average frequency of fsync writes to log files. |
| | Temporary Tables Automatically Created on Hard Disk when MySQL Statements Are Executed | The number of temporary tables that are automatically created on the hard disk when the database executes SQL statements. |
| | MySQL_COMDML | The number of SQL statements that the database executes per second. The following SQL statements are included: <ul style="list-style-type: none"> ◦ Insert ◦ Delete ◦ Insert_Select ◦ Replace ◦ Replace_Select ◦ Select ◦ Update |
| | MySQL_RowDML | The numbers of operations that InnoDB performs per second. The following items are included: <ul style="list-style-type: none"> ◦ The number of physical writes to log files per second. ◦ The number of rows that are read, updated, deleted, and inserted from InnoDB tables per second. |
| | | |

| Category | Metric | Description |
|----------|---|---|
| | MyISAM Read/Write Frequency | The numbers of operations that MyISAM performs per second. The following items are included: <ul style="list-style-type: none"> ◦ The number of MyISAM reads and writes from the buffer pool per second. ◦ The number of MyISAM reads and writes from the hard disk per second. |
| | MyISAM Key Buffer Read/Write/Usage Ratio (%) | The read hit ratio, write hit ratio, and usage of the MyISAM key buffer per second. |

8.9.2. Set a monitoring frequency

The ApsaraDB RDS console provides a variety of performance metrics for which you can set a monitoring frequency.

Prerequisites

The ApsaraDB RDS instance runs one of the following MySQL versions and RDS editions:

- MySQL 5.6
- MySQL 5.7 on RDS High-availability Edition with local SSDs
- MySQL 5.7 on RDS High-availability Edition with standard SSDs

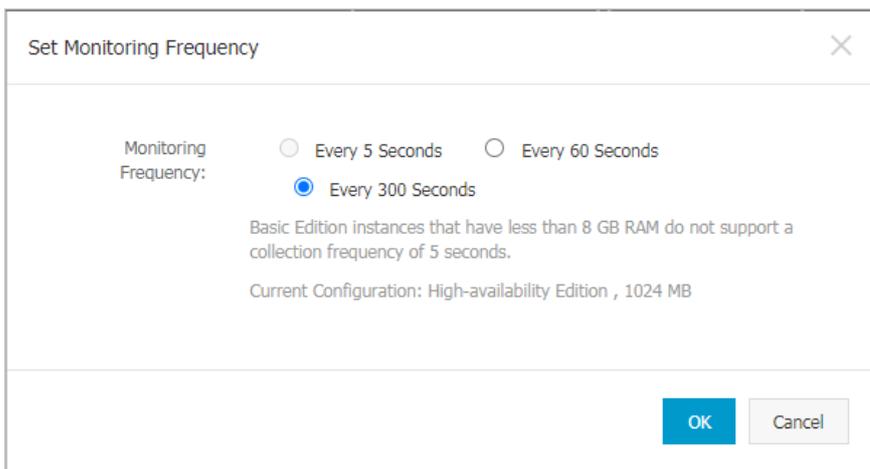
Context

ApsaraDB RDS provides the following monitoring frequencies:

- Every 5 seconds for the first seven days. After the seventh day, performance metrics are monitored every minute.
- Every 60 seconds.
- Every 300 seconds.

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Monitoring and Alerts**.
5. On the **Monitoring** tab, click **Set Monitoring Frequency**.
6. In the **Set Monitoring Frequency** dialog box, select a new monitoring frequency.



Note If the RDS instance runs the RDS Basic Edition or its memory capacity is less than 8 GB, the Every 5 Seconds monitoring frequency is not supported.

7. Click **OK**.

8.10. Data security

8.10.1. Configure a whitelist

To ensure database security and reliability, you must modify the whitelist of an ApsaraDB RDS instance before you enable the instance. You must add the IP addresses or CIDR blocks that are used for database access to the whitelist.

Context

The whitelist improves the access security of your ApsaraDB RDS instance. We recommend that you maintain the whitelist on a regular basis. The whitelist configuration process does not affect the normal operations of the ApsaraDB RDS instance.

To configure a whitelist, perform the following operations:

- Configure a whitelist: Add IP addresses to allow them to connect to the ApsaraDB RDS instance.
- Configure an ECS security group: Add an ECS security group for the ApsaraDB RDS instance to allow ECS instances in the group to connect to the ApsaraDB RDS instance.

Precautions

- The default whitelist can be modified or cleared, but cannot be deleted.
- You can add up to 1,000 IP addresses or CIDR blocks to a whitelist. If you want to add a large number of IP addresses, we recommend that you merge them into CIDR blocks, such as 192.168.1.0/24.

Configure a standard IP address whitelist

1. For more information, see [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find an instance. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
3. In the left-side navigation pane, click **Data Security**.
4. On the **Whitelist Settings** tab, click **Edit** corresponding to the **default** whitelist.

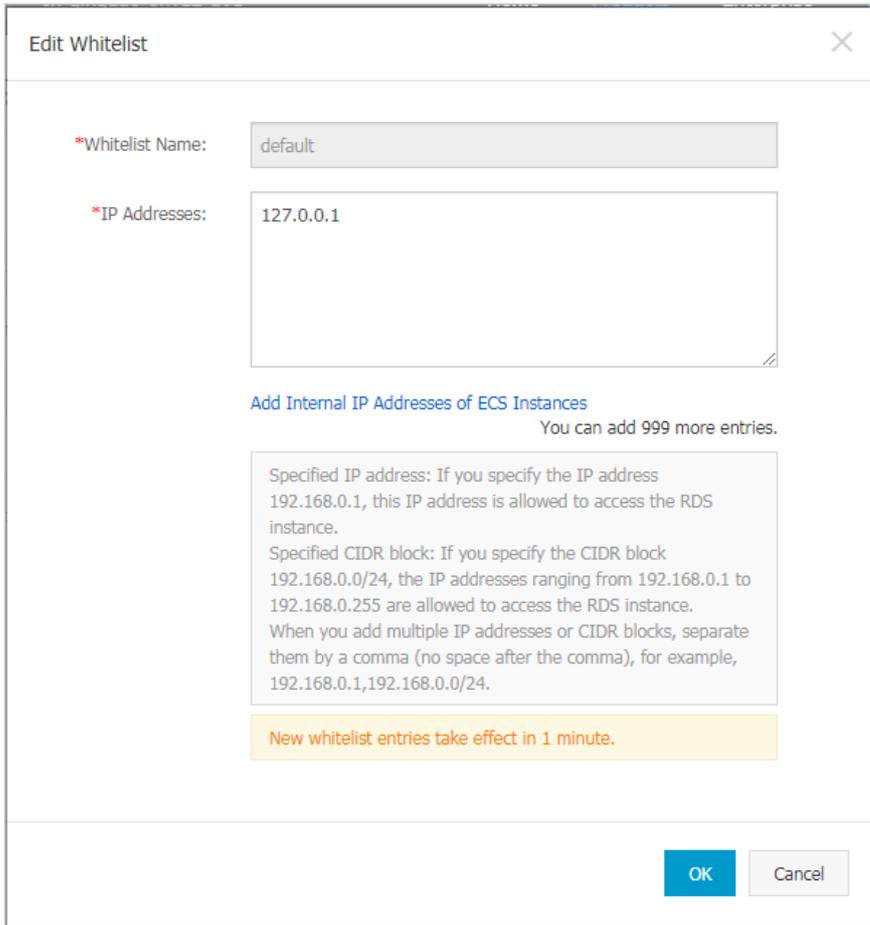


Note

- If you want to connect an ECS instance to an ApsaraDB RDS instance by using an internal endpoint, you must make sure that the two instances are in the same region and have the same network type. Otherwise, the connection fails.
- You can also click **Create Whitelist** to create a new whitelist.

5. In the **Edit Whitelist** dialog box, enter the IP addresses or CIDR blocks that are allowed to access your ApsaraDB RDS instance, and then click **OK**.
 - If you add the CIDR block 10.10.10.0/24, all IP addresses in the 10.10.10.X format are allowed to access the ApsaraDB RDS instance.
 - If you enter more than one IP address or CIDR block, you must separate them with commas (,). Do not add spaces before or after the commas. Example: 192.168.0.1,172.16.213.9.
 - If you click **Add Internal IP Addresses of ECS Instances**, the IP addresses of all of the ECS instances that are created in your Alibaba Cloud account appear. Then, you can select the required IP addresses and add them to the whitelist.

 **Note** If you add a new IP address or CIDR block to the **default** whitelist, the default address 127.0.0.1 is deleted.



8.10.2. Configure SSL encryption

This topic describes how to enhance endpoint security. You can enable Secure Sockets Layer (SSL) encryption and install SSL certificates that are issued by certificate authorities (CAs) to the required application services. SSL is used at the transport layer to encrypt network connections and enhance the security and integrity of communication data. However, SSL increases the response time.

Prerequisites

Your ApsaraDB RDS instance runs one of the following MySQL versions and RDS editions:

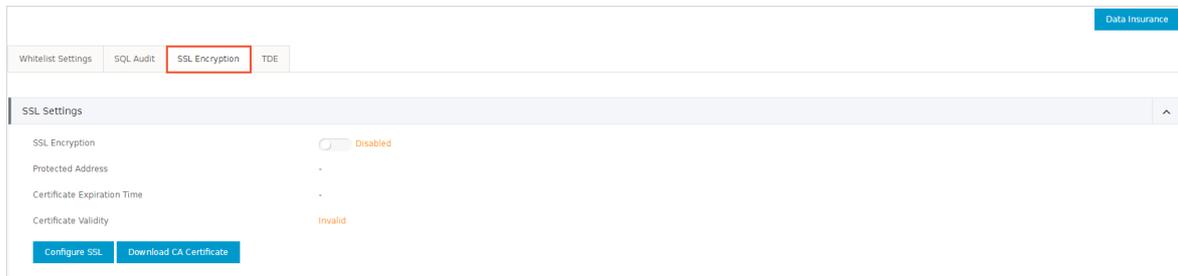
- MySQL 8.0 on RDS High-availability Edition (with local SSDs)
- MySQL 5.7 on RDS High-availability Edition (with local SSDs)
- MySQL 5.6 on RDS High-availability Edition (with local SSDs)

Precautions

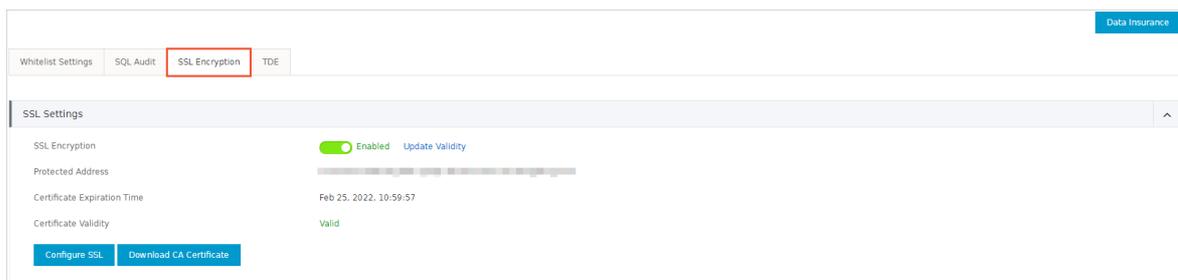
- An SSL CA certificate is valid for one year. You must update the validity period of the SSL CA certificate in your application or client within one year. Otherwise, your application or client that uses encrypted network connections cannot connect to the ApsaraDB RDS instance.
- SSL encryption may cause a significant increase in CPU utilization. We recommend that you enable SSL encryption only when you want to encrypt connections from the Internet. In most cases, connections that use an internal endpoint do not require SSL encryption.
- Read/write splitting endpoints do not support SSL encryption.
- If you disable SSL encryption, the ApsaraDB RDS instance restarts. Proceed with caution.

Enable SSL encryption

1. Log on to the ApsaraDB for RDS console.
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. Click the **SSL Encryption** tab.



6. In the **SSL Settings** section, turn on **SSL Encryption**.
7. In the **Configure SSL** dialog box, select the endpoint for which you want to enable SSL encryption and click **OK**.
8. Click **Download CA Certificate** to download the SSL CA certificate files in a compressed package.



The downloaded package contains the following files:

- o P7B file: used to import CA certificates to the Windows system.
- o PEM file: used to import CA certificates to other operating systems or applications.
- o JKS file: the Java truststore file. The password is `apsaradb`. It is used to import the CA certificate chain to Java programs.

Note When the JKS file is used in Java, you must modify the default JDK security configuration in JDK 7 and JDK 8. Open the `/jre/lib/security/java.security` file on the host where your application resides, and modify the following configurations:

```
jdk.tls.disabledAlgorithms=SSLv3, RC4, DH keySize < 224
jdk.certpath.disabledAlgorithms=MD2, RSA keySize < 1024
```

If you do not modify the JDK security configuration, the following error is reported. Similar errors are also caused by the Java security configuration.

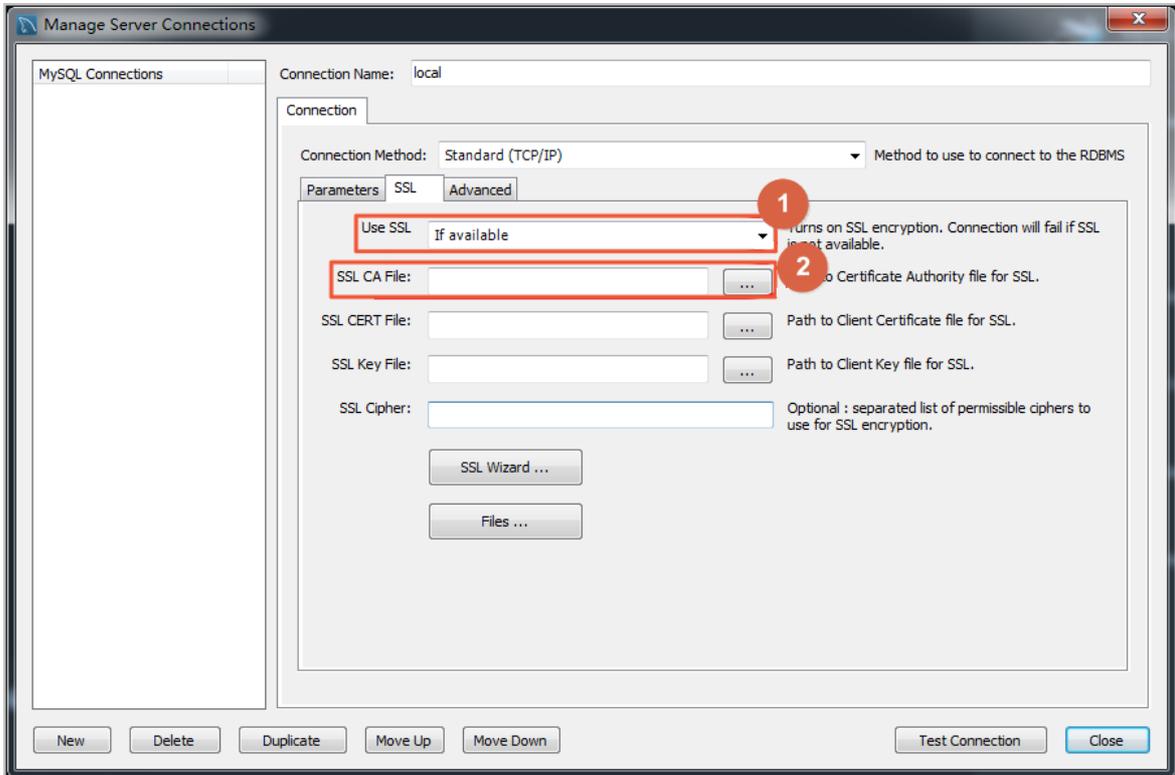
```
javax.net.ssl.SSLHandshakeException: DHPublicKey does not comply to algorithm constraints
```

Configure an SSL CA certificate

After you enable SSL encryption, configure the SSL CA certificate on your application or client before they can connect to the ApsaraDB RDS instance. This section describes how to configure an SSL CA certificate. MySQL Workbench and Navicat are used in the example. If you are using other applications or clients, see the related instructions.

Configure a certificate on MySQL Workbench

1. Start MySQL Workbench.
2. Choose **Database > Manage Connections**.
3. In the **Connection** section, click the **SSL** tab and configure the following parameters.



- ①: Enable Use SSL.
- ②: Import the SSL CA certificate file.

Configure a certificate on Navicat

1. Start Navicat.
2. Right-click the database and select **Edit Connection**.
3. Click the **SSL** tab. Select the path of the SSL CA certificate file in the PEM format.
4. Click **OK**.

Note If the `connection is being used` error is reported, the previous session is still connected. Restart Navicat.

5. Double-click the database to test whether the database is connected.

Update the validity period of an SSL CA certificate

Note Update Validity causes the ApsaraDB RDS instance to restart. Proceed with caution.

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.

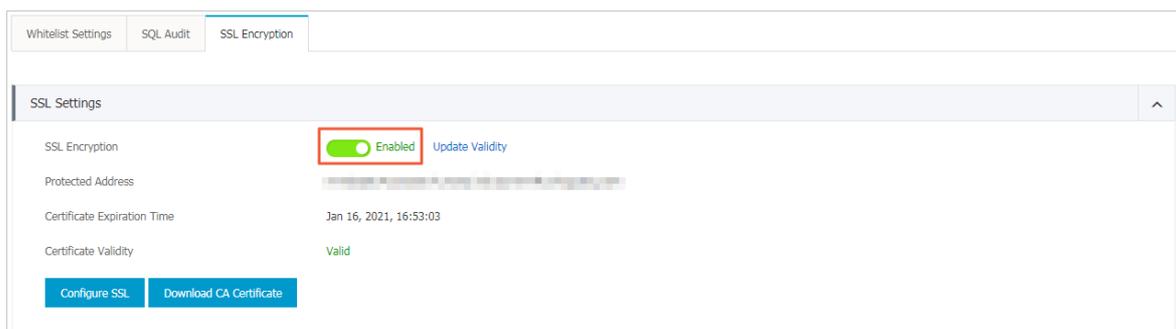
5. Click the **SSL Encryption** tab.
6. Click **Update Validity**.

Disable SSL encryption

Note

- If you disable SSL encryption, the ApsaraDB RDS instance restarts. To reduce the impact on your business, the system triggers a primary/secondary switchover. We recommend that you disable SSL encryption during off-peak hours.
- After you disable SSL encryption, access performance increases, but security decreases. We recommend that you disable SSL encryption only in secure environments.

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. Click the **SSL Encryption** tab.
6. In the **SSL Settings** section, turn off **SSL Encryption**. In the message that appears, click **OK**.



8.10.3. Configure TDE

This topic describes how to configure Transparent Data Encryption (TDE) for your ApsaraDB RDS for MySQL instance. TDE encrypts and decrypts data files in real time. It encrypts data files when they are written to disks, and decrypts data files when they are loaded to the memory from disks. TDE does not increase the size of data files. You can use TDE without the need to make changes to applications.

Prerequisites

- Your ApsaraDB RDS instance runs one of the following MySQL versions and RDS editions:
 - MySQL 5.7 on RDS High-availability Edition (with local SSDs)
 - MySQL 5.6 on RDS High-availability Edition (with local SSDs)
 - MySQL 8.0 on RDS High-availability Edition (with local SSDs)
- Key Management Service (KMS) is activated. If KMS is not activated, you can activate it when you enable TDE.

Context

The key used for TDE is created and managed by KMS. ApsaraDB RDS does not provide the key or certificates that are required for encryption. For specific zones, you can use the keys that are automatically generated by Apsara Stack, or you can use your own key materials to generate data keys and authorize your ApsaraDB RDS for MySQL instance to use these keys.

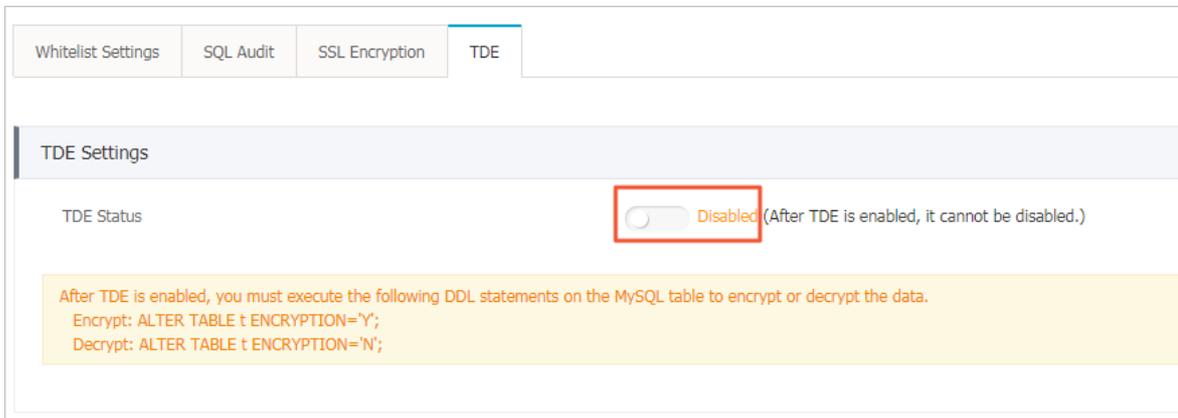
Precautions

- When TDE is being enabled, your ApsaraDB RDS for MySQL instance is restarted and all services are disconnected. Make appropriate service arrangements before you enable TDE. Proceed with caution.
- You cannot disable TDE after it is enabled.
- You cannot change the key used for encryption after TDE is enabled.
- If you want to restore the data to your computer after TDE is enabled, you must decrypt data on your ApsaraDB RDS for MySQL instance. For more information, see the "[Decrypt a table](#)" section of this topic.
- After TDE is enabled, CPU utilization significantly increases.
- If you use an existing custom key for encryption, take note of the following items:
 - If you disable a key, set a key deletion plan, or delete the key materials, the key becomes unavailable.
 - If you revoke the key that is authorized for an ApsaraDB RDS for MySQL instance, the instance becomes unavailable after it is restarted.
 - You must use an Apsara Stack tenant account or an account that has the `AliyunSTSAssumeRoleAccess` permission.

 **Note** For more information, see topics about key management in *Key Management Service User Guide*.

Use a key that is automatically generated by Apsara Stack

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. Click the **TDE** tab.
6. In the **TDE Settings** section, turn on **TDE Status**.



7. In the dialog box that appears, select **Use an Automatically Generated Key** and click **OK**.

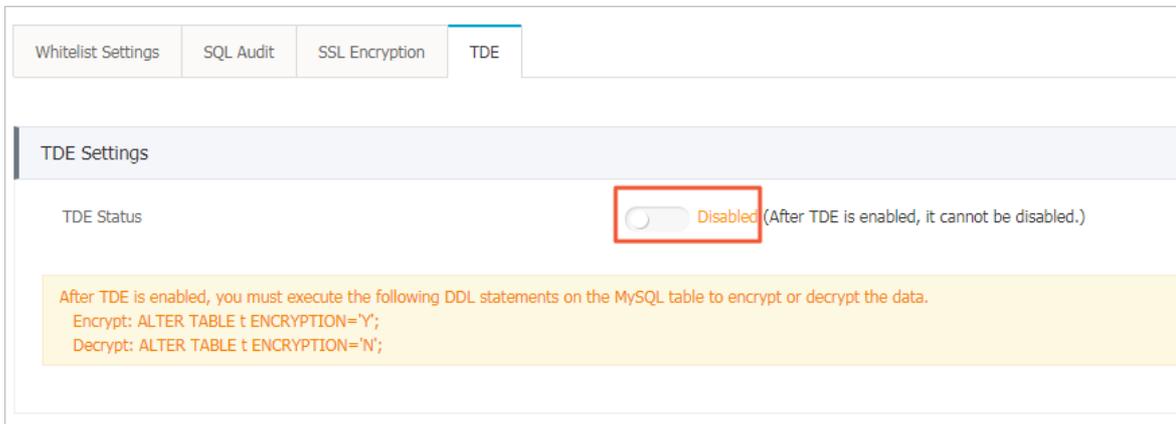
 **Note** If the instance runs MySQL 5.7 on RDS High-availability Edition, you can select one of the following encryption methods:

- SM4 encryption
- AES_256_CBC encryption

Use an existing custom key

1. [Log on to the ApsaraDB for RDS console](#).

2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. Click the **TDE** tab.
6. In the **TDE Settings** section, turn on **TDE Status**.



7. In the dialog box that appears, select **Use an Existing Custom Key** and click **OK**.

Note If you do not have a custom key, click **create a key** to go to the KMS console and import the key materials. For more information, see *Create a key in Key Management Service User Guide*.

Encrypt a table

Log on to the database and execute the following statement to encrypt a table:

- MySQL 5.6

```
alter table <tablename> engine=innodb,block_format=encrypted;
```

- MySQL 5.7 and MySQL 8.0

```
alter table <tablename> encryption='Y';
```

Decrypt a table

Execute the following statement to decrypt a table that is encrypted by using TDE:

- MySQL 5.6

```
alter table <tablename> engine=innodb,block_format=default;
```

- MySQL 5.7 and MySQL 8.0

```
alter table <tablename> encryption='N';
```

FAQ

- Q: Can common database tools such as Navicat be used after TDE is enabled?
A: Yes, after you enable TDE, you can still use common database tools such as Navicat.
- Q: Why is data still displayed in plaintext after it is encrypted?

A: After you enable TDE, your data is stored in ciphertext. However, when the data is queried, it is decrypted and loaded into memory as plaintext. TDE encrypts backup files to prevent data leaks. The encrypted backup files cannot be used to restore data to your computer. If you want to restore these backup files to your computer, you must decrypt them. For more information, see the "[Decrypt a table](#)" section of this topic.

8.10.4. SQL audit

You can use the SQL audit feature to audit SQL executions and check the details. SQL audit does not affect instance performance.

Context

 **Note** You cannot view the logs that are generated before you enable SQL audit.

You can view the incremental data of your ApsaraDB RDS for MySQL instance in SQL audit logs or binlogs. However, these two methods differ in the following aspects:

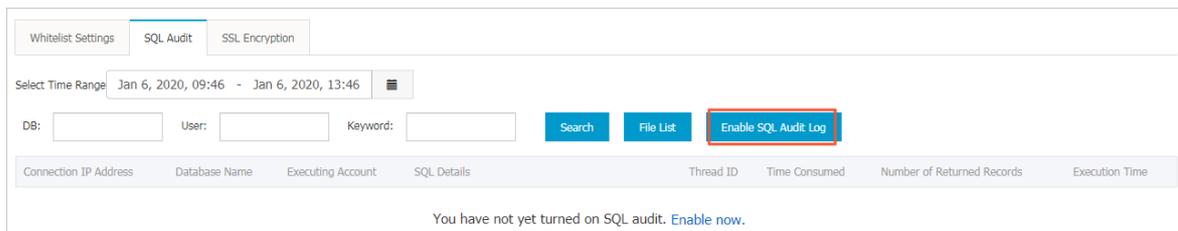
- SQL audit logs are similar to audit logs in MySQL and record all DML and DDL operations by using network protocol analysis. SQL audit does not parse the actual parameter values. Therefore, a small amount of information may be lost if a large number of SQL statements are executed to query data. The incremental data obtained by using this method may be inaccurate.
- Binlogs record all add, delete, and modify operations and the incremental data used for data restoration. Binlogs are temporarily stored in your ApsaraDB RDS instance after they are generated. The system transfers full binlog files to OSS on a regular basis. OSS then stores the files for seven days. However, a binlog file cannot be transferred if data is being written to it. Such binlog files cannot be uploaded to OSS after you click **Upload Binlogs** on the **Backup and Restoration** page. Binlogs are not generated in real time, but you can obtain accurate incremental data from them.

Precautions

- SQL audit is disabled by default. SQL audit does not affect instance performance.
- SQL audit logs are retained for 30 days.
- Log files exported from SQL audit are retained for two days. The system clears files that are retained for more than two days.

Enable SQL audit

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. Click the **SQL Audit** tab.



6. Click **Enable SQL Audit**.
7. In the message that appears, click **Confirm**.

After SQL audit is enabled, you can query SQL information based on conditions such as the time range, database, user, and keyword.

Disable SQL audit

Note If SQL audit is disabled, all SQL audit logs are deleted. We recommend that you export and store audit logs to your computer before you disable SQL audit.

You can disable SQL audit to avoid charges when you do not need it. To disable SQL audit, perform the following operations:

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. Click the **SQL Audit** tab.

6. Click **Export File** to export and store the SQL audit content to your computer.
7. After the file is exported, click **Disable SQL Audit**.
8. In the message that appears, click **Confirm**.

8.11. Service availability

8.11.1. Configure automatic or manual switchover

This topic describes how to automatically or manually switch over services between primary and secondary instances. After a switchover, the original primary instance becomes a secondary instance.

Context

- **Automatic switchover:** the default switchover mode. If the primary instance experiences a fault, your ApsaraDB RDS services are automatically switched over to the secondary instance.

Note You can click **Switch Primary/Secondary Instance** on the **Service Availability** page of an ApsaraDB RDS for MySQL instance with standard or enhanced SSDs to disable automatic switchover. This facilitates troubleshooting when errors occur on the primary instance.

- **Manual switchover:** allows you to manually switch over services between primary and secondary instances.

Note Data is synchronized between the primary and secondary instances in real time. You can connect only to the primary instance. The secondary instances serve only as backups and do not allow external access.

Precautions

- Services may be disconnected during a switchover. Make sure that your applications are configured with automatic reconnection policies to avoid service disruptions.
- If the primary instance is attached with read-only instances, data on the read-only instances shows a latency of several minutes after a switchover. This is because it takes time to re-establish replication connections and synchronize incremental data.

Manually switch over services between primary and secondary instances

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Service Availability**.
5. Click **Switch Primary/Secondary Instance** on the right side of the page.

 **Note** Services may be disconnected once or twice during the switchover. Make sure that your applications are configured with automatic reconnection policies to avoid service disruptions.

6. In the dialog box that appears, click **OK**.

FAQ

Q: Can I connect to secondary instances?

A: No, you cannot connect to secondary instances. You can connect only to primary instances. Secondary instances serve only as backups and do not allow external access.

8.11.2. Change the data replication mode

You can set the data replication mode between primary and secondary ApsaraDB RDS instances to improve database availability.

Prerequisites

Your ApsaraDB RDS instance runs one of the following MySQL versions and RDS editions:

- MySQL 5.6 on RDS High-availability Edition (with local SSDs)
- MySQL 5.7 on RDS High-availability Edition (with local SSDs)
- MySQL 8.0 on RDS High-availability Edition (with local SSDs)

Data replication modes

- Semi-synchronous

After an update that is initialized by your application is complete on the primary instance, the log is synchronized to all the secondary instances. After the secondary instances receive the log, the update transaction is considered committed. Your database system does not need to wait for the log to be replayed.

If the secondary instances are unavailable or a network exception occurs between the primary and secondary instances, semi-synchronous replication degrades to the asynchronous mode.

- Asynchronous

When your application initiates a request to add, delete, or modify data, the primary instance responds to your application immediately after it completes the operation. At the same time, the primary instance starts to asynchronously replicate data to its secondary instances. During asynchronous data replication, the unavailability of secondary instances does not affect the operations on the primary instance. Data remains consistent even if the primary instance is unavailable.

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.

4. In the left-side navigation pane, click **Service Availability**.
5. Click **Change Data Replication Mode** on the right side of the page.
6. In the dialog box that appears, select a data replication mode and click **OK**.

FAQ

Q: Which data replication mode is recommended?

A: You can select a data replication mode based on your business requirements. If you require quick responses, we recommend that you select the asynchronous mode. In other scenarios, you can select the semi-synchronous mode.

8.12. Database backup and restoration

8.12.1. Automatic backup

ApsaraDB RDS automatic backup supports full physical backups. ApsaraDB RDS automatically backs up data based on pre-configured policies. This topic describes how to configure a policy for automatic backup.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Backup and Restoration**.
5. Click the **Backup Settings** tab.
6. Click **Edit**.

 **Note** To ensure data security, the system compares the new backup cycle and time with the original settings, and selects the most recent time point to back up the data. Therefore, the next backup may still be performed based on the original backup cycle and time. For example, if the backup time is set to 19:00-20:00 every Wednesday and you modify the time to 19:00-20:00 every Thursday before 19:00 this Wednesday, the system still backs up data during 19:00-20:00 this Wednesday.

Backup Settings
✕

Data Retention Period: Days

Backup Cycle: Monday Tuesday Wednesday
 Thursday Friday Saturday Sunday

Backup Time: ▼

Log Backup: Enabled Disabled

Log Retention Period: Days

OSS Dump Status Enabled Disabled
After database dump is enabled, new backups will be automatically dumped to the specified OSS bucket

OSS Dumped Data Data Backup Log Backup

OSS Bucket: ▼

Restore Individual Database/Table Enabled Disabled
After Restore Individual Database/Table is enabled, the backup format will be changed to support restoring individual databases and tables. This feature cannot be disabled.

7. Configure the following parameters.

| Parameter | Description |
|------------------------------|--|
| Data Retention Period | The number of days for which data backup files are retained. Valid values: 7 to 730. Default value: 7. |
| Backup Cycle | The backup cycle. You can select one or multiple days within a week. |
| Backup Time | A period of time within a day. Unit: hours. We recommend that you back up data during off-peak hours. |
| Log Backup | Specifies whether to enable log backup. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 5px;"> <p> Notice If you disable log backup, all the log backup files are deleted, and you cannot restore data to a saved point in time.</p> </div> |
| Log Retention Period | The number of days for which log backup files are retained. Valid values: 7 to 730. Default value: 7. |

| Parameter | Description |
|-----------------------------------|--|
| Restore Individual Database/Table | <p>Specifies whether to enable restoration of individual databases or tables. You cannot disable this feature after it is enabled.</p> <p>Note Restoration of individual databases or tables can be enabled only on instances with local SSDs. For more information about this feature, see Restore individual databases and tables for an ApsaraDB RDS for MySQL instance.</p> |

8. After you configure the preceding parameters, click **OK**.

8.12.2. Manual backup

Manual backup supports both full physical backups and full logical backups. This topic describes how to manually back up ApsaraDB RDS data.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. Click **Back Up Instance** in the upper-right corner.

5. Set the backup mode and backup policy, and click **OK**.

- Note** Two backup methods are available:
- Physical backup: directly backs up all files in all databases.
 - Logical backup: extracts data from the databases by using SQL statements and backs up the data in the text format. If you select logical backup, you must select a backup policy:
 - Instance Backup: backs up the entire instance.
 - Single-Database Backup: backs up one of the databases in the instance.

8.12.3. Restore individual databases and tables for an ApsaraDB RDS for MySQL instance

This topic describes how to restore the individual databases and tables that are accidentally deleted from an ApsaraDB RDS for MySQL instance. You can restore these databases and tables from backup files.

Prerequisites

- Your ApsaraDB RDS instance runs one of the following MySQL versions and RDS editions:
 - MySQL 5.7 on RDS High-availability Edition with local SSDs
 - MySQL 5.6 on RDS High-availability Edition
- The number of tables on the ApsaraDB RDS instance does not exceed 50,000.
- If you want to restore individual databases and tables of your ApsaraDB RDS instance to the same instance, the ApsaraDB RDS instance meets the following requirements:
 - The ApsaraDB RDS instance is in the Running state and is not locked.
 - The ApsaraDB RDS instance does not have ongoing migration tasks.
 - If you want to restore individual databases and tables to a point in time, the log backup feature is enabled.
 - If you want to restore an instance from a backup set, at least one backup set is available.

 **Note** For more information about how to restore databases at the instance level, see [Restore data to a new instance \(formerly known as cloning an instance\)](#).

Precautions

- If you restore individual databases and tables to the original ApsaraDB RDS instance, a primary/secondary switchover is triggered. This may cause a network interruption. Make sure that your application is configured to automatically reconnect to the original ApsaraDB RDS instance. If you restore individual databases and tables to a new ApsaraDB RDS instance, no primary/secondary switchover is triggered.
- The Restore Individual Database/Table feature restores only the selected tables. You must select all of the tables that you want to restore. The restoration fails in the following scenarios:
 - The selected tables are deleted during the specified period. The specified period spans from the point in time when the last backup set is generated to the point in time to which you want to restore the selected tables.
 - The restoration involves a table that you have not selected. For example, you selected Table B, but Table B was renamed from Table A before the specified point in time. In this case, the restoration fails because you did not select Table A.
- You can select a maximum of 50 databases or tables at a time.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Backup and Restoration**.
5. In the upper-right corner of the page, click **Restore Individual Database/Table**.

 **Note** If the **Restore Individual Database/Table** button is not displayed, you can check whether all the requirements that are specified in the "Prerequisites" section of this topic are met.

6. Configure the following parameters.

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|--|--|
| Restore To | Current Instance: restores individual databases and tables to the original ApsaraDB RDS instance. |
| Restore Method | <ul style="list-style-type: none"> ◦ By Backup Set ◦ By Time: restores data to a point in time within the specified log backup retention period. <p>Note The By Time option appears only when the log backup feature is enabled.</p> |
| Backup Set | <p>Select the backup set from which you want to restore individual databases and tables.</p> <p>Note This parameter appears only when you set the Restore Method parameter to By Backup Set.</p> |
| Restore Time | <p>Select the point in time to which you want to restore individual databases and tables.</p> <p>Note This parameter appears only when you set the Restore Method parameter to By Time.</p> |
| Databases and Tables to Restore | Select the databases or tables that you want to restore. |
| Selected Databases and Tables | <ul style="list-style-type: none"> ◦ This section displays the selected databases and tables. You can specify new names for these databases and tables. ◦ This section also displays the total size of the selected databases and tables and the remaining storage space. Make sure that the remaining storage space is sufficient before the restoration. |

7. Click **OK**.

FAQ

- After the backup file format is changed from TAR to xstream, are the existing backup files in the TAR format still available?

Yes, the original backup files in the TAR format are still available.

- Why does the Restore Individual Database/Table feature suddenly become unavailable?

Check whether the number of tables on your ApsaraDB RDS instance exceeds 50,000. If the number exceeds 50,000, the Restore Individual Database/Table feature is unavailable.

8.12.4. Download data and log backup files

This topic describes how to download unencrypted data and log backup files in the ApsaraDB RDS console to archive the files and restore data to an on-premises database.

Limits

| Database engine | Download of data backup files | Download of log backup files |
|--|-------------------------------|------------------------------|
| MySQL 5.6 on RDS High-availability Edition (with local SSDs) | Supported | Supported |
| MySQL 5.7 on RDS High-availability Edition (with local SSDs) | Supported | Supported |
| MySQL 8.0 on RDS High-availability Edition (with local SSDs) | Supported | Supported |

Procedure

1. Log on to the ApsaraDB for RDS console.
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Backup and Restoration**.
5. Click the **Data Backup** or **Log Backup** tab.
 - o To download data backup files, click the **Data Backup** tab.
 - o To download log backup files, click the **Log Backup** tab.
6. Select a time range to which you want to restore the instance.
7. Find the data or log backup file that you want to download, and click **Download** in the **Actions** column.

Note

- o If the Download button is unavailable, see the "**Limits**" section of this topic.
- o If you want to use a data backup file to restore data, select the backup file that is the closest to the time for restoration.
- o If you want to use a log backup file to restore data to an on-premises database, take note of the following items:
 - The instance No. of the log backup file must be the same as that of the data backup file.
 - The start time of the log backup file must be later than the end time of the data backup file that you select and earlier than the point in time to which you want to restore the data of your instance.

8. In the message that appears, click **Download**.

| Download method | Description |
|-----------------|--|
| Download | Use a browser to download the backup file. |

| Download method | Description |
|-------------------|---|
| Copy Internal URL | Copy the internal URL to download the file. If your Elastic Compute Service (ECS) and ApsaraDB RDS instances reside within the same region, you can log on to the ECS instance and use the internal URL to download the file. This method is fast and secure. |
| Copy Public URL | Copy the public URL to download the file. If you want to use other tools to download the file, use the public URL. |

Note If you use a Linux operating system, you can run the following command to download the file:

```
wget -c '<The URL that is used to download the backup file>' -O <The name of the backup file>
```

- The `-c` option enables resumable download.
- The `-O` option saves the downloaded file by using a specified name. We recommend that you use the file name contained in the download URL.
- If the URL contains more than one parameter, enclose the download URL in a pair of single quotation marks (').

```
[root@iZbp... ~]# wget -c 'http://rdslog-hz-...cn-hangzhou.aliyuncs.com/.../mysql-bin.000457' -O mysql-bin.000457
```

8.12.5. Upload binlogs

Context

This topic describes how to upload binlog files to OSS.

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Backup and Restoration** to go to the **Backup and Restoration** page.
5. In the upper-right corner of the page, click **Upload Binlogs**.
6. In the message that appears, click **Confirm**.

8.12.6. Restore data to a new instance (formerly known as cloning an instance)

A cloned instance is a new instance that has the same content as the primary instance, including data and settings. This feature allows you to restore data of the primary instance or create multiple instances that are the same as the primary instance.

Prerequisites

Before you clone an instance, make sure that the following requirements are met:

- The primary instance is in the running state.
- The primary instance does not have an ongoing migration task.

- Data backup and log backup are enabled.
- The primary instance has at least one completed backup set before you clone the instance by backup set.

Context

You can specify a backup set or a point in time within the backup retention period to clone an instance.

Note

- A cloned instance copies only the content of the primary instance, but not the content of read-only instances. The copied data includes database information, account information, and instance settings such as whitelist settings, backup settings, parameter settings, and alert threshold settings.
- The database engine of a cloned instance must be the same as that of the primary instance. Other settings can be different, such as the instance edition, zone, network type, instance type, and storage capacity. If you want to restore the data of a primary instance, we recommend that you select a higher instance type and more storage capacity than those of the primary instance. This can speed up the data restoration process.
- The account type of a cloned instance must be the same as that of the primary instance. The account password of the cloned instance can be changed.

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Backup and Restoration**.
5. In the backup list, find a backup and click **Restore** in the **Actions** column.
6. In the dialog box that appears, select **Restore Database** and click **OK**.
7. On the **Restore Instance** page, configure the following parameters.

| Section | Parameter | Description |
|-------------------------|----------------------|---|
| Region | Region | The region where the ApsaraDB RDS instance resides. |
| Restore Database | Restore Mode | The data restore mode of the primary instance. Valid values: <ul style="list-style-type: none"> ◦ By Time ◦ By Backup Set |
| | Restore Time | The point in time to which you want to restore the database. <div style="background-color: #e0f2f1; padding: 5px; margin-top: 5px;">  Note When Restore Mode is set to By Time, you must specify this parameter. </div> |
| | Backup Set | The backup set for restoration. <div style="background-color: #e0f2f1; padding: 5px; margin-top: 5px;">  Note When Restore Mode is set to By Backup Set, you must specify this parameter. </div> |
| | Instance Name | The name of the cloned instance. |

| Section | Parameter | Description |
|----------------|------------------|--|
| Specifications | Database Engine | The engine of the database, which cannot be modified. |
| | Engine Version | The version of the database engine, which cannot be modified. |
| | Edition | The edition of the database. The actual values are displayed in the console. |
| | Storage Type | The storage type of the database. The actual values are displayed in the console. |
| | Instance Type | The type of the cloned instance. Note We recommend that you select a higher instance type and more storage capacity than those of the primary instance. This can speed up the data restoration process. |
| Network Type | Storage Capacity | The storage capacity of the instance, which includes the space to store data, system files, binlog files, and transaction files. The available storage capacity is displayed in the console. Note ApsaraDB RDS instances with local SSDs in the dedicated instance family occupy exclusive resources. The storage capacities are determined based on instance types. |
| | Network Type | The network type of the instance. ApsaraDB RDS instances support the following network types: <ul style="list-style-type: none"> ◦ Classic Network: Cloud services in the classic network are not isolated. Unauthorized access to a cloud service is blocked only by the security group or whitelist policy of the service. ◦ VPC: A virtual private cloud (VPC) helps you build an isolated network environment on Apsara Stack. You can customize route tables, CIDR blocks, and gateways in a VPC. We recommend that you select VPC for improved security. |
| | VPC | Select a VPC. Note When Network Type is set to VPC , you must specify this parameter. |
| | vSwitch | Select a vSwitch. Note When Network Type is set to VPC , you must specify this parameter. |

8. Click **Submit**.

8.13. CloudDBA

8.13.1. Introduction to CloudDBA

CloudDBA is a cloud service for database self-detection, self-repair, self-optimization, self-maintenance, and self-security ensuring based on machine learning and expert experience. CloudDBA helps you ensure stable, secure, and efficient databases without worrying about the management complexity and services failures caused by manual operations.

Features

In ApsaraDB RDS for MySQL, CloudDBA provides the following features:

- **Diagnostics**
You can diagnose your instance and view the visualized diagnostic results.
- **Instance sessions**
You can view sessions, collect session statistics, analyze SQL statements, and optimize the execution of SQL statements.
- **Real-time monitoring**
You can view the real-time monitoring information of your instance, such as the queries per second (QPS), transactions per second (TPS), number of connections, and network traffic.
- **Storage analysis**
You can view the space utilization, trends, exceptions, tablespaces, and data spaces.
- **Deadlock analysis**
You can view and analyze the last deadlock in a database.
- **Dashboard**
You can view and compare performance trends, customize monitoring dashboards, check exceptions, and view instance topologies.
- **Slow query logs**
You can view the trends and statistics of slow queries.
- **Diagnostic reports**
You can use this feature to generate diagnostics reports or view automatically generated reports about instance health, alerts, and slow query logs.

8.13.2. Diagnostics

In ApsaraDB RDS for MySQL, CloudDBA provides the diagnostics feature. This feature diagnoses your ApsaraDB RDS for MySQL instance and visualizes the results.

Go to the Diagnostics page

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, choose **CloudDBA > Diagnostics**.
5. Click the **Diagnostics** tab.

 **Note** For more information, see Diagnostics in *Database Autonomy Service User Guide*.

8.13.3. Session management

In ApsaraDB RDS for MySQL, CloudDBA provides the session management feature, which allows you to view and manage sessions of an instance.

Navigate to the Session Management page

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, choose **CloudDBA > Diagnostics**.
5. Click the **Session Management** tab.

 **Note** For more information, see Instance sessions in *Database Autonomy Service User Guide*.

8.13.4. Real-time monitoring

In ApsaraDB RDS for MySQL, CloudDBA provides the real-time monitoring feature. This feature allows you to view the real-time performance of your ApsaraDB RDS for MySQL instance.

Go to the Real-time Monitoring page

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, choose **CloudDBA > Diagnostics**.
5. Click the **Real-time Monitoring** tab.

 **Note** For more information, see Real-time monitoring in *Database Autonomy Service User Guide*.

8.13.5. Storage analysis

In ApsaraDB RDS for MySQL, CloudDBA provides the storage analysis feature. This feature allows you to check and solve storage exceptions in a timely manner to ensure database stability.

Context

You can use the storage analysis feature of CloudDBA to view the disk space usage of your ApsaraDB RDS for MySQL instance and the number of remaining days when disk space is available. It also provides information about the space usage, fragmentation, and exception diagnostic results of a table.

Go to the Storage Analysis page

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, choose **CloudDBA > Diagnostics**.
5. Click the **Storage Analysis** tab.

 **Note** For more information, see Storage analysis in *Database Autonomy Service User Guide*.

8.13.6. Deadlock analysis

In ApsaraDB RDS for MySQL, CloudDBA provides the deadlock analysis feature. This feature allows you to view and analyze the last deadlock in a database.

Open the Deadlock Analysis page

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, choose **CloudDBA > Diagnostics**.
5. Click the **Deadlock Analysis** tab.

 **Note** For more information, see Deadlock analysis in *DAS User Guide*.

8.13.7. Dashboard

In ApsaraDB RDS for MySQL, CloudDBA provides the dashboard feature. This feature allows you to view performance trends in specific ranges, compare performance trends, and customize charts to view performance trends.

Go to the Dashboard page

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, choose **CloudDBA > Dashboard**.

 **Note** For more information, see Performance trends in *Database Autonomy Service User Guide*.

8.13.8. Slow query logs

In ApsaraDB RDS for MySQL, CloudDBA provides the slow query logs feature. This feature allows you to view the trends and execution details of slow queries and obtain optimization suggestions for your ApsaraDB RDS for MySQL instance.

Go to the Slow Query Logs page

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, choose **CloudDBA > Slow Query Logs**.

 **Note** For more information, see Slow query logs in *Database Autonomy Service User Guide*.

8.13.9. Diagnostic reports

In ApsaraDB RDS for MySQL, CloudDBA provides the diagnostic reports feature. This feature allows you to create and view diagnostic reports.

Open the Report page

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, choose **CloudDBA > Diagnostic Reports**.

 **Note** For more information, see [View diagnostic reports](#) in *Database Autonomy Service User Guide*.

8.14. Logs

All ApsaraDB RDS instances support log management. You can query details about the error logs and slow query logs of an ApsaraDB RDS instance by using the ApsaraDB RDS console. The logs help you locate faults.

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Logs**.
5. On the **Logs** page, click the **Error Logs**, **Slow Query Logs**, **Slow Query Log Summary**, or **Primary/Secondary Switching Logs** tab, select a time range, and then click **Search**.

| Log type | Description |
|----------------------------------|---|
| Error Logs | Records database running errors that occurred within the last month. |
| Slow Log Details | Records SQL statements within the last month that took longer than one second to execute. Duplicated SQL statements are removed.  Note Slow query logs in the ApsaraDB RDS console are updated once every minute. However, you can query real-time slow query logs from the <code>mysql.slow_log</code> table. |
| Slow Query Log Summary | Records and analyzes SQL statements within the last month that took longer than one second to execute. Analysis reports of slow query logs are provided. |
| Primary/Secondary Switching Logs | Records the primary/secondary instance switching logs. This feature is applicable to ApsaraDB RDS for MySQL instances on High-availability Edition. |

8.15. Use mysqldump to migrate MySQL data

This topic describes how to use mysqldump to migrate data from an on-premises database to an ApsaraDB RDS for MySQL instance.

Prerequisites

An ECS instance is created.

Context

mysqldump is easy to use but requires extensive downtime. This tool is suitable for scenarios where the amount of data is small or extensive downtime is allowed.

ApsaraDB RDS for MySQL is fully compatible with the native database service. The procedure of migrating data from the original database to an ApsaraDB RDS for MySQL instance is similar to that of migrating data from one MySQL server to another.

Before you migrate data, you must create an account that is used to migrate data from the on-premises MySQL database. You must grant the read and write permissions on the on-premises MySQL databases to the account.

Procedure

1. Run the following command to create a migration account for the on-premises database:

```
CREATE USER 'username'@'host' IDENTIFIED BY 'password';
```

Parameter description:

- username: the name of the account to be created.
- host: the host from which the account is authorized to log on to the on-premises MySQL database. If you want to allow access from a local host, set this parameter to localhost. If you want to allow access from all hosts, set this parameter to a percent sign (%).
- password: the password of the account.

For example, you can run the following command to create an account with the username William and the password Changme123. The account is authorized to log on to the on-premises MySQL database from all hosts.

```
CREATE USER 'William'@'%' IDENTIFIED BY 'Changme123';
```

2. Run the following command to grant permissions to the migration account in the on-premises database:

```
GRANT SELECT ON databasename.tablename TO 'username'@'host' WITH GRANT OPTION; GRANT REPLICATION SLAVE ON databasename.tablename TO 'username'@'host' WITH GRANT OPTION; GRANT REPLICATION SLAVE ON databasename.tablename TO 'username'@'host' WITH GRANT OPTION;
```

Parameter description:

- privileges: the operation permissions granted to the account, such as SELECT, INSERT, and UPDATE. To authorize the account to perform all operations, set this parameter to ALL.
- databasename: the name of the on-premises MySQL database. If you want to grant all database permissions to the account, set this parameter to an asterisk (*).
- tablename: the name of the table whose data you want to migrate. If you want to grant all table permissions to the account, set this parameter to an asterisk (*).
- username: the name of the account.
- host: the host from which the account is authorized to log on to the on-premises MySQL database. If you want to allow access from a local host, set this parameter to localhost. If you want to allow access from all hosts, set this parameter to a percent sign (%).
- WITH GRANT OPTION: authorizes the account to use the GRANT statement. This parameter is optional.

For example, you can execute the following statement to grant all permissions on tables and databases to the William account. The account is authorized to log on to the database from all hosts.

```
GRANT ALL ON *.* TO 'William'@'%';
```

3. Use the data export tool of mysqldump to export data from the database as a data file.

 **Notice** Do not update data during data export. In this step, only data is exported. Stored procedures, triggers, and functions are not exported.

```
mysqldump -h localhost -u userName -p --opt --default-character-set=utf8 --hex-blob dbName --skip-triggers > /tmp/dbName.sql
```

Parameter description:

- localhost: the IP address of the host where the on-premises MySQL database resides.
- userName: the account that is used to migrate data from the on-premises MySQL database.
- dbName: the name of the on-premises MySQL database.
- /tmp/dbName.sql: the name of the exported data file.

4. Use mysqldump to export stored procedures, triggers, and functions.

 **Notice** Skip this step if no stored procedures, triggers, or functions are used in the database. When stored procedures, triggers, and functions are exported, you must remove the DEFINER to ensure compatibility with ApsaraDB RDS for MySQL.

```
mysqldump -h localhost -u userName -p --opt --default-character-set=utf8 --hex-blob dbName -R | sed -e 's/DEFINER[ ]*=[ ]*[*]*\*/' > /tmp/triggerProcedure.sql
```

Parameter description:

- localhost: the IP address of the host where the on-premises MySQL database resides.
- userName: the account that is used to migrate data from the on-premises MySQL database.
- dbName: the name of the on-premises MySQL database.
- /tmp/triggerProcedure.sql: the name of the exported stored procedure file.

5. Upload the data file and stored procedure file to the ECS instance.

In this example, the files are uploaded to the following paths:

```
/tmp/dbName.sql
```

```
/tmp/triggerProcedure.sql
```

6. Log on to the ECS console and import both the data file and the stored procedure file to the destination ApsaraDB RDS for MySQL instance.

 **Note** For information about how to log on to the ECS instance, see topics in the [Connect to an instance](#) section of ECS User Guide.

```
mysql -h intranet4example.mysql.rds.aliyuncs.com -u userName -p dbName < /tmp/dbName.sql
```

```
mysql -h intranet4example.mysql.rds.aliyuncs.com -u userName -p dbName < /tmp/triggerProcedure.sql
```

Parameter description:

- intranet4example.mysql.rds.aliyuncs.com: the endpoint of the ApsaraDB RDS for MySQL instance. An internal endpoint is used in this example.
- userName: the migration account of the ApsaraDB RDS for MySQL database.
- dbName: the name of the on-premises MySQL database from which you want to import data.
- /tmp/dbName.sql: the name of the data file that you want to import.
- /tmp/triggerProcedure.sql: the name of the stored procedure file that you want to import.

9. ApsaraDB RDS for SQL Server

9.1. What is ApsaraDB RDS?

ApsaraDB Relational Database Service (RDS) is a stable, reliable, and scalable online database service. Based on the distributed file system and high-performance storage, ApsaraDB RDS provides a set of solutions for disaster recovery, backup, restoration, monitoring, and migration.

ApsaraDB RDS supports four database engines, which are MySQL, SQL Server, PolarDB, and PostgreSQL. You can create database instances based on these engines to meet your business requirements. This topic describes the SQL Server engine.

ApsaraDB RDS for SQL Server

ApsaraDB RDS for SQL Server provides strong support for a variety of enterprise applications under the high-availability architecture. ApsaraDB RDS for SQL Server can also restore data to a specific point in time, which reduces costs.

ApsaraDB RDS for SQL Server provides basic features such as whitelist configuration, backup and restoration, transparent data encryption, data migration, and management for instances, accounts, and databases.

9.2. Log on to the ApsaraDB RDS console

This topic describes how to log on to the ApsaraDB RDS console.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

 **Note** When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top navigation bar, choose **Products > Database Services > ApsaraDB RDS**.

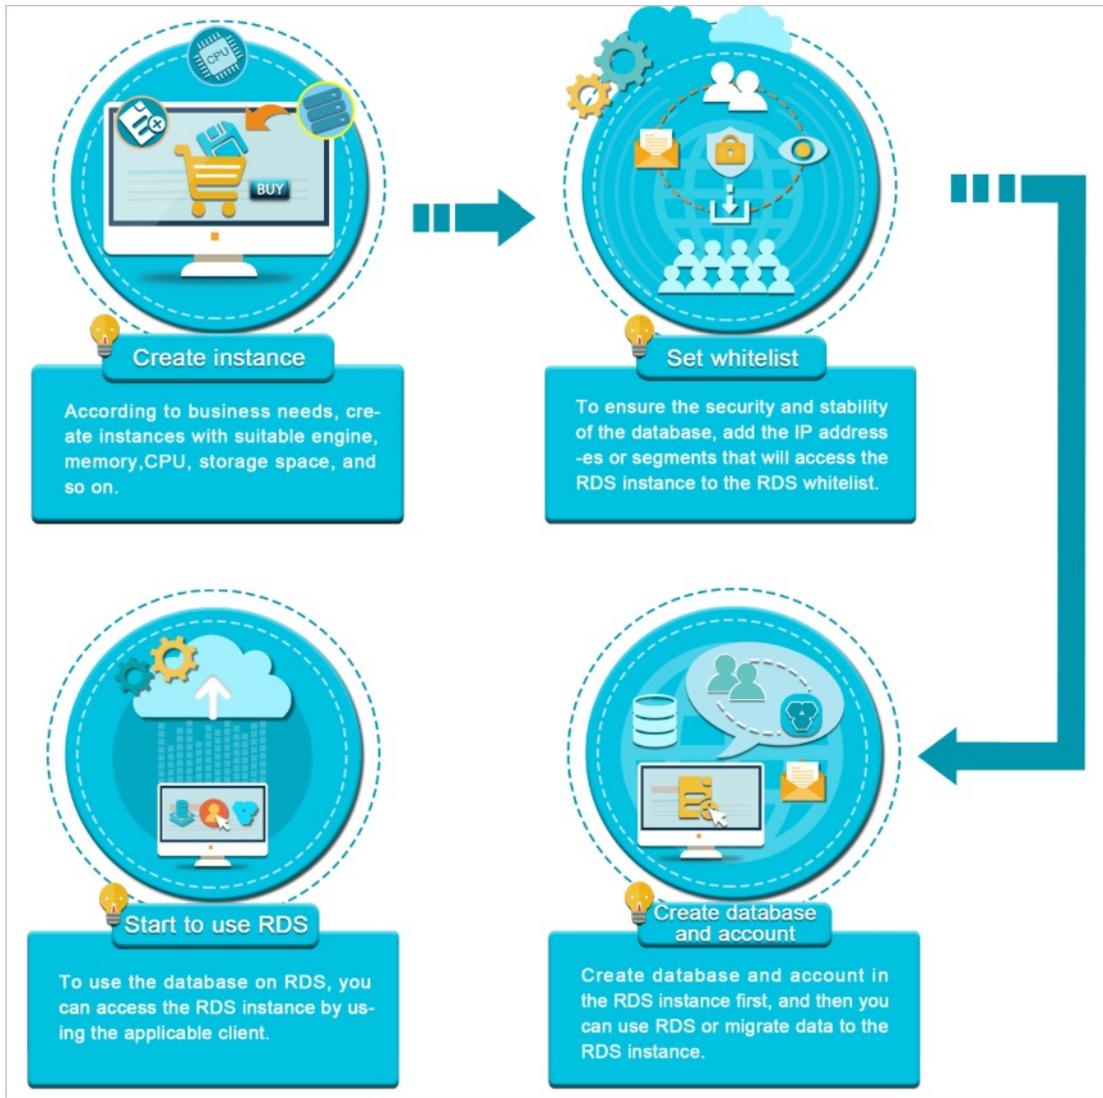
9.3. Quick Start

9.3.1. Procedure

ApsaraDB RDS quick start covers the following topics: creating an ApsaraDB RDS instance, configuring an IP address whitelist, creating a database, creating an account, and connecting to the instance.

The following figure shows the operations that you must perform before you use an ApsaraDB RDS instance.

Quick start flowchart



9.3.2. Create an instance

This topic describes how to create an instance in the ApsaraDB RDS console.

Prerequisites

An Apsara Stack tenant account is created.

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, click **Create Instance** in the upper-right corner.
3. Configure the following parameters.

| Section | Parameter | Description |
|----------------|------------------------|--|
| Basic Settings | Organization | The organization to which the instance belongs. |
| | Resource Set | The resource set to which the instance belongs. |
| Region | Region | The region in which you want to create the instance. Services in different regions cannot communicate over an internal network. After the instance is created, the region cannot be changed. |
| | Zone of Primary Node | The zone where the primary instance is deployed. |
| | Deployment Method | Specifies whether to deploy the primary and secondary instances in separate zones. ApsaraDB RDS supports Multi-zone Deployment and Single-zone Deployment . If you select Multi-zone Deployment , you must configure Zone of Secondary Node . |
| | Zone of Secondary Node | The zone where the secondary instance is deployed. This parameter is available only when Deployment Method is set to Multi-zone Deployment . <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e6f2ff;"> <p> Note If you select the same zone for both the primary and secondary instances, the deployment is equivalent to single-zone deployment.</p> </div> |
| | Quantity | The number of ApsaraDB RDS instances that you want to create. Default value: 1. |
| | Instance Name | The name of the instance. <ul style="list-style-type: none"> ◦ The name must be 2 to 64 characters in length. ◦ The name must start with a letter. ◦ The name can contain letters, digits, and the following special characters: _ - : ◦ The name cannot start with http:// or https://. |
| | Connection Type | The connection type of the instance. ApsaraDB RDS instances support the following connection types: <ul style="list-style-type: none"> ◦ Internet: ApsaraDB RDS instances of this connection type can be connected over the Internet. ◦ Internal Network: ApsaraDB RDS instances of this connection type can be connected over an internal network. <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e6f2ff;"> <p> Note The value of this parameter cannot be changed after the instance is created. Proceed with caution.</p> </div> |
| | Database Engine | The database engine of the instance. Select SQLServer . |

| Specification Section | Parameter | Description |
|-----------------------|------------------|--|
| | Engine Version | The version of the database engine. Valid values: <ul style="list-style-type: none"> ◦ 2012_ent_ha: SQL Server 2012 EE ◦ 2012_std_ha: SQL Server 2012 SE ◦ 2016_ent_ha: SQL Server 2016 EE ◦ 2016_std_ha: SQL Server 2016 SE ◦ 2017_ent_ha: SQL Server 2017 EE |
| | Edition | The edition of the instance. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| | Storage Type | The storage type of the instance. The storage type is automatically set to cloud ssd . |
| | Encrypted | Specifies whether to encrypt the standard SSD. This parameter is available only when Storage Type is set to cloud ssd . If you select Encrypted, you must specify the Encryption Key parameter. If you do not have a key, you must first create one in Key Management Service (KMS). For more information, see Create a CMK in <i>KMS User Guide</i> . |
| | Encryption Key | The key that is used to encrypt the standard SSD. This parameter is available only when you select Encrypted . |
| | Instance Type | The instance type of the instance. Memory size determines the maximum number of connections and the input/output operations per second (IOPS). The actual values are displayed in the console. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| | Storage Capacity | The storage capacity of the instance, which includes the space to store data, system files, binlog files, and transaction files. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| Network | Network Type | The network type of the instance. ApsaraDB RDS instances support the following network types: <ul style="list-style-type: none"> ◦ Classic Network: Cloud services in the classic network are not isolated. Unauthorized access to a cloud service is blocked only by the security group or whitelist policy of the service. ◦ VPC: A virtual private cloud (VPC) helps you build an isolated network environment on Apsara Stack. You can customize route tables, CIDR blocks, and gateways in a VPC. We recommend that you select VPC for improved security. |
| | VPC | Select a VPC. <div style="background-color: #e0f2f1; padding: 5px; margin-top: 10px;"> ? Note When Network Type is set to VPC, you must specify this parameter. </div> |
| | vSwitch | Select a vSwitch. <div style="background-color: #e0f2f1; padding: 5px; margin-top: 10px;"> ? Note When Network Type is set to VPC, you must specify this parameter. </div> |

| Section | Parameter | Description |
|---------|-----------------------------|---|
| | IP Address Whitelist | The IP addresses that are allowed to connect to the instance. |

4. Click **Submit**.

9.3.3. Configure an IP address whitelist

To ensure database security and reliability, you must modify the whitelist of an ApsaraDB RDS instance before you enable the instance. You must add the IP addresses or CIDR blocks that are used for database access to the whitelist.

Context

The whitelist improves the access security of your ApsaraDB RDS instance. We recommend that you maintain the whitelist on a regular basis. The whitelist configuration process does not affect the normal operations of the ApsaraDB RDS instance.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. On the **Whitelist Settings** tab, click **Edit** corresponding to the **default** whitelist.

Note

- If you want to connect an ECS instance to an ApsaraDB RDS instance by using an internal endpoint, you must make sure that the two instances are in the same region and have the same network type. Otherwise, the connection fails.
- You can also click **Create Whitelist** to create a new whitelist.

6. In the **Edit Whitelist** dialog box, enter the IP addresses or CIDR blocks that are allowed to access your ApsaraDB RDS instance, and then click **OK**.
 - If you add the CIDR block 10.10.10.0/24, all IP addresses in the 10.10.10.X format are allowed to access the ApsaraDB RDS instance.
 - If you enter more than one IP address or CIDR block, you must separate them with commas (,). Do not add spaces before or after the commas. Example: 192.168.0.1,172.16.213.9.
 - If you click **Add Internal IP Addresses of ECS Instances**, the IP addresses of all of the ECS instances that are created in your Alibaba Cloud account appear. Then, you can select the required IP addresses and add them to the whitelist.

 **Note** If you add a new IP address or CIDR block to the **default** whitelist, the default address 127.0.0.1 is deleted.

Edit Whitelist

*Whitelist Name: default

*IP Addresses: 127.0.0.1

[Add Internal IP Addresses of ECS Instances](#)
You can add 999 more entries.

Specified IP address: If you specify the IP address 192.168.0.1, this IP address is allowed to access the RDS instance.
Specified CIDR block: If you specify the CIDR block 192.168.0.0/24, the IP addresses ranging from 192.168.0.1 to 192.168.0.255 are allowed to access the RDS instance.
When you add multiple IP addresses or CIDR blocks, separate them by a comma (no space after the comma), for example, 192.168.0.1,192.168.0.0/24.

New whitelist entries take effect in 1 minute.

OK Cancel

9.3.4. Connect to an instance

This topic describes how to use Data Management (DMS) to connect to an ApsaraDB RDS instance.

Prerequisites

- A database is created. For more information, see [Create a database](#).
- A database account is created. For more information, see [Create an account](#).

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. Click **Log On to DB** in the upper-right corner of the page.
5. In the **Login instance** dialog box of the DMS console, check values of **Database type**, **Instance Area**, and **Connection string address**. If the information is correct, enter **Database account** and **Database password**, as shown in the following figure.

| Parameter | Description |
|---------------------------|---|
| Database type | The engine of the database. By default, the engine of the database to be connected is displayed. |
| Instance Area | The region where the instance is deployed. By default, the region of the current instance is displayed. |
| Connection string address | The endpoint of the instance. By default, the endpoint of the current instance is displayed. |
| Database account | The account of the database to be connected. |
| Database password | The password of the account used to connect to the database. |

6. Click **Login**.

Note

- If you want the browser to remember the password, select **Remember password** and click **Login**.
- If you cannot connect to the instance, check the IP address whitelist settings. For more information, see [Configure a whitelist](#).

9.3.5. Create an account

This topic describes how to create an account on an ApsaraDB RDS for SQL Server instance.

Prerequisites

The instance is in the **Running** state.

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Accounts**.
5. On the right side of the page, click **Create Account**.
6. Enter the information of the account that you want to create.

| Parameter | Description |
|-----------------------------|---|
| Database Account | Enter the name of the account. The name must be 2 to 16 characters in length and can contain lowercase letters, digits, and underscores (_). It must start with a lowercase letter and end with a lowercase letter or digit. |
| Account Type | <ul style="list-style-type: none"> ◦ Privileged Account: You can select the Privileged Account option only if you create an account on your ApsaraDB RDS instance for the first time. Each ApsaraDB RDS instance can have only a single privileged account. The privileged account of an ApsaraDB RDS instance cannot be deleted. ◦ Standard Account: You can select the Standard Account option only after a privileged account is created on your ApsaraDB RDS instance. Each ApsaraDB RDS instance can have more than one standard account. You must manually grant permissions on databases to each standard account. |
| Authorized Databases | <p>Select the authorized databases of the account when the Standard Account type is selected. If no databases are created, you can leave this parameter empty.</p> <p>You can perform the following steps to grant permissions on more than one database to the account:</p> <ol style="list-style-type: none"> In the Unauthorized Databases section, select the databases on which you want to grant permissions to the account. Click Add to add the selected databases to the Authorized Databases section. In the Authorized Databases section, specify the permissions that the account is granted on each authorized database. The permissions can be Read/Write, Read-only, or Owner. You can also click Set All to Read/Write, Set All to Read-only, or Set All to Owner to set the permissions of the account on all authorized databases. <div style="background-color: #e1f5fe; padding: 10px; margin-top: 10px;"> <p> Note</p> <ul style="list-style-type: none"> ■ The account is authorized to create tables, delete tables, and modify schemas in a database only when it has the Owner permission on the database. ■ The account has permissions on all databases and does not require authorization if you select the Privileged Account type. </div> |
| Password | <p>Enter the password of the account. The password must meet the following requirements:</p> <ul style="list-style-type: none"> ◦ The password is 8 to 32 characters in length. ◦ The password contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. ◦ Special characters include <code>!@#\$%^&*()_+ -=</code> |
| Re-enter Password | Enter the password of the account again. |

| Parameter | Description |
|-------------|---|
| Description | Enter a description that helps identify the account. The description can be up to 256 characters in length. |

7. Click **Create**.

9.3.6. Create a database

This topic describes how to create a database on an ApsaraDB RDS for SQL Server instance in the ApsaraDB RDS console.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Databases**.
5. In the upper-right corner of the page, click **Create Database**.
6. Configure the parameters for the database that you want to create.

| Parameter | Description |
|--------------------------|--|
| Database Name | Enter the name of the database. The name must be 2 to 64 characters in length. It can contain lowercase letters, digits, underscores (_), and hyphens (-). It must start with a lowercase letter and end with a lowercase letter or digit. |
| Supported Character Sets | Select the character set that is supported by the database. You can also select all and then select a character set from the drop-down list that appears. |
| Description | Enter a description of the database to facilitate subsequent management. The description can be up to 256 characters in length. |

7. Click **Create**.

9.4. Instances

9.4.1. Create an instance

This topic describes how to create an instance in the ApsaraDB RDS console.

Prerequisites

An Apsara Stack tenant account is created.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, click **Create Instance** in the upper-right corner.
3. Configure the following parameters.

| Section | Parameter | Description |
|----------------|------------------------|--|
| Basic Settings | Organization | The organization to which the instance belongs. |
| | Resource Set | The resource set to which the instance belongs. |
| Region | Region | The region in which you want to create the instance. Services in different regions cannot communicate over an internal network. After the instance is created, the region cannot be changed. |
| | Zone of Primary Node | The zone where the primary instance is deployed. |
| | Deployment Method | Specifies whether to deploy the primary and secondary instances in separate zones. ApsaraDB RDS supports Multi-zone Deployment and Single-zone Deployment . If you select Multi-zone Deployment , you must configure Zone of Secondary Node . |
| | Zone of Secondary Node | The zone where the secondary instance is deployed. This parameter is available only when Deployment Method is set to Multi-zone Deployment . <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e6f2ff;"> <p> Note If you select the same zone for both the primary and secondary instances, the deployment is equivalent to single-zone deployment.</p> </div> |
| | Quantity | The number of ApsaraDB RDS instances that you want to create. Default value: 1. |
| | Instance Name | The name of the instance. <ul style="list-style-type: none"> ◦ The name must be 2 to 64 characters in length. ◦ The name must start with a letter. ◦ The name can contain letters, digits, and the following special characters: _ - : ◦ The name cannot start with http:// or https://. |
| | Connection Type | The connection type of the instance. ApsaraDB RDS instances support the following connection types: <ul style="list-style-type: none"> ◦ Internet: ApsaraDB RDS instances of this connection type can be connected over the Internet. ◦ Internal Network: ApsaraDB RDS instances of this connection type can be connected over an internal network. <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e6f2ff;"> <p> Note The value of this parameter cannot be changed after the instance is created. Proceed with caution.</p> </div> |
| | Database Engine | The database engine of the instance. Select SQLServer . |

| Specification Section | Parameter | Description |
|-----------------------|------------------|--|
| | Engine Version | The version of the database engine. Valid values: <ul style="list-style-type: none"> ◦ 2012_ent_ha: SQL Server 2012 EE ◦ 2012_std_ha: SQL Server 2012 SE ◦ 2016_ent_ha: SQL Server 2016 EE ◦ 2016_std_ha: SQL Server 2016 SE ◦ 2017_ent_ha: SQL Server 2017 EE |
| | Edition | The edition of the instance. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| | Storage Type | The storage type of the instance. The storage type is automatically set to cloud ssd . |
| | Encrypted | Specifies whether to encrypt the standard SSD. This parameter is available only when Storage Type is set to cloud ssd . If you select Encrypted, you must specify the Encryption Key parameter. If you do not have a key, you must first create one in Key Management Service (KMS). For more information, see Create a CMK in <i>KMS User Guide</i> . |
| | Encryption Key | The key that is used to encrypt the standard SSD. This parameter is available only when you select Encrypted . |
| | Instance Type | The instance type of the instance. Memory size determines the maximum number of connections and the input/output operations per second (IOPS). The actual values are displayed in the console. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| | Storage Capacity | The storage capacity of the instance, which includes the space to store data, system files, binlog files, and transaction files. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| Network | Network Type | The network type of the instance. ApsaraDB RDS instances support the following network types: <ul style="list-style-type: none"> ◦ Classic Network: Cloud services in the classic network are not isolated. Unauthorized access to a cloud service is blocked only by the security group or whitelist policy of the service. ◦ VPC: A virtual private cloud (VPC) helps you build an isolated network environment on Apsara Stack. You can customize route tables, CIDR blocks, and gateways in a VPC. We recommend that you select VPC for improved security. |
| | VPC | Select a VPC. <div style="background-color: #e0f2f7; padding: 5px; margin-top: 10px;"> ? Note When Network Type is set to VPC, you must specify this parameter. </div> |
| | vSwitch | Select a vSwitch. <div style="background-color: #e0f2f7; padding: 5px; margin-top: 10px;"> ? Note When Network Type is set to VPC, you must specify this parameter. </div> |

| Section | Parameter | Description |
|---------|----------------------|---|
| | IP Address Whitelist | The IP addresses that are allowed to connect to the instance. |

4. Click **Submit**.

9.4.2. View basic information of an instance

This topic describes how to view the details of an ApsaraDB RDS instance, such as its basic information, internal network connection information, status, and configurations.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. Use one of the following methods to go to the **Basic Information** page of an instance:
 - On the **Instances** page, click the ID of an instance to go to the **Basic Information** page.
 - On the **Instances** page, click **Manage** in the **Actions** column corresponding to an instance to go to the **Basic Information** page.

9.4.3. Restart an instance

This topic describes how to manually restart an ApsaraDB RDS for MySQL instance. This applies if the number of connections exceeds a specific threshold or if an instance has performance issues.

Prerequisites

The instance is in the **Running** state.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. Click **Restart Instance** in the upper-right corner.

 **Note** When you restart an instance, applications are disconnected from the instance. We recommend that you make appropriate service arrangements before you restart an instance. Proceed with caution.

5. In the **Restart Instance** message, click **Confirm**.

9.4.4. Change the specifications of an instance

This topic describes how to change specifications such as the instance type and storage space if they do not meet the requirements of your application. When the specification changes take effect, a 30-second network interruption may occur. Business operations that involve databases, accounts, and networks are interrupted. We recommend that you change the specifications during off-peak hours or make sure that your applications are configured with automatic reconnection policies.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.

3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the **Configuration Information** section, click **Change Specifications**.
5. On the **Change Specifications** page, specify **Instance Type** and **Storage Capacity**.
6. After you configure the preceding parameters, click **Submit**.

9.4.5. Set a maintenance window

This topic describes how to set the maintenance window of an ApsaraDB RDS for SQL Server instance. The backend system performs maintenance on the ApsaraDB RDS instance during the maintenance window. This ensures the stability of the ApsaraDB RDS instance. The default maintenance window is from 02:00 (UTC+8) to 06:00 (UTC+8). We recommend that you set the maintenance window to off-peak hours of your business to avoid impacts on your business.

Context

- An instance enters the **Maintaining Instance** state before the maintenance window to ensure stability during the maintenance process. When the instance is in this state, access to data in the database and query operations such as performance monitoring are not affected. However, except for account and database management and IP address whitelist configuration, modification operations such as upgrade, downgrade, and restart are temporarily unavailable.
- During the maintenance window, one or two network interruptions may occur. Make sure that your applications are configured with automatic reconnection policies.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the **Configuration Information** section, click **Configure** to the right of **Maintenance Window**.
5. Select a maintenance window and click **Save**.

 **Note** The maintenance window is displayed in UTC+8.

9.4.6. Configure primary/secondary switchover

ApsaraDB RDS provides the primary/secondary switchover feature to ensure the high availability of databases. The primary/secondary switchover is performed when the primary instance becomes unavailable. You can also manually switch your business to the secondary instance.

Prerequisites

The instance is in the **Running** state.

Context

An ApsaraDB RDS for SQL Server instance has a secondary instance. Data is synchronized in real time between the primary and secondary instances. You can access only the primary instance. The secondary instance serves only as a backup instance and does not allow external access. If the primary instance cannot be accessed, your business automatically switches over to the secondary instance. After the switchover, the primary instance becomes the secondary instance.

 Notice

- During a switchover, a network interruption may occur. Make sure that your applications are configured with automatic reconnection policies.
- During a switchover, a 1-minute data quality protection mechanism is enabled for data synchronization. If the primary and secondary database states are incorrect or if the latency for data synchronization exceeds 1 minute due to SQL Server errors, the HA system does not automatically perform the primary/secondary switchover. You must determine whether to perform the switchover.
- If an instance is intermittently unavailable due to excessive mirroring event waits, the switchover is not performed. The instance automatically becomes available again.

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Service Availability**.
5. In the **Availability Information** section, click **Switch Primary/Secondary Instance**.
6. In the dialog box that appears, click **OK**.

Result

After the switchover is complete, the original primary instance becomes the secondary instance for the next primary/secondary switchover.

9.4.7. Release an instance

This topic describes how to manually release an instance.

Context

- Only instances in the running state can be manually released.
- After an instance is released, the instance data is immediately deleted. We recommend that you back up your data before you release an instance.

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. Find the instance that you want to release and choose **More > Release Instance**.
3. In the **Release Instance** message, click **Confirm**.

9.4.8. Read-only instances

9.4.8.1. Overview of read-only ApsaraDB RDS for SQL Server instances

This topic provides an overview of read-only ApsaraDB RDS for SQL Server instances. If a large number of read requests overwhelm the primary instance, your business may be interrupted. In this case, you can create one or more read-only instances to offload read requests from the primary instance. This scales the read capability of your database system and increases the throughput of your application.

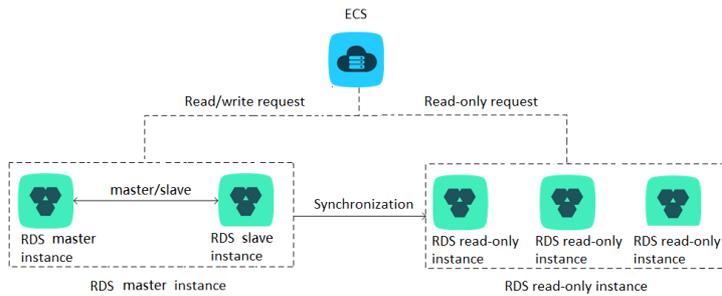
Overview

When a read-only instance is created, the data is replicated from the secondary instance. The data is consistent with that of the primary instance. Data updates of the primary instance are synchronized to all read-only instances.

Note

- Only ApsaraDB RDS instances that run SQL Server 2017 EE support read-only instances.
- Each read-only instance works in a single-node architecture, where no instances are provided as backups.

The following figure shows the topology of read-only instances.



Features

- The specifications of a read-only instance can differ from the specifications of the primary instance, and can be changed at any time. We recommend that you select specifications of a read-only instance that are higher than or equal to those of the primary instance. If the specifications of a read-only instance are lower than those of the primary instance, the read-only instance may have high latency or workloads.
- Read-only instances do not require database or account maintenance, because their database and account information is synchronized with the primary instance.
- A read-only instance automatically replicates the IP address whitelists of the primary instance. However, the IP address whitelists for the read-only instance are independent of those of the primary instance. For information about how to modify the whitelists of a read-only instance, see [Configure a whitelist](#).
- You can monitor up to 20 system performance metrics, such as the disk capacity, input/output operations per second (IOPS), number of connections, CPU utilization, and network traffic.

Limits

- You can create up to seven read-only instances.
- You cannot configure backup policies or manually create backups for read-only instances, because these are already configured or created on the primary instance.
- You cannot create a temporary instance by using a backup set or from a point in time. In addition, you cannot overwrite a read-only instance by using a backup set.
- After a read-only instance is created, you cannot use a data backup file to restore it in overwrite mode.
- You cannot migrate data to read-only RDS instances.
- You cannot create or delete databases on read-only instances.
- You cannot create or delete accounts, authorize accounts, or change the passwords of accounts on read-only instances.

FAQ

Can I manage the accounts created on the primary instance from its read-only instances?

No, although accounts created on the primary instance are replicated to its read-only instances, you cannot manage the accounts on the read-only instances. The accounts have only read permissions on the read-only instances.

9.4.8.2. Create a read-only ApsaraDB RDS for SQL Server instance

This topic describes how to create a read-only instance for your primary ApsaraDB RDS for SQL server instance. This allows your database system to process a large number of read requests and increases the throughput of your application. Each read-only ApsaraDB RDS instance is a replica of the primary instance. Data updates on the primary instance are synchronized to all the read-only instances.

Prerequisites

The primary instance runs SQL Server 2017 EE.

Precautions

- You can create read-only instances for the primary ApsaraDB RDS instance. However, you cannot convert existing ApsaraDB RDS instances into read-only instances.
- While you create a read-only instance, the system replicates data from a secondary instance. Therefore, the operation of your primary instance is not interrupted.
- You can create up to seven read-only instances.
- For more information about read-only ApsaraDB RDS instances, see [Overview of read-only ApsaraDB RDS for SQL Server instances](#).

Create a read-only instance

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the **Distributed by Instance Role** section on the right side of the page, click **Create Read-only Instance**.
5. Configure the following parameters and click **Submit**.

| Section | Parameter | Description |
|---------------|------------------------|---|
| Region | Region | The region in which you want to create the instance. |
| | Database Engine | The database engine of the read-only instance, which is the same as that of the primary instance and cannot be changed. |
| | Engine Version | The engine version of the read-only instance, which is the same as that of the primary instance and cannot be changed. |
| | Edition | Set the value to Read-only . |

| Section | Parameter | Description |
|----------------|------------------|--|
| Specifications | Instance Type | <p>The instance type of the read-only instance. The instance type of the read-only instance can be different from that of the primary instance, and can be changed at any time to facilitate flexible upgrade and downgrade. For more information, see Instance types in <i>ApsaraDB RDS Product Information</i>.</p> <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p>Note To ensure sufficient I/O throughput for data synchronization, we recommend that you select at least the same instance type as the primary instance for read-only instances.</p> </div> |
| | Storage Capacity | <p>The storage space of the read-only instance. To ensure sufficient I/O throughput for data synchronization, we recommend that you select at least the same instance type and storage space as the primary instance for the read-only instance. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i>.</p> |
| Network Type | Network Type | <p>The network type of the read-only instance, which is the same as that of the primary instance and cannot be changed.</p> |
| | VPC | <p>Select a VPC if the network type is set to VPC.</p> |
| | vSwitch | <p>Select a vSwitch if the network type is set to VPC.</p> |

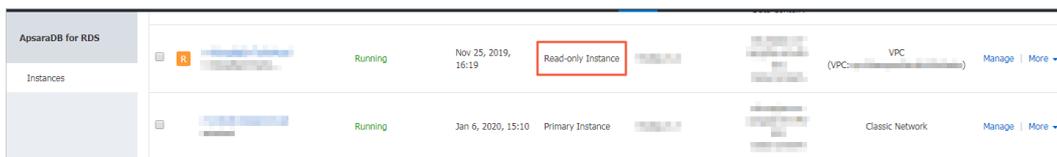
9.4.8.3. View details of read-only instances

This topic describes how to view details of read-only instances. You can go to the Basic Information page of a read-only instance from the Instances page or the read-only instance list of the primary instance. Read-only instances are managed in the same way as primary instances. The read-only instance management page shows the management operations that can be performed.

View instance details by using a read-only instance

1. Log on to the [ApsaraDB for RDS console](#).
2. On the **Instances** page, click the ID of a read-only instance. The **Basic Information** page appears. In the instance list, Instance Role of read-only instances is displayed as Read-only Instance, as shown in [View a read-only instance](#).

View a read-only instance



View instance details by using the primary instance

1. Log on to the [ApsaraDB for RDS console](#).

2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. On the **Basic Information** page, move the pointer over the number below **Read-only Instance** in the **Distributed by Instance Role** section. The ID of the read-only instance is displayed.
5. Click the ID of the read-only instance to go to the read-only instance management page.

9.5. Accounts

9.5.1. Create an account

This topic describes how to create an account on an ApsaraDB RDS for SQL Server instance.

Prerequisites

The instance is in the **Running** state.

Procedure

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Accounts**.
5. On the right side of the page, click **Create Account**.
6. Enter the information of the account that you want to create.

| Parameter | Description |
|-------------------------|---|
| Database Account | Enter the name of the account. The name must be 2 to 16 characters in length and can contain lowercase letters, digits, and underscores (_). It must start with a lowercase letter and end with a lowercase letter or digit. |
| Account Type | <ul style="list-style-type: none"> ◦ Privileged Account: You can select the Privileged Account option only if you create an account on your ApsaraDB RDS instance for the first time. Each ApsaraDB RDS instance can have only a single privileged account. The privileged account of an ApsaraDB RDS instance cannot be deleted. ◦ Standard Account: You can select the Standard Account option only after a privileged account is created on your ApsaraDB RDS instance. Each ApsaraDB RDS instance can have more than one standard account. You must manually grant permissions on databases to each standard account. |

| Parameter | Description |
|----------------------|---|
| Authorized Databases | <p>Select the authorized databases of the account when the Standard Account type is selected. If no databases are created, you can leave this parameter empty.</p> <p>You can perform the following steps to grant permissions on more than one database to the account:</p> <ol style="list-style-type: none"> In the Unauthorized Databases section, select the databases on which you want to grant permissions to the account. Click Add to add the selected databases to the Authorized Databases section. In the Authorized Databases section, specify the permissions that the account is granted on each authorized database. The permissions can be Read/Write, Read-only, or Owner. You can also click Set All to Read/Write, Set All to Read-only, or Set All to Owner to set the permissions of the account on all authorized databases. <div style="background-color: #e1f5fe; padding: 10px; margin-top: 10px;"> <p> Note</p> <ul style="list-style-type: none"> ■ The account is authorized to create tables, delete tables, and modify schemas in a database only when it has the Owner permission on the database. ■ The account has permissions on all databases and does not require authorization if you select the Privileged Account type. </div> |
| Password | <p>Enter the password of the account. The password must meet the following requirements:</p> <ul style="list-style-type: none"> ○ The password is 8 to 32 characters in length. ○ The password contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. ○ Special characters include <code>!@#\$%^&*()_+ -=</code> |
| Re-enter Password | Enter the password of the account again. |
| Description | Enter a description that helps identify the account. The description can be up to 256 characters in length. |

7. Click **Create**.

9.5.2. Reset the password

You can use the ApsaraDB RDS console to reset the password of your database account.

Prerequisites

The instance is in the **Running** state.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.

4. In the left-side navigation pane, click **Accounts**.
5. Find an account and click **Reset Password** in the Actions column.
6. In the dialog box that appears, enter and confirm the new password, and then click **OK**.

- Note** The password must meet the following requirements:
- The password must be 8 to 32 characters in length.
 - The password must contain at least three of the following characters: uppercase letters, lowercase letters, digits, and special characters.
 - Special characters include ! @ # \$ % ^ & * () _ + - =

9.6. Databases

9.6.1. Create a database

This topic describes how to create a database on an ApsaraDB RDS for SQL Server instance.

Terms

- **Instance:** a virtualized database server on which you can create and manage more than one database.
- **Database:** a set of data that is stored in an organized manner and can be shared by a number of users. A database provides the minimal redundancy and is independent of applications. In simple words, a database is a data warehouse that is used to store data.
- **Character set:** a collection of letters, special characters, and encoding rules that are used in a database.

Prerequisites

An ApsaraDB RDS for SQL Server instance is created. For more information, see [Create an instance](#).

Procedure

For more information, see [Create a database](#).

9.6.2. Delete a database

This topic describes how to delete a database from an ApsaraDB RDS for SQL Server instance. You can delete a database by using the ApsaraDB RDS console or an SQL statement.

Use the console to delete a database

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Databases**.
5. Find the database that you want to delete and click **Delete** in the **Actions** column.
6. In the message that appears, click **Confirm**.

Execute an SQL statement to delete a database

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic**

Information page.

- Click **Log On to DB** in the upper-right corner of the page.
- In the **Login instance** dialog box of the **DMS** console, check values of **Database type**, **Instance Area**, and **Connection string address**. If the information is correct, enter **Database account** and **Database password**, as shown in the following figure.

| Parameter | Description |
|----------------------------------|---|
| Database type | The engine of the database. By default, the engine of the database to be connected is displayed. |
| Instance Area | The region where the instance is deployed. By default, the region of the current instance is displayed. |
| Connection string address | The endpoint of the instance. By default, the endpoint of the current instance is displayed. |
| Database account | The account of the database to be connected. |
| Database password | The password of the account used to connect to the database. |

- Click **Login**.

Note

- If you want the browser to remember the password, select **Remember password** and click **Login**.
- If you cannot connect to the instance, check the IP address whitelist settings. For more information, see [Configure a whitelist](#).

- The **SQLConsole** page appears after you log on to the instance. Execute a statement in the following format to delete a database:

```
drop database <database name>;
```

Note If the instance runs SQL Server 2012 or later on RDS High-availability Edition, you can also use the following stored procedure. This stored procedure deletes the specified database, removes the associated image, and closes the connection to the database.

```
EXEC sp_rds_drop_database 'database name'
```

8. Click **execute**.

9.6.3. Change the character set collation and the time zone of system databases

This topic describes how to change the character set collation and the time zone of system databases. System databases include master, msdb, tempdb, and model.

Prerequisites

- The instance runs SQL Server 2012, 2016, or 2017.
- No database other than system databases exists on the instance.

Note If you have just deleted databases from the instance, the deletion task may be pending in the secondary instance. Before you change the character set collation and the time zone, make sure that the primary and secondary instances do not contain databases.

Precautions

- The default character set collation is Chinese_PRC_CI_AS.
- The default time zone is China Standard Time.
- You can view the available character set collations and time zones in the console.
- The instance is in the unavailable state during the change process. It takes about 1 minute to change the time zone, and 2 to 10 minutes to change the character set collation.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Databases**.
5. On the Databases page, click **Change Character Set Collation and Time Zone**.

Note If you fail to find this button on the page, make sure that the requirements in [Prerequisites](#) are met.

6. In the dialog box that appears, select **Time Zone**, **Character Set Collation**, or both of them, and click **OK**.

UTC offsets of time zones

| Time zone | UTC offset | Description |
|---------------------------|-------------|-------------|
| Afghanistan Standard Time | (UTC+04:30) | Kabul |

| Time zone | UTC offset | Description |
|---------------------------------|-------------|---|
| Alaskan Standard Time | (UTC-09:00) | Alaska |
| Arabian Standard Time | (UTC+04:00) | Abu Dhabi, Muscat |
| Atlantic Standard Time | (UTC-04:00) | Atlantic Time (Canada) |
| AUS Central Standard Time | (UTC+09:30) | Darwin |
| AUS Eastern Standard Time | (UTC+10:00) | Canberra, Melbourne, Sydney |
| Belarus Standard Time | (UTC+03:00) | Minsk |
| Canada Central Standard Time | (UTC-06:00) | Saskatchewan |
| Cape Verde Standard Time | (UTC-01:00) | Cabo Verde Is. |
| Gen. Australia Standard Time | (UTC+09:30) | Adelaide |
| Central America Standard Time | (UTC-06:00) | Central America |
| Central Asia Standard Time | (UTC+06:00) | Astana |
| Central Brazilian Standard Time | (UTC-04:00) | Cuiaba |
| Central Europe Standard Time | (UTC+01:00) | Belgrade, Bratislava, Budapest, Ljubljana, Prague |
| Central European Standard Time | (UTC+01:00) | Sarajevo, Skopje, Warsaw, Zagreb |
| Central Pacific Standard Time | (UTC+11:00) | Solomon Islands, New Caledonia |
| Central Standard Time | (UTC-06:00) | Central Time (US and Canada) |
| Central Standard Time (Mexico) | (UTC-06:00) | Guadalajara, Mexico City, Monterrey |
| China Standard Time | (UTC+08:00) | Beijing, Chongqing, Hong Kong, Urumqi |
| E. Africa Standard Time | (UTC+03:00) | Nairobi |
| E. Australia Standard Time | (UTC+10:00) | Brisbane |
| E. Europe Standard Time | (UTC+02:00) | Chisinau |
| E. South America Standard Time | (UTC-03:00) | Brasilia |
| Eastern Standard Time | (UTC-05:00) | Eastern Time (US and Canada) |
| Georgian Standard Time | (UTC+04:00) | Tbilisi |
| GMT Standard Time | (UTC) | Dublin, Edinburgh, Lisbon, London |
| Greenland Standard Time | (UTC-03:00) | Greenland |
| Greenwich Standard Time | (UTC) | Monrovia, Reykjavik |
| GTB Standard Time | (UTC+02:00) | Athens, Bucharest |

| Time zone | UTC offset | Description |
|---------------------------------|-------------|--|
| Hawaiian Standard Time | (UTC-10:00) | Hawaii |
| India Standard Time | (UTC+05:30) | Chennai, Kolkata, Mumbai, New Delhi |
| Jordan Standard Time | (UTC+02:00) | Amman |
| Korea Standard Time | (UTC+09:00) | Seoul |
| Middle East Standard Time | (UTC+02:00) | Beirut |
| Mountain Standard Time | (UTC-07:00) | Mountain Time (US and Canada) |
| Mountain Standard Time (Mexico) | (UTC-07:00) | Chihuahua, La Paz, Mazatlan |
| US Mountain Standard Time | (UTC-07:00) | Arizona |
| New Zealand Standard Time | (UTC+12:00) | Auckland, Wellington |
| Newfoundland Standard Time | (UTC-03:30) | Newfoundland |
| Pacific SA Standard Time | (UTC-03:00) | Santiago |
| Pacific Standard Time | (UTC-08:00) | Pacific Time (US and Canada) |
| Pacific Standard Time (Mexico) | (UTC-08:00) | Baja California |
| Russian Standard Time | (UTC+03:00) | Moscow, St. Petersburg, Volgograd |
| SA Pacific Standard Time | (UTC-05:00) | Bogota, Lima, Quito, Rio Branco |
| SE Asia Standard Time | (UTC+07:00) | Bangkok, Hanoi, Jakarta |
| China Standard Time | (UTC+08:00) | Kuala Lumpur, Singapore |
| Tokyo Standard Time | (UTC+09:00) | Osaka, Sapporo, Tokyo |
| US Eastern Standard Time | (UTC-05:00) | Indiana (East) |
| UTC | UTC | Coordinated Universal Time |
| UTC-02 | (UTC-02:00) | Coordinated Universal Time-02 |
| UTC-08 | (UTC-08:00) | Coordinated Universal Time-08 |
| UTC-09 | (UTC-09:00) | Coordinated Universal Time-09 |
| UTC-11 | (UTC-11:00) | Coordinated Universal Time-11 |
| UTC+12 | (UTC+12:00) | Coordinated Universal Time+12 |
| W. Australia Standard Time | (UTC+08:00) | Perth |
| W. Central Africa Standard Time | (UTC+01:00) | West Central Africa |
| W. Europe Standard Time | (UTC+01:00) | Amsterdam, Berlin, Bern, Rome, Stockholm, Vienna |

9.7. Database connection

9.7.1. Change the endpoint and port number of an instance

This topic describes how to view and change the endpoint and port number of an instance.

View the endpoint and port number

1. Log on to the ApsaraDB for RDS console.
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Connection**.
You can view **Internal Endpoint** and **Internal Port** of the instance on the **Database Connection** page. If you apply for a public endpoint, you can also view **Public Endpoint** and **Public Port**.

Change the endpoint and port number

1. Log on to the ApsaraDB for RDS console.
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Connection**.
5. Click **Change Endpoint**.
6. In the **Change Endpoint** dialog box, select a connection type from the **Connection Type** drop-down list.
7. Set **Endpoint** and **Port**, and click **OK**.

Change Endpoint

Connection Type: Internal Endpoint

Endpoint: [redacted].mysql.rds.intra.env17e.shuguang.com
The endpoint must be 8 to 64 characters and can contain letters, digits, and hyphen (-). It must start with a lowercase letter.

Port: 3306
Port Range: 1000 to 65534

OK Cancel

Note

- The prefix of an endpoint must be 8 to 64 characters in length and can contain only letters, digits, and hyphens (-). It must start with a lowercase letter.
- The port number must be a value within the range of 1000 to 65534.

9.7.2. Connect to an instance

This topic describes how to use Data Management (DMS) to connect to an ApsaraDB RDS instance.

Prerequisites

- A database is created. For more information, see [Create a database](#).
- A database account is created. For more information, see [Create an account](#).

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. Click **Log On to DB** in the upper-right corner of the page.
5. In the **Login instance** dialog box of the **DMS** console, check values of **Database type**, **Instance Area**, and **Connection string address**. If the information is correct, enter **Database account** and **Database password**, as shown in the following figure.

The screenshot shows a 'Login instance' dialog box with the following fields and controls:

- Database type**: A dropdown menu with a red box around it.
- Instance Area**: A dropdown menu with a red box around it.
- Connection string address**: A text input field with a red box around it.
- * Database account**: A text input field with the placeholder text 'Please enter a database account'.
- * Database password**: A text input field.
- Remember password**: A checkbox with a help icon.
- Buttons**: 'Test connection', 'Login', and 'Cancel'.

| Parameter | Description |
|----------------------------------|---|
| Database type | The engine of the database. By default, the engine of the database to be connected is displayed. |
| Instance Area | The region where the instance is deployed. By default, the region of the current instance is displayed. |
| Connection string address | The endpoint of the instance. By default, the endpoint of the current instance is displayed. |

| Parameter | Description |
|-------------------|--|
| Database account | The account of the database to be connected. |
| Database password | The password of the account used to connect to the database. |

6. Click **Login**.

 **Note**

- If you want the browser to remember the password, select **Remember password** and click **Login**.
- If you cannot connect to the instance, check the IP address whitelist settings. For more information, see [Configure a whitelist](#).

9.8. Monitoring and alerting

9.8.1. Set a monitoring frequency

This topic describes how to set the monitoring frequency of an ApsaraDB RDS for SQL Server instance.

Context

ApsaraDB RDS provides the following monitoring frequencies:

- Every 60 seconds
- Every 300 seconds

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Monitoring and Alerts**.
5. On the **Resource Monitoring** tab, click **Set Monitoring Frequency**.
6. In the **Set Monitoring Frequency** dialog box, select the required monitoring frequency.
7. Click **OK**.

9.8.2. View resource and engine monitoring data

The ApsaraDB RDS console provides a variety of performance metrics to monitor the status of your instances.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Monitoring and Alerts**.
5. On the **Monitoring and Alerts** page, select **Resource Monitoring** or **Engine Monitoring**, and select a time range to view the corresponding monitoring data. The following table describes the metrics.

| Monitoring type | Metric | Description |
|----------------------------|---|--|
| Resource Monitoring | Disk Space (unit: MB) | The disk usage of the instance, which includes the following items: <ul style="list-style-type: none"> Instance Size Data Usage Log Size Temporary File Size Other System File Size |
| | IOPS | The number of input/output operations per second (IOPS) for the instance. |
| | Total Connections | The total number of current connections of the instance. |
| | MSSQL Instance CPU Utilization (percentage in the operating system) | The CPU utilization of the instance. This includes the CPU utilization for the operating system. Unit: %. |
| | SQLServer Average Input/Output Traffic | The inbound and outbound traffic of the instance per second. Unit: KB. |
| Engine Monitoring | Average Transaction Frequency | The number of transactions processed per second. |
| | Average QPS | The number of SQL statements executed per second. |
| | Buffer Hit Ratio (%) | The read hit ratio of the buffer pool. |
| | Page Write Frequency at Check Point | The number of checkpoints written to pages per second. |
| | Login Frequency | The number of logons to the instance per second. |
| | Average Frequency of Whole Table Scans | The number of full table scans per second. |
| | SQL Compilations per Second | The number of SQL statements compiled per second. |
| | Lock Timeout Times | The number of lock timeouts on the instance per second. |
| | Deadlock Frequency | The number of deadlocks on the instance per second. |
| | Lock Wait Frequency | The number of lock waits on the instance per second. |

9.9. Data security

9.9.1. Configure an IP address whitelist

To ensure database security and reliability, you must modify the whitelist of an ApsaraDB RDS instance before you enable the instance. You must add the IP addresses or CIDR blocks that are used for database access to the whitelist.

Context

The whitelist improves the access security of your ApsaraDB RDS instance. We recommend that you maintain the whitelist on a regular basis. The whitelist configuration process does not affect the normal operations of the ApsaraDB RDS instance.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. On the **Whitelist Settings** tab, click **Edit** corresponding to the **default** whitelist.

Note

- If you want to connect an ECS instance to an ApsaraDB RDS instance by using an internal endpoint, you must make sure that the two instances are in the same region and have the same network type. Otherwise, the connection fails.
- You can also click **Create Whitelist** to create a new whitelist.

6. In the **Edit Whitelist** dialog box, enter the IP addresses or CIDR blocks that are allowed to access your ApsaraDB RDS instance, and then click **OK**.
 - If you add the CIDR block 10.10.10.0/24, all IP addresses in the 10.10.10.X format are allowed to access the ApsaraDB RDS instance.
 - If you enter more than one IP address or CIDR block, you must separate them with commas (.). Do not add spaces before or after the commas. Example: 192.168.0.1,172.16.213.9.
 - If you click **Add Internal IP Addresses of ECS Instances**, the IP addresses of all of the ECS instances that are created in your Alibaba Cloud account appear. Then, you can select the required IP addresses and add them to the whitelist.

 **Note** If you add a new IP address or CIDR block to the **default** whitelist, the default address 127.0.0.1 is deleted.

Whitelist Name: default

IP Addresses: 127.0.0.1

[Add Internal IP Addresses of ECS Instances](#)
You can add 999 more entries.

Specified IP address: If you specify the IP address 192.168.0.1, this IP address is allowed to access the RDS instance.
Specified CIDR block: If you specify the CIDR block 192.168.0.0/24, the IP addresses ranging from 192.168.0.1 to 192.168.0.255 are allowed to access the RDS instance.
When you add multiple IP addresses or CIDR blocks, separate them by a comma (no space after the comma), for example, 192.168.0.1,192.168.0.0/24.

New whitelist entries take effect in 1 minute.

OK Cancel

9.9.2. Configure SSL encryption

This topic describes how to enhance endpoint security. You can enable Secure Sockets Layer (SSL) encryption and install SSL certificates that are issued by certificate authorities (CAs) to the required application services. SSL is used at the transport layer to encrypt network connections and enhance the security and integrity of communication data. However, SSL increases the response time.

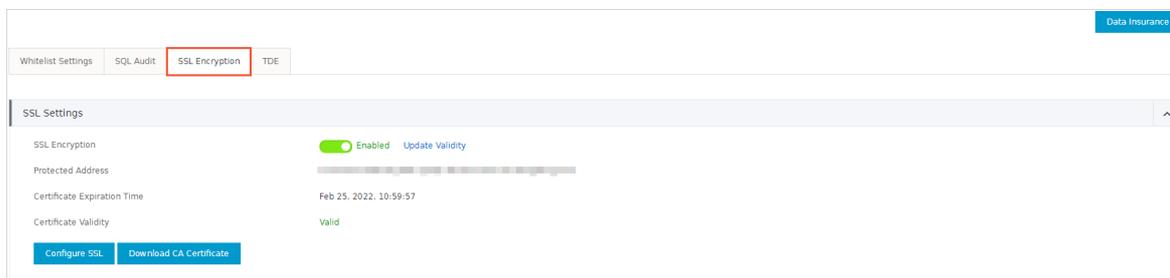
Precautions

- An SSL CA certificate is valid for one year. You must update the validity period of the SSL CA certificate in your application or client within one year. Otherwise, your application or client that uses encrypted network connections cannot connect to the ApsaraDB RDS instance.
- SSL encryption may cause a significant increase in CPU utilization. We recommend that you enable SSL encryption only when you want to encrypt connections from the Internet. In most cases, connections that use an internal endpoint do not require SSL encryption.
- SSL encryption cannot be disabled after it is enabled. Proceed with caution.

Enable SSL encryption

1. [Log on to the ApsaraDB for RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. Click the **SSL Encryption** tab.

6. In the SSL Settings section, turn on **SSL Encryption**.
7. In the **Configure SSL** dialog box, select the endpoint for which you want to enable SSL encryption and click **OK**.
8. Click **Download CA Certificate** to download the SSL CA certificate files in a compressed package.



The downloaded package contains three files:

- o P7B file: used to import CA certificates to the Windows operating system.
- o PEM file: used to import CA certificates to other operating systems or applications.
- o JKS file: the Java truststore file. The password is `apsaradb`. It is used to import the CA certificate chain to Java programs.

Note When the JKS file is used in Java, you must modify the default JDK security configuration in JDK 7 and JDK 8. Open the `/jre/lib/security/java.security` file on the host where your application resides, and modify the following configurations:

```
jdk.tls.disabledAlgorithms=SSLv3, RC4, DH keySize < 224
jdk.certpath.disabledAlgorithms=MD2, RSA keySize < 1024
```

If you do not modify the JDK security configuration, the following error is reported. Typically, other similar errors are also caused by invalid Java security configurations.

```
javax.net.ssl.SSLHandshakeException: DHPublicKey does not comply to algorithm constraints
```

Configure an SSL CA certificate

After you enable SSL encryption, configure the SSL CA certificate on your application or client before they can connect to the ApsaraDB RDS instance. This section describes how to configure an SSL CA certificate. MySQL Workbench and Navicat are used in the example. For more information, see the instructions for the other applications or clients.

Configure a certificate on MySQL Workbench

1. Start MySQL Workbench.
2. Choose **Database > Manage Connections**.
3. Enable **Use SSL** and import the SSL CA certificate files.

Configure a certificate on Navicat

1. Start Navicat.
2. Right-click the database and select **Edit Connection**.
3. Click the **SSL** tab. Select the path of the PEM-formatted CA certificate, as shown in the following figure.
4. Click **OK**.

Note If the `connection is being used` error is reported, the previous session is still connected. Restart Navicat.

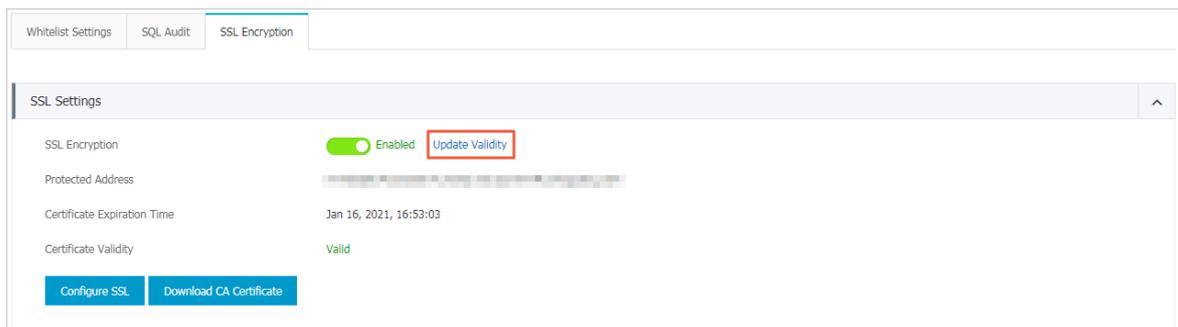
- Double-click the database to test whether the database is connected.

Update the validity period of an SSL CA certificate

Note

- Update Validity** causes the ApsaraDB RDS instance to restart. Proceed with caution.
- After you update the validity period, you must download and configure the SSL CA certificate again.

- Log on to the [ApsaraDB for RDS console](#).
- On the **Instances** page, find the target instance.
- Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
- In the left-side navigation pane, click **Data Security**.
- Click the **SSL Encryption** tab.
- Click **Update Validity**.



- In the message that appears, click **OK**.

9.9.3. Configure TDE

This topic describes how to configure Transparent Data Encryption (TDE) for your ApsaraDB RDS for SQL Server instance. TDE allows your ApsaraDB RDS instance to encrypt the data that will be written into the disk and decrypt the data that will be read from the disk to the memory. TDE does not increase the sizes of data files. When you use TDE, you do not need to modify the application that uses the ApsaraDB RDS instance.

Precautions

- Instance-level TDE can be enabled but cannot be disabled. Database-level TDE can be enabled or disabled.
- The keys used for data encryption are generated and managed by Key Management Service (KMS). ApsaraDB RDS does not provide the keys or certificates used for data encryption. If you want to restore data to your computer after TDE is enabled, you must decrypt the data on your ApsaraDB RDS instance. For more information, see [Decrypt data](#).
- TDE increases CPU utilization.

Prerequisites

- Your ApsaraDB RDS instance runs SQL Server EE.
- KMS is activated. If KMS is not activated, you can activate it as prompted when you enable TDE.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. Click the **TDE** tab.
6. Turn on **TDE Status**.
7. In the dialog box that appears, click **Confirm**.

 **Note** If you have not enabled KMS, you are prompted to do so when you enable TDE. After you enable KMS, you can turn on **TDE Status** to enable TDE.

8. Click **Configure TDE**. In the Database TDE Settings dialog box, select the databases you want to encrypt from the **Unselected Databases** list, click the  icon to add them to the **Selected Databases** list, and then click **OK**.

Decrypt data

If you want to decrypt a database that is encrypted by using TDE, you need only to remove the database from the **Selected Databases** section in the **Database TDE Settings** dialog box.

9.10. Database backup and restoration

9.10.1. Configure an automatic backup policy

Automatic backup supports full physical backups. ApsaraDB RDS automatically backs up data based on pre-configured policies. This topic describes how to configure a policy for automatic backup.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Backup and Restoration**.
5. On the **Backup and Restoration** page, click the **Backup Settings** tab.
6. Click **Edit**.
7. In the dialog box that appears, configure the automatic backup policy.

| Parameter | Description |
|------------------------------|--|
| Data Retention Period | The number of days for which you want to retain data backup files. Valid values: 7 to 730. Default value: 7. |

| Parameter | Description |
|------------------|---|
| Backup Cycle | <p>The cycle based on which you want to create a backup. You can select one or more days within a week.</p> <p>Note For data security purposes, we recommend that you back up your ApsaraDB RDS instance at least twice a week.</p> |
| Backup Time | The period of time for which you want to back up data. Unit: hours. |
| Backup Frequency | <p>The frequency at which you want to back up logs. The following options are available:</p> <ul style="list-style-type: none"> Same as Data Backup Every 30 Minutes <p>The total size of log backup files remains the same regardless of the backup frequency.</p> |

8. Click **OK**.

9.10.2. Manually back up an instance

This topic describes how to manually back up an ApsaraDB RDS instance.

Procedure

- Log on to the [ApsaraDB for RDS console](#).
- On the **Instances** page, find the target instance.
- Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
- On the **Basic Information** page, click **Back Up Instance** in the upper-right corner.
- In the **Back Up Instance** dialog box, select **Automatic Backup** or **Full Backup** from the **Select Backup Mode** drop-down list.

- Note** ApsaraDB RDS supports the following backup methods:
- Automatic Backup:** After you select Automatic Backup, the system immediately performs an incremental or full backup based on the instance.
 - Full Backup:** After you select Full Backup, the system immediately performs a full backup.

6. Click **OK**.

Result

After the backup is complete, you can view the backup task on the **Data Backup** tab of the **Backup and Restoration** page.

9.10.3. Shrink transaction logs

ApsaraDB RDS for SQL Server allows you to shrink transaction logs to reduce the log file size.

Prerequisites

The instance is in the **Running** state.

Procedure

1. Log on to the [ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Backup and Restoration**.
5. In the upper-right corner of the page, click **Shrink Transaction Log**. In the message that appears, click **OK**.

 **Note** The shrinkage takes about 20 minutes to complete. ApsaraDB RDS for SQL Server shrinks transaction logs during each backup.

9.11. Migrate full backup data to ApsaraDB RDS for SQL Server

This topic describes how to migrate full backup files of an on-premises database from Object Storage Service (OSS) to ApsaraDB RDS for SQL Server.

Prerequisites

- Your ApsaraDB for RDS instance has sufficient storage space. If the space is insufficient, you must increase it before you migrate data to the instance.
- The destination database on your ApsaraDB for RDS instance has a different name from the on-premises database.
- A privileged account is created on your ApsaraDB for RDS instance. For more information, see [Create an account](#).
- An Object Storage Service (OSS) bucket is created in the region where your ApsaraDB for RDS instance is created. For more information, see [Create buckets in the OSS User Guide](#).
- The DBCC CHECKDB statement is executed, and the execution result indicates that no allocation or consistency errors occur.

 **Note** If no allocation or consistency errors occur, the following execution result is returned:

```
...  
CHECKDB found 0 allocation errors and 0 consistency errors in database 'xxx'.  
DBCC execution completed. If DBCC printed error messages, contact your system administrator.
```

Precautions

- Full backup files cannot be migrated to an ApsaraDB for RDS instance of an earlier SQL Server version. For example, if the on-premises database runs SQL Server 2016 and your ApsaraDB for RDS instance runs SQL Server 2012, you cannot migrate full backup files of the on-premises database to your ApsaraDB for RDS instance.
- Differential or log backup files are not supported.
- The names of full backup files cannot contain special characters, such as `@` signs and vertical bars (`|`). If the file names contain special characters, the migration fails.
- After the service account of your ApsaraDB for RDS instance is granted the access permission on the OSS bucket, the system creates a role named **AliyunRDSImportRole** in RAM. Do not modify or delete this role. Otherwise, you cannot download full backup files when you migrate data to your ApsaraDB for RDS instance. In this case, you must re-authorize the service account of your ApsaraDB for RDS instance.
- Before the migration is complete, do not delete the backup files from the OSS bucket. Otherwise, the migration fails.
- The names of backup files can be suffixed only with bak, diff, tm, or log. If you do not use the script in this topic

to generate a backup file, you must name the backup file by using one of the following suffixes:

- bak: indicates a full backup file.
- diff: indicates a differential backup file.
- tm or log: indicates a log backup file.

Back up the on-premises database

 **Note** Before you perform a full backup, stop writing data to the on-premises database. The data written during the backup process is not backed up.

1. Download the [backup script](#). Double-click the backup script to open it by using the Microsoft SQL Server Management Studio (SSMS) client.
2. Configure the following parameters.

| Parameter | Description |
|------------------------|--|
| @backup_databases_list | The databases that you want to back up. Separate them with semicolons (;) or commas (,). |
| @backup_type | The backup type. Valid values: <ul style="list-style-type: none"> ◦ FULL: full backup ◦ DIFF: differential backup ◦ LOG: log backup |
| @backup_folder | The directory in which you want to store the backup files on your computer. If the specified directory does not exist, the system creates a directory. |
| @is_run | Specifies whether to perform a backup. Valid values: <ul style="list-style-type: none"> ◦ 1: performs a backup. ◦ 0: performs no backup but a check. |

3. Run the backup script.

Upload full backup files to the OSS bucket

After the on-premises database is backed up, you must upload full backup files to the OSS bucket. You can use one of the following methods:

- Use the OSS console

If the size of backup files is smaller than 5 GB, you can upload the files in the OSS console. For more information, see [Upload objects in the OSS User Guide](#).

- Call an OSS API operation

You can call an OSS API operation to upload the full backup files in resumable mode. For more information, see [Multipart upload-relevant operations in the OSS Developer Guide](#).

Create a migration task

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Backup and Restoration**.
5. In the upper-right corner of the page, click **Migrate OSS Backup Data to RDS**.

6. Click **Next** twice until the Import Data step appears.
7. Configure the following parameters.

| Parameter | Description |
|------------------------|---|
| Database Name | Enter the name of the destination database on your ApsaraDB for RDS instance. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 5px;"> ? Note The name of the database must meet the requirements of SQL Server. </div> |
| OSS Bucket | Select the OSS bucket that stores the backup files. |
| OSS Subfolder Name | Enter the name of the OSS subfolder that stores the backup files. |
| OSS File | Click the search icon to search for backup files by using the prefix-based fuzzy match. The system displays the name, size, and update time of each backup file. Select the backup file that you want to migrate to your ApsaraDB for RDS instance. |
| Cloud Migration Method | One-time Full Backup File Migration: uploads full backup data to your ApsaraDB for RDS instance. Select this option if you want to migrate only a single full backup file. |

8. Click **OK**.

Wait for the migration task to complete. You can click **Refresh** to view the latest status of the migration task. If the migration fails, fix the error based on the message displayed in the Task Description column. For more information, see [Common errors](#).

View the migration task

In the left-side navigation pane, click **Backup and Restoration**. Click the **Backup Data Upload History** tab. The system displays the migration tasks in the last week.

Common errors

Each record of a migration task contains a task description, which helps you identify the error cause and fix the error. The following list describes common errors:

- A database with the same name as the on-premises database exists on your ApsaraDB for RDS instance.
 - Error message: The database (xxx) is already exist on RDS, please backup and drop it, then try again.
 - Cause: The on-premises database is named the same as an existing database on your ApsaraDB for RDS instance. For data security purposes, ApsaraDB RDS for SQL Server does not allow such a database to be migrated.
 - Solution: If you need to overwrite the database in your ApsaraDB for RDS instance with the on-premises database, you must backup the database, delete it from your ApsaraDB for RDS instance, and then migrate the on-premises database to your ApsaraDB for RDS instance.
- A differential backup file is used.
 - Error message: Backup set (xxx.bak) is a Database Differential backup, we only accept a FULL Backup.
 - Cause: The file that you uploaded is a differential backup file, but not a full backup file. The migration solution for full backup data supports only full backup files.
- A log backup file is used.
 - Error message: Backup set (xxx.trn) is a Transaction Log backup, we only accept a FULL Backup.
 - Cause: The file that you uploaded is a log backup file, but not a full backup file. The migration solution for full backup data supports only full backup files.
- The backup file fails the verification.

- Error message: Failed to verify xxx.bak, backup file was corrupted or newer edition than RDS.
 - Cause: The backup file is damaged, or the on-premises database runs an SQL Server version later than your ApsaraDB for RDS instance. For example, if the on-premises database runs SQL Server 2016 and your ApsaraDB for RDS instance runs SQL Server 2012, the error message is returned.
 - Solution: If the backup file is damaged, perform a full backup on the on-premises database again. If the database engine version does not meet the requirements, select an ApsaraDB for RDS instance that runs the same version as or a later version than the on-premises database.
- DBCC CHECKDB fails to be executed.
 - Error message: DBCC checkdb failed.
 - Cause: Allocation or consistency errors occurred in the on-premises database.
 - Solution: Execute the following statement in the on-premises database.

 **Note** Data loss may occur when you use this statement to fix errors.

```
DBCC CHECKDB (DBName, REPAIR_ALLOW_DATA_LOSS) WITH NO_INFOMSGS, ALL_ERRORMSG
```

- The remaining storage space of your ApsaraDB for RDS instance is insufficient. (1)
 - Error message: Not Enough Disk Space for restoring, space left (xxx MB) < needed (xxx MB).
 - Cause: The remaining storage space of your ApsaraDB for RDS instance does not meet the migration requirements.
 - Solution: Increase the storage space of your ApsaraDB for RDS instance.
- The remaining storage space of your ApsaraDB for RDS instance is insufficient. (2)
 - Error message: Not Enough Disk Space, space left xxx MB < bak file xxx MB.
 - Cause: The remaining storage space of your ApsaraDB for RDS instance is smaller than the size of the backup file.
 - Solution: Increase the storage space of your ApsaraDB for RDS instance.
- No privileged account exists.
 - Error message: Your RDS doesn't have any init account yet, please create one and grant permissions on RDS console to this migrated database (XXX).
 - Cause: No privileged account is created on your ApsaraDB for RDS instance, and the database permissions are not granted to accounts. However, when this error message is returned, the backup file has been restored to your ApsaraDB for RDS instance, and the migration task is successful.
 - Solution: Create a privileged account. For more information, see [Create an account](#).

10. ApsaraDB RDS for PostgreSQL

10.1. What is ApsaraDB RDS?

ApsaraDB RDS is a stable, reliable, and scalable online database service. Based on the distributed file system and high-performance storage, ApsaraDB RDS allows you to perform database operations and maintenance with its set of solutions for disaster recovery, backup, restoration, monitoring, and migration.

ApsaraDB RDS supports four database engines, which are MySQL, SQL Server, PolarDB, and PostgreSQL. You can create database instances based on these database engines to meet your business requirements. This topic describes the PostgreSQL engine.

ApsaraDB RDS for PostgreSQL

ApsaraDB RDS for PostgreSQL is the most advanced open source database. It is fully compatible with SQL and supports a diverse range of data formats such as JSON, IP, and geometric data. In addition to features such as transactions, subqueries, multi-version concurrency control (MVCC), and data integrity check, ApsaraDB RDS for PostgreSQL integrates a series of features including high availability, backup, and restoration to ease operation and maintenance loads.

10.2. Limits on ApsaraDB RDS for PostgreSQL

Before you use ApsaraDB RDS for PostgreSQL, you must understand its limits and take the necessary precautions.

The following table describes the limits on ApsaraDB RDS for PostgreSQL.

| Operation | Limit |
|-------------------------------|--|
| Root permissions of databases | Superuser permissions are not provided. |
| Database replication | ApsaraDB RDS for PostgreSQL provides a primary/secondary replication architecture except in the Basic Edition. The secondary instances in the architecture are hidden and cannot be accessed by your applications. |
| Instance restart | Instances must be restarted by using the ApsaraDB RDS console or API operations. |

10.3. Log on to the ApsaraDB RDS console

This topic describes how to log on to the ApsaraDB RDS console.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

Note When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.

4. In the top navigation bar, choose **Products > Database Services > ApsaraDB RDS**.

10.4. Quick Start

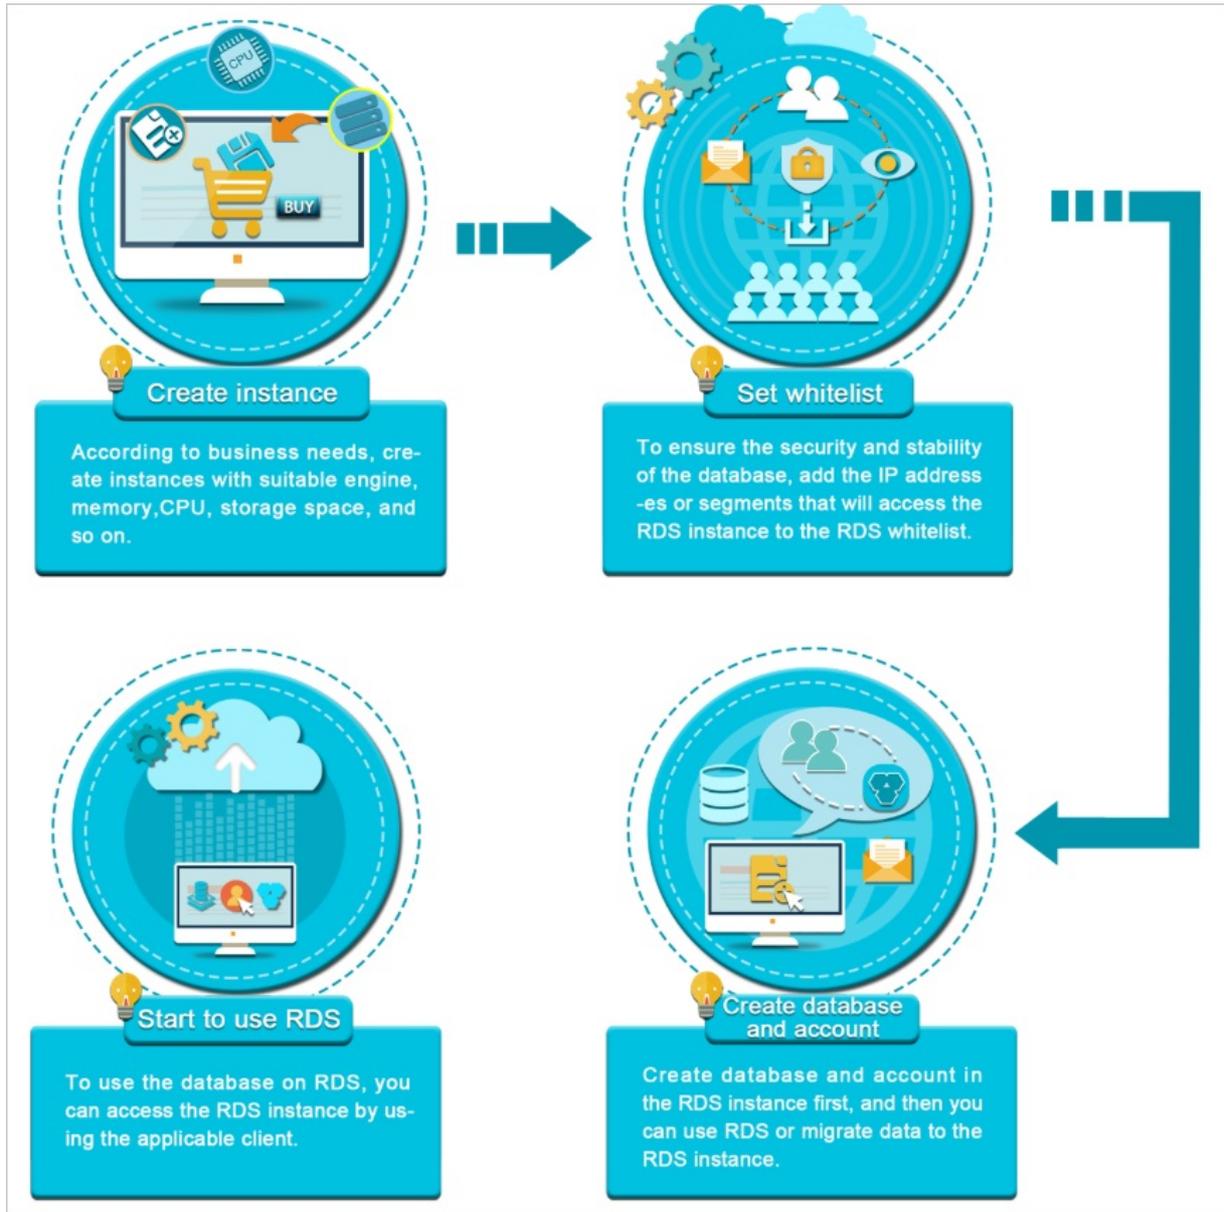
10.4.1. Procedure

ApsaraDB RDS quick start covers the following topics: creating an ApsaraDB RDS instance, configuring a whitelist, creating a database, creating an account, and connecting to the instance.

Flowchart for an ApsaraDB RDS instance

If you are using ApsaraDB RDS for the first time, you can start with [Limits](#).

The following figure shows the operations that you must perform before you use an ApsaraDB RDS instance.



10.4.2. Create an instance

This topic describes how to create one or more ApsaraDB RDS for PostgreSQL instances in the ApsaraDB RDS console.

Procedure

1. Log on to the ApsaraDB for RDS console.
2. On the **Instances** page, click **Create Instance** in the upper-right corner.
3. Configure the following parameters.

| Section | Parameter | Description |
|---------|--------------|---|
| Basic | Organization | The organization to which the instance belongs. |
| | | |

| Settings Section | Parameter | Description |
|-----------------------|-------------------------------|---|
| | Resource Set | The resource set to which the instance belongs. |
| Region | Region | The region in which you want to create the instance. Services in different regions cannot communicate over an internal network. After the instance is created, the region cannot be changed. |
| | Zone of Primary Node | The zone where the primary instance is deployed. |
| | Deployment Method | Specifies whether to deploy the primary and secondary instances in separate zones. ApsaraDB RDS supports Multi-zone Deployment and Single-zone Deployment . If you select Multi-zone Deployment , you must configure Zone of Secondary Node . |
| | Zone of Secondary Node | The zone where the secondary instance is deployed. This parameter is available only when Deployment Method is set to Multi-zone Deployment . <div style="background-color: #e0f2f7; padding: 5px;"> ? Note If you select the same zone for both the primary and secondary instances, the deployment is equivalent to single-zone deployment. </div> |
| Specifications | Quantity | The number of ApsaraDB RDS instances that you want to create. Default value: 1. |
| | Instance Name | The name of the instance. <ul style="list-style-type: none"> ◦ The name must be 2 to 64 characters in length. ◦ The name must start with a letter. ◦ The name can contain letters, digits, and the following special characters: _ - : ◦ The name cannot start with http:// or https://. |
| | Connection Type | The connection type of the instance. ApsaraDB RDS instances support the following connection types: <ul style="list-style-type: none"> ◦ Internet: ApsaraDB RDS instances of this connection type can be connected over the Internet. ◦ Internal Network: ApsaraDB RDS instances of this connection type can be connected over an internal network. <div style="background-color: #e0f2f7; padding: 5px;"> ? Note The value of this parameter cannot be changed after the instance is created. Proceed with caution. </div> |
| | Database Engine | The database engine of the instance. Select PostgreSQL . |
| | Engine Version | The version of the database engine. Valid values: <ul style="list-style-type: none"> ◦ 9.4 ◦ 10.0 ◦ 11.0 ◦ 12.0 |

| Section | Parameter | Description |
|-----------------------------|-------------------------|---|
| | Edition | The edition of the instance. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| | Storage Type | The storage type of the instance. Local and standard SSDs are supported. Note PostgreSQL 9.4 instances support only local SSDs. |
| | Encrypted | Specifies whether to encrypt the standard SSD. This parameter is available only when Storage Type is set to cloud ssd . If you select Encrypted , you must specify the Encryption Key parameter. If you do not have a key, you must first create one in Key Management Service (KMS). For more information, see Create a CMK in <i>KMS User Guide</i> . |
| | Encryption Key | The key that is used to encrypt the standard SSD. This parameter is available only when you select Encrypted . |
| | Instance Type | The instance type of the instance. Memory size determines the maximum number of connections and the input/output operations per second (IOPS). The actual values are displayed in the console. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| | Storage Capacity | The storage capacity of the instance, which includes the space to store data, system files, binlog files, and transaction files. Valid values: 20 to 600. Unit: GB. The value must be in 1 GB increments. |
| | Network | Network Type |
| IP Address Whitelist | | The IP addresses that are allowed to connect to the instance. |

4. Click **Submit**.

10.4.3. Configure an IP address whitelist

This topic describes how to configure a whitelist for an ApsaraDB RDS instance. Only entities that are listed in a whitelist can access your ApsaraDB RDS instance.

Context

Whitelists make your ApsaraDB RDS instance more secure and do not interrupt the operations of your ApsaraDB RDS instance when you configure whitelists. We recommend that you perform maintenance on your whitelists on a regular basis.

To configure a whitelist, perform the following operations:

- Configure a whitelist: Add IP addresses to allow them to connect to the ApsaraDB RDS instance.

 **Note** The IP address whitelist labeled **default** contains only the default IP address 0.0.0.0/0, which allows all entities to access your ApsaraDB RDS instance.

- Configure an ECS security group: Add an ECS security group for the ApsaraDB RDS instance to allow ECS instances in the group to connect to the ApsaraDB RDS instance.

Procedure

1. Log on to the [ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. On the **Whitelist Settings** tab, click **Edit** corresponding to the **default** whitelist.

 **Note** You can also click **Create Whitelist** to create a whitelist.

6. In the **Edit Whitelist** dialog box, enter the IP addresses or CIDR blocks used to access the instance and click **OK**. The following section describes the rules:
 - If you enter the CIDR block 10.10.10.0/24 in the IP Addresses field, all IP addresses in the 10.10.10.X format can access your ApsaraDB RDS instance.
 - If you enter more than one IP address or CIDR block, you must separate them with commas (,). Do not add spaces before or after the commas. Example: 192.168.0.1,172.16.213.9.
 - If you click **Add Internal IP Addresses of ECS Instances**, the IP addresses of all ECS instances created within your Alibaba Cloud account are displayed. You can select the required IP addresses to add them to the IP address whitelist.

10.4.4. Create a database and an account

Before you start to use ApsaraDB RDS, you must create databases and accounts on an ApsaraDB RDS instance. This topic describes how to create a database and an account on an ApsaraDB RDS for PostgreSQL instance.

Account types

ApsaraDB RDS for PostgreSQL instances support two types of accounts: privileged accounts and standard accounts. The following table describes these account types.

| Account type | Description |
|--------------|-------------|
|--------------|-------------|

| Account type | Description |
|---------------------------|---|
| Privileged account | <ul style="list-style-type: none"> You can create and manage privileged accounts only by using the ApsaraDB RDS console or API operations. If your ApsaraDB RDS instance uses local SSDs, you can create only a single privileged account. If your ApsaraDB RDS instance uses standard or enhanced SSDs, you can create more than one privileged account. A privileged account allows you to manage all the standard accounts and databases that are created on your ApsaraDB RDS instance. A privileged account has more permissions that allow you to manage your ApsaraDB RDS instance at more fine-grained levels. For example, you can grant the query permissions on different tables to different users. A privileged account has the permissions to disconnect accounts that are created on your ApsaraDB RDS instance. |
| Standard account | <ul style="list-style-type: none"> You can create and manage standard accounts by using the ApsaraDB RDS console, API operations, or SQL statements. You can create more than one standard account on your ApsaraDB RDS instance. You must grant the permissions on specific databases to a standard account. A standard account does not have the permissions to create, manage, or disconnect other accounts on your ApsaraDB RDS instance. |

Precautions

- If your ApsaraDB RDS instance uses local SSDs, you can create one privileged account in the ApsaraDB RDS console. After the privileged account is created, it cannot be deleted. You can also create and manage more than one standard account by using SQL statements.
- If your ApsaraDB RDS instance uses standard or enhanced SSDs, you can create more than one privileged account and standard account in the ApsaraDB RDS console. You can also create and manage more than one standard account by using SQL statements.
- To migrate data from an on-premises database to your ApsaraDB RDS instance, you must create a database and an account on the ApsaraDB RDS instance. Make sure that the created database has the same properties as the on-premises database. Also make sure that the created account has the same permissions on the created database as the account that is authorized to manage the on-premises database.
- Follow the least privilege principle to create accounts and grant them read-only permissions or read and write permissions on databases. If necessary, you can create more than one account and grant them only the permissions on specific databases. If an account does not need to write data to a database, grant only the read-only permissions on that database to the account.
- For security purposes, we recommend that you specify strong passwords for the accounts on your ApsaraDB RDS instance and change the passwords on a regular basis.

Create a privileged account on an instance that uses local SSDs

- Log on to the [ApsaraDB RDS console](#).
- On the **Instances** page, find the target instance.
- Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
- In the left-side navigation pane, click **Accounts**.
- On the Accounts page, click **Create Privileged Account** and configure the following parameters.

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|-------------------|---|
| Database Account | <ul style="list-style-type: none"> The name of the account must be 2 to 16 characters in length. The name can contain lowercase letters, digits, and underscores (_). The name must start with a letter and end with a letter or digit. |
| Password | <ul style="list-style-type: none"> The password of the account must be 8 to 32 characters in length. The password must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. Special characters include !@#%&^&*()_+ -= |
| Re-enter Password | Enter the password of the account again. |

6. Click **Create**.

Create a privileged or standard account on an instance that uses standard or enhanced SSDs

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Accounts**.
5. On the Accounts page, click **Create Account** and configure the following parameters.

| Parameter | Description |
|-------------------|---|
| Database Account | <ul style="list-style-type: none"> The name of the account must be 2 to 16 characters in length. The name can contain lowercase letters, digits, and underscores (_). The name must start with a letter and end with a letter or digit. |
| Account Type | Select Privileged Account or Standard Account . |
| Password | <ul style="list-style-type: none"> The password of the account must be 8 to 32 characters in length. The password must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. Special characters include !@#%&^&*()_+ -= |
| Re-enter Password | Enter the password of the account again. |
| Description | This parameter is optional. You can enter relevant description to make the instance identifiable. The description can be up to 256 characters in length. |

6. Click **Create**.

Create a database and a standard account

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.

- Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
- Click **Log On to DB** in the upper-right corner of the page.
- In the **Login instance** dialog box of the **DMS** console, check values of **Database type**, **Instance Area**, and **Connection string address**. If the information is correct, enter **Database account** and **Database password**, as shown in the following figure.

| Parameter | Description |
|----------------------------------|---|
| Database type | The engine of the database. By default, the engine of the database to be connected is displayed. |
| Instance Area | The region where the instance is deployed. By default, the region of the current instance is displayed. |
| Connection string address | The endpoint of the instance. By default, the endpoint of the current instance is displayed. |
| Database account | The account of the database to be connected. |
| Database password | The password of the account used to connect to the database. |

- Click **Login**. If you want the browser to remember the password, select **Remember password** before you click **Login**.

Note If a message prompts you that the connection fails, the problem may be caused by an improperly configured whitelist. Reconfigure the whitelist in the console. For more information, see [Configure an IP address whitelist](#).

- The **SQLConsole** page appears after you log on to the instance. Execute a statement in the following format to create a database:

```
CREATE DATABASE name
[[ WITH ] [ OWNER [=] user_name ]
[ TEMPLATE [=] template ]
[ ENCODING [=] encoding ]
[ LC_COLLATE [=] lc_collate ]
[ LC_CTYPE [=] lc_ctype ]
[ TABLESPACE [=] tablespace_name ]
[ CONNECTION LIMIT [=] connlimit ]]
```

For example, if you want to create a database named test, execute the following statement:

```
create database test;
```

- Click **execute**.
- In the SQL window, execute a statement in the following format to create a standard account:

```
CREATE USER name [[ WITH ] option [ ... ]
where option can be:
SUPERUSER | NOSUPERUSER
| CREATEDB | NOCREATEDB
| CREATEROLE | NOCREATEROLE
| CREATEUSER | NOCREATEUSER
| INHERIT | NOINHERIT
| LOGIN | NOLOGIN
| REPLICATION | NOREPLICATION
| CONNECTION LIMIT connlimit
| [ ENCRYPTED | UNENCRYPTED ] PASSWORD 'password'
| VALID UNTIL 'timestamp'
| IN ROLE role_name [, ...]
| IN GROUP role_name [, ...]
| ROLE role_name [, ...]
| ADMIN role_name [, ...]
| USER role_name [, ...]
| SYSID uid
```

For example, if you want to create a user account named test2 whose password is 123456, execute the following statement:

```
create user test2 password '123456';
```

- Click **execute**.

10.4.5. Connect to an ApsaraDB RDS for PostgreSQL instance

This topic describes how to use Data Management (DMS) or the pgAdmin 4 client to connect to an ApsaraDB RDS instance.

Context

You can log on to DMS from the ApsaraDB RDS console and then connect to an ApsaraDB RDS instance.

DMS is a data management service that integrates data, schema, and server management, access security, BI charts, data trends, data tracking, and performance optimization. DMS can be used to manage relational and non-relational databases, such as MySQL, SQL Server, PostgreSQL, MongoDB, and Redis. It can also be used to manage Linux servers.

You can also use a client to connect to an ApsaraDB RDS instance. ApsaraDB RDS for PostgreSQL is fully compatible with PostgreSQL. You can connect to an ApsaraDB RDS for PostgreSQL instance in a similar manner as you would connect to an open source PostgreSQL instance. In this topic, the pgAdmin 4 client is used to connect to an ApsaraDB RDS instance.

Use DMS to connect to an ApsaraDB RDS instance

For more information about how to use DMS to connect to an ApsaraDB RDS instance, see [Log on to an ApsaraDB for RDS instance by using DMS](#).

Use the pgAdmin 4 client to connect to an ApsaraDB RDS instance

1. Add the IP address of the pgAdmin client to an IP address whitelist of the ApsaraDB RDS instance. For more information about how to configure a whitelist, see [Configure an IP address whitelist](#).
2. Start the pgAdmin 4 client.

 **Note** For information about how to download the pgAdmin 4 client, visit [pgAdmin 4 \(Windows\)](#).

3. Right-click **Servers** and choose **Create > Server**, as shown in the following figure.
4. On the **General** tab of the **Create - Server** dialog box, enter the name of the server, as shown in the following figure.
5. Click the **Connection** tab and enter the information of the instance, as shown in the following figure.

| Parameter | Description |
|-------------------|---|
| Host name/address | The internal endpoint of the ApsaraDB RDS instance. For more information about how to view the internal endpoint, see View and modify the internal endpoint and port number . |
| Port | The internal port number that is used to connect to the ApsaraDB RDS instance. For more information about how to view the internal port number, see View and modify the internal endpoint and port number . |
| Username | The name of the privileged account on the ApsaraDB RDS instance. For more information about how to obtain a privileged account, see Create a database and an account . |
| Password | The password of the privileged account of the ApsaraDB RDS instance. |

6. Click **Save**.
7. If the connection information is correct, choose **Servers > Server Name > Databases > postgres**. If the following page appears, the connection is established.

 **Notice** The postgres database is the default system database of the ApsaraDB RDS instance. Do not perform operations on this database.

10.5. Instances

10.5.1. Create an instance

This topic describes how to create one or more ApsaraDB RDS for PostgreSQL instances in the ApsaraDB RDS console.

Procedure

1. Log on to the [ApsaraDB for RDS console](#).
2. On the **Instances** page, click **Create Instance** in the upper-right corner.
3. Configure the following parameters.

| Section | Parameter | Description |
|----------------|------------------------|--|
| Basic Settings | Organization | The organization to which the instance belongs. |
| | Resource Set | The resource set to which the instance belongs. |
| Region | Region | The region in which you want to create the instance. Services in different regions cannot communicate over an internal network. After the instance is created, the region cannot be changed. |
| | Zone of Primary Node | The zone where the primary instance is deployed. |
| | Deployment Method | Specifies whether to deploy the primary and secondary instances in separate zones. ApsaraDB RDS supports Multi-zone Deployment and Single-zone Deployment . If you select Multi-zone Deployment , you must configure Zone of Secondary Node . |
| | Zone of Secondary Node | The zone where the secondary instance is deployed. This parameter is available only when Deployment Method is set to Multi-zone Deployment .  Note If you select the same zone for both the primary and secondary instances, the deployment is equivalent to single-zone deployment. |
| | Quantity | The number of ApsaraDB RDS instances that you want to create. Default value: 1. |
| | Instance Name | The name of the instance. <ul style="list-style-type: none"> ◦ The name must be 2 to 64 characters in length. ◦ The name must start with a letter. ◦ The name can contain letters, digits, and the following special characters: _ - : ◦ The name cannot start with http:// or https://. |
| | Connection Type | The connection type of the instance. ApsaraDB RDS instances support the following connection types: <ul style="list-style-type: none"> ◦ Internet: ApsaraDB RDS instances of this connection type can be connected over the Internet. ◦ Internal Network: ApsaraDB RDS instances of this connection type can be connected over an internal network.  Note The value of this parameter cannot be changed after the instance is created. Proceed with caution. |

| Section | Parameter | Description |
|----------------|-------------------------|---|
| Specifications | Database Engine | The database engine of the instance. Select PostgreSQL . |
| | Engine Version | The version of the database engine. Valid values: <ul style="list-style-type: none"> ◦ 9.4 ◦ 10.0 ◦ 11.0 ◦ 12.0 |
| | Edition | The edition of the instance. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| | Storage Type | The storage type of the instance. Local and standard SSDs are supported. <div style="background-color: #e0f2f1; padding: 5px; margin-top: 5px;"> ? Note PostgreSQL 9.4 instances support only local SSDs. </div> |
| | Encrypted | Specifies whether to encrypt the standard SSD. This parameter is available only when Storage Type is set to cloud ssd . If you select Encrypted, you must specify the Encryption Key parameter. If you do not have a key, you must first create one in Key Management Service (KMS). For more information, see Create a CMK in <i>KMS User Guide</i> . |
| | Encryption Key | The key that is used to encrypt the standard SSD. This parameter is available only when you select Encrypted . |
| | Instance Type | The instance type of the instance. Memory size determines the maximum number of connections and the input/output operations per second (IOPS). The actual values are displayed in the console. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| | Storage Capacity | The storage capacity of the instance, which includes the space to store data, system files, binlog files, and transaction files. Valid values: 20 to 600. Unit: GB. The value must be in 1 GB increments. |

| Section | Parameter | Description |
|---------|----------------------|---|
| Network | Network Type | <p>The network type of the instance. ApsaraDB RDS instances support the following network types:</p> <ul style="list-style-type: none"> ◦ Classic Network: Cloud services in the classic network are not isolated. Unauthorized access to a cloud service is blocked only by the security group or whitelist policy of the service. ◦ VPC: A virtual private cloud (VPC) helps you build an isolated network environment on Apsara Stack. You can customize route tables, CIDR blocks, and gateways in a VPC. We recommend that you select VPC for improved security. <div style="background-color: #e6f2ff; padding: 10px; border: 1px solid #d9e1f2;"> <p> Note</p> <ul style="list-style-type: none"> ▪ If you configure multi-zone deployment, you must create vSwitches for the zones of primary and secondary instances in the specified VPC. ▪ If you select VPC, you must specify a VPC and a vSwitch. </div> |
| | IP Address Whitelist | The IP addresses that are allowed to connect to the instance. |

4. Click **Submit**.

10.5.2. Create an ApsaraDB RDS for PostgreSQL instance that uses standard or enhanced SSDs

Cloud disks are block-level data storage products provided by Alibaba Cloud for Elastic Compute Service (ECS). They provide low latency, high performance, durability, and reliability. This topic describes how to create one or more instances that use standard or enhanced SSDs in the ApsaraDB RDS console.

Prerequisites

The instance runs PostgreSQL 10.0 or later.

Context

An ApsaraDB RDS instance with standard or enhanced SSDs uses a distributed triplicate mechanism to ensure 99.9999999% data reliability. If service disruptions occur within a zone due to hardware faults, data in that zone is copied to an unaffected disk in another zone to ensure data availability.

Procedure

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, click **Create Instance** in the upper-right corner.
3. Configure the following parameters.

| Section | Parameter | Description |
|----------------|------------------------|--|
| Basic Settings | Organization | The organization to which the instance belongs. |
| | Resource Set | The resource set to which the instance belongs. |
| Region | Region | The region in which you want to create the instance. Services in different regions cannot communicate over an internal network. After the instance is created, the region cannot be changed. |
| | Zone of Primary Node | The zone where the primary instance is deployed. |
| | Deployment Method | Specifies whether to deploy the primary and secondary instances in separate zones. ApsaraDB RDS supports Multi-zone Deployment and Single-zone Deployment . If you select Multi-zone Deployment , you must configure Zone of Secondary Node . |
| | Zone of Secondary Node | The zone where the secondary instance is deployed. This parameter is available only when Deployment Method is set to Multi-zone Deployment . <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e6f2ff;"> ? Note If you select the same zone for both the primary and secondary instances, the deployment is equivalent to single-zone deployment. </div> |
| Specifications | Quantity | The number of ApsaraDB RDS instances that you want to create. Default value: 1. |
| | Instance Name | The name of the instance. <ul style="list-style-type: none"> ◦ The name must be 2 to 64 characters in length. ◦ The name must start with a letter. ◦ The name can contain letters, digits, and the following special characters: _ - : ◦ The name cannot start with http:// or https://. |
| | Connection Type | The connection type of the instance. ApsaraDB RDS instances support the following connection types: <ul style="list-style-type: none"> ◦ Internet: ApsaraDB RDS instances of this connection type can be connected over the Internet. ◦ Internal Network: ApsaraDB RDS instances of this connection type can be connected over an internal network. <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e6f2ff;"> ? Note The value of this parameter cannot be changed after the instance is created. Proceed with caution. </div> |
| | Database Engine | The database engine of the instance. Select PostgreSQL . |
| | Engine Version | The version of the database engine. Set the value to 10.0 or a later version number. |
| | Edition | The edition of the instance. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |

| Section | Parameter | Description |
|----------------|-----------------------------|--|
| | Storage Type | The storage type of the instance. Set the value to cloud ssd . |
| | Encrypted | Specifies whether to encrypt the standard SSD. If you select Encrypted , you must specify the Encryption Key parameter. If you do not have a key, you must first create one in Key Management Service (KMS). For more information, see Configure data encryption . Note Disk encryption provides maximum protection for your data with minimal impact on your business or applications. Both the snapshots generated from encrypted disks and the disks created from those snapshots are automatically encrypted. |
| | Encryption Key | The key that is used to encrypt the standard SSD. This parameter is available only when you select Encrypted . |
| | Instance Type | The instance type of the instance. Memory size determines the maximum number of connections and the input/output operations per second (IOPS). The actual values are displayed in the console. For more information, see Instance types in <i>ApsaraDB RDS Product Introduction</i> . |
| | Storage Capacity | The storage capacity of the instance, which includes the space to store data, system files, binlog files, and transaction files. Valid values: 20 to 600. Unit: GB. The value must be in 1 GB increments. |
| Network | Network Type | The network type of the instance. ApsaraDB RDS instances support the following network types: <ul style="list-style-type: none"> ◦ Classic Network: Cloud services in the classic network are not isolated. Unauthorized access to a cloud service is blocked only by the security group or whitelist policy of the service. ◦ VPC: A virtual private cloud (VPC) helps you build an isolated network environment on Apsara Stack. You can customize route tables, CIDR blocks, and gateways in a VPC. We recommend that you select VPC for improved security. Note <ul style="list-style-type: none"> ▪ If you configure multi-zone deployment, you must create vSwitches for the zones of primary and secondary instances in the specified VPC. ▪ If you select VPC, you must specify a VPC and a vSwitch. |
| | IP Address Whitelist | The IP addresses that are allowed to connect to the instance. |

4. Click **Submit**.

10.5.3. View basic information of an instance

This topic describes how to view the details of an ApsaraDB RDS instance, such as basic information, internal network connection information, status, and configurations.

Procedure

1. [Log on to the ApsaraDB RDS console](#).

2. Use one of the following methods to go to the **Basic Information** page of an instance:
 - On the **Instances** page, click the ID of the instance to go to the **Basic Information** page.
 - On the **Instances** page, find the instance and click **Manage** in the corresponding **Actions** column. The **Basic Information** page appears.

10.5.4. Restart an instance

This topic describes how to manually restart an ApsaraDB RDS instance. This applies if the number of connections exceeds the specified threshold or if an instance has performance issues.

Prerequisites

The instance is in the **Running** state.

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the upper-right corner of the page, click **Restart Instance**.

 **Note** When you restart an instance, applications are disconnected from the instance. We recommend that you make appropriate service arrangements before you restart an instance. Proceed with caution.

5. In the message that appears, click **Confirm**.

10.5.5. Change the specifications of an instance

This topic describes how to change the specifications of an ApsaraDB RDS instance. You can upgrade or downgrade an ApsaraDB RDS instance to meet your business needs.

Prerequisites

The instance is in the **Running** state and is not in the backing up or restoring state.

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the **Configure Information** section of the **Basic Information** page, click **Change Specifications**.
5. On the **Change Specifications** page, set **Edition**, **Instance Type**, and **Storage Capacity**.
6. Click **Submit**.

10.5.6. Set a maintenance window

This topic describes how to set a maintenance window for an ApsaraDB RDS instance.

Context

The backend system performs maintenance on the ApsaraDB RDS instances during the maintenance window. This ensures the stability of the ApsaraDB RDS instance. The default maintenance window is from 02:00 (UTC+8) to

06:00 (UTC+8). We recommend that you set the maintenance window to off-peak hours of your business to avoid impacts on your business.

Precautions

- An instance enters the **Maintaining Instance** state before the maintenance window to ensure stability during the maintenance process. When the instance is in this state, access to data in the database and query operations such as performance monitoring are not affected. However, except for account and database management and IP address whitelist configuration, modification operations such as upgrade, downgrade, and restart are temporarily unavailable.
- During the maintenance window, one or two network interruptions may occur. Make sure that your applications are configured with automatic reconnection policies.

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the **Configuration Information** section, click **Configure** to the right of **Maintenance Window**.
5. Select a maintenance window and click **Save**.

 **Note** The maintenance window is displayed in UTC+8.

10.5.7. Configure primary/secondary switchover

ApsaraDB RDS provides the primary/secondary switchover feature to ensure the high availability of databases. The primary/secondary switchover is performed when the primary instance becomes unavailable. You can also manually switch your business to the secondary instance. This topic describes how to manually switch over services between a primary instance and its secondary instance.

Context

Data is synchronized in real time between the primary and secondary instances. You can access only the primary instance. The secondary instance serves only as a backup instance and does not allow external access. After the switchover, the original primary instance becomes the secondary instance.

 **Note** Network interruptions may occur during a switchover. Make sure that your applications are configured with automatic reconnection policies.

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Service Availability**.
5. In the **Availability Information** section, click **Switch Primary/Secondary Instance**.
6. In the **Switch Primary/Secondary Instance** message, click **OK**.

 Note

- During the switchover, operations such as managing databases and accounts and changing network types cannot be performed. Therefore, we recommend that you select Switch Within Maintenance Window.
- For more information about how to set a maintenance window, see [Set a maintenance window](#).

10.5.8. Release an instance

This topic describes how to manually release an instance.

Precautions

- Only instances in the running state can be released.
- After an instance is released, the instance data is immediately deleted. We recommend that you back up your data before you release an instance.
- When you release a primary instance, all of its read-only instances are also released.

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. Find the instance that you want to release and choose **More > Release Instance** in the Actions column.
3. In the **Release Instance** message, click **Confirm**.

10.5.9. Modify parameters of an instance

This topic describes how to view and modify the values of some parameters and query parameter modification records in the console.

Precautions

- To ensure instance stability, you can modify only specific parameters in the ApsaraDB RDS console.
- When you modify parameters on the **Editable Parameters** tab, refer to the **Value Range** column corresponding to each parameter.
- After specific parameters are modified, you must restart your instance for the changes to take effect. The necessity of restart is displayed in the **Force Restart** column on the **Editable Parameters** tab. We recommend that you modify the parameters of an instance during off-peak hours and make sure that your applications are configured with automatic reconnection policies.

Modify parameters

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Parameters**.
5. Perform the following operations:

Export the parameter settings of the instance to your computer.

On the **Editable Parameters** tab, click **Export Parameters**. The parameter settings of the instance are exported as a TXT file to your computer.

Modify and import the parameter settings.

- i. After you modify parameters in the exported parameter file, click **Import Parameters** and copy the parameter settings to the field.
- ii. Click **OK**.
- iii. In the upper-right corner of the page, click **Apply Changes**.

 **Note**

- If the new parameter value takes effect only after an instance restart, the system prompts you to restart the instance. We recommend that you restart the instance during off-peak hours and make sure that your applications are configured with automatic reconnection policies.
- Before the new parameter values are applied, you can click **Cancel Changes** to cancel them.

Modify a single parameter.

- i. On the **Editable Parameters** tab, find the parameter that you want to modify and click the  icon in the **Actual Value** column.
- ii. Enter a new value based on the prompted value range.
- iii. Click **Confirm**.
- iv. In the upper-right corner of the page, click **Apply Changes**.

 **Note**

- If the new parameter value takes effect only after an instance restart, the system prompts you to restart the instance. We recommend that you restart the instance during off-peak hours and make sure that your applications are configured with automatic reconnection policies.
- Before the new parameter value is applied, you can click **Cancel Changes** to cancel it.

View the parameter modification history

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Parameters**.
5. On the page that appears, click the **Edit History** tab.
6. Select a time range and click **Search**.

10.5.10. Read-only instances

10.5.10.1. Overview of read-only ApsaraDB RDS for PostgreSQL instances

This topic provides an overview of read-only ApsaraDB RDS for PostgreSQL instances. If a large number of read requests overwhelm the primary instance, your business may be interrupted. In this case, you can create one or more read-only instances to offload read requests from the primary instance. This scales the read capability of your database system and increases the throughput of your application.

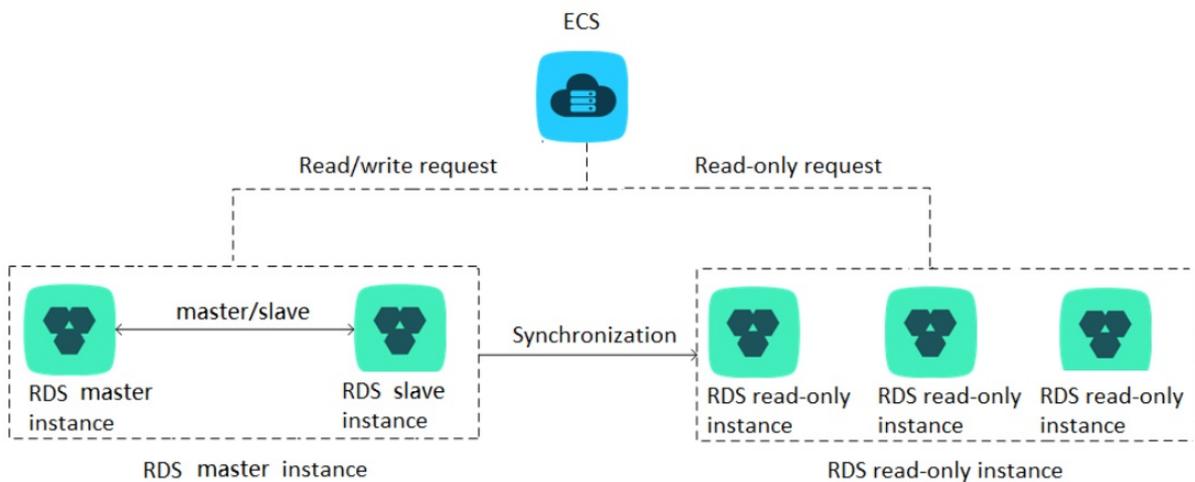
Overview

When a read-only instance is created, the data is replicated from the secondary instance. The data is consistent with that of the primary instance. Data updates of the primary instance are automatically synchronized to all read-only instances immediately after the primary instance completes operations.

Note

- Only ApsaraDB RDS instances that run PostgreSQL 10.0 support read-only instances.
- The specifications of the primary instance must have at least eight CPU cores and 32 GB of memory.
- Each read-only instance works in a single-node architecture, where no instances are provided as backups.

The following figure shows the topology of read-only instances.



Features

- **Region and zone:** Read-only instances reside within the same region as the primary instance, but possibly in different zones.
- **Specifications and storage space:** The specifications and storage space of read-only instances cannot be lower than those of the primary instance.
- **Network type:** The network type of a read-only instance can differ from that of the primary instance.
- **Account and database management:** Read-only instances do not require database or account maintenance, because their database and account information is synchronized with the primary instance.
- **IP address whitelist:** A read-only instance automatically replicates the IP address whitelists of the primary instance. However, the IP address whitelists for the read-only instance are independent of those of the primary instance. For information about how to modify the IP address whitelists of a read-only instance, see [Configure an IP address whitelist](#).
- **Monitoring and alerts:** You can monitor system performance metrics, such as the disk capacity, IOPS, number of connections, and CPU utilization.

Limits

- **Number of read-only instances:** A maximum of five read-only instances can be created on a primary instance.
- **Instance backup:** Read-only instances do not support backup settings or manual backups because backups have been configured or created on the primary instance.
- **Data migration:** You cannot migrate data to read-only instances.
- **Database management:** You cannot create or delete databases on read-only instances.
- **Account management:** You cannot create or delete accounts, authorize accounts, or change the passwords of accounts on read-only instances.

FAQ

Q: Can I manage accounts created on the primary instance from its read-only instances?

A: No, although accounts created on the primary instance are replicated to its read-only instances, you cannot manage the accounts on the read-only instances. The accounts have only read permissions on the read-only instances.

10.5.10.2. Create a read-only ApsaraDB RDS for PostgreSQL instance

This topic describes how to create a read-only instance for your primary ApsaraDB RDS for PostgreSQL instance. This allows your database system to process a large number of read requests and increases the throughput of your application. The data on each read-only instance is a copy of that of the primary instance. Data updates to the primary instance are synchronized to all of its read-only instances.

Prerequisites

- The primary instance runs PostgreSQL 10.0.
- The specifications of the primary instance must have at least eight CPU cores and 32 GB of memory.

Precautions

- You can create read-only instances only for your primary instance. You cannot change existing instances to read-only instances.
- When you create a read-only instance, the system replicates data from a secondary instance. Therefore, operations on your primary instance are not interrupted.
- A read-only instance does not inherit the parameter settings of the primary instance. The system generates default parameter settings for the read-only instance, and you can modify the settings in the ApsaraDB RDS console.
- The instance type and storage capacity of a read-only instance cannot be lower than that of the primary instance.
- You can create up to five read-only instances.

Create a read-only instance

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the Distributed by Instance Role section of the Basic Information page, click **Create Read-only Instance**.
5. On the Create Read-only Instance page, configure parameters and click **Submit**. The following table describes the parameters.

| Section | Parameter | Description |
|---------|-----------------|---|
| Region | Region | The region where the instance is deployed. |
| | Database Engine | The database engine of the read-only instance, which is the same as that of the primary instance and cannot be changed. |
| | Engine Version | The engine version of the read-only instance, which is the same as that of the primary instance and cannot be changed. |

| Section | Parameter | Description |
|----------------|------------------|--|
| Specifications | Edition | The edition of the read-only instance, which is the same as that of the primary instance and cannot be changed. |
| | Instance Type | The instance type of the read-only instance. The instance type of the read-only instance can be different from that of the primary instance, and can be changed at any time to facilitate flexible upgrade and downgrade. To ensure sufficient I/O throughput for data synchronization, we recommend that you select at least the same instance type as the primary instance for the read-only instance. |
| | Storage Capacity | The storage capacity of the read-only instance. To ensure sufficient I/O throughput for data synchronization, we recommend that you select at least the same storage capacity as the primary instance for the read-only instance. |
| Network Type | Network Type | The network type of the read-only instance, which is the same as that of the primary instance and cannot be changed. |
| | VPC | Select a VPC if the network type is set to VPC. |
| | vSwitch | Select a vSwitch if the network type is set to VPC. |

10.5.10.3. View a read-only ApsaraDB RDS for PostgreSQL instance

This topic describes how to view details of a read-only ApsaraDB RDS for PostgreSQL instance. You can go to the Basic Information page of a read-only instance from the Instances page or the read-only instance list of the primary instance. Read-only instances are managed in the same manner as primary instances. The Basic Information page shows the management operations that can be performed.

View instance details of a read-only instance by using its ID

1. [Log on to the ApsaraDB RDS console.](#)
2. On the **Instances** page, find the instance that you want to view.
3. Click the ID of the instance or click **Manage** in the corresponding Actions column to go to the Basic Information page.

View details of a read-only instance by using the primary instance

1. [Log on to the ApsaraDB RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. On the **Basic Information** page, move the pointer over the number below **Read-only Instance** in the **Distributed by Instance Role** section. The ID of the read-only instance is displayed.
5. Click the ID of the read-only instance to go to the Basic Information page of the read-only instance.

View the latency of a read-only instance

When a read-only instance synchronizes data from its primary RDS instance, latency may occur. You can navigate to the Basic Information page of the read-only instance to view the latency of data synchronization to the instance.

The screenshot displays the console interface for an ApsaraDB RDS instance. The 'Basic Information' section includes details like Instance ID, Name, Region, and Type. The 'Status' section shows the instance is 'Running'. The 'Configuration Information' section lists hardware and software specifications. The 'Usage Statistics' section is expanded to show a table of latency metrics for a read-only instance.

| Delay for Read-only Instance | | | |
|---|--|--|---|
| Delay for Sending Write-Ahead Logging Data: 0MB | Delay for Writing Write-Ahead Logging Data: 0MB | Delay for Syncing Write-Ahead Logging Data: 0MB | Delay for Applying Write-Ahead Logging Data: 0MB |
| | Delay for Writing Write-Ahead Logging Data: 0.000103Second | Delay for Syncing Write-Ahead Logging Data: 0.000152Second | Delay for Applying Write-Ahead Logging Data: 0.0002Second |

10.6. Database connection

10.6.1. Connect to an ApsaraDB RDS for PostgreSQL instance

This topic describes how to use Data Management (DMS) or the pgAdmin 4 client to connect to an ApsaraDB RDS instance.

Context

You can log on to DMS from the ApsaraDB RDS console and then connect to an ApsaraDB RDS instance.

DMS is a data management service that integrates data, schema, and server management, access security, BI charts, data trends, data tracking, and performance optimization. DMS can be used to manage relational and non-relational databases, such as MySQL, SQL Server, PostgreSQL, MongoDB, and Redis. It can also be used to manage Linux servers.

You can also use a client to connect to an ApsaraDB RDS instance. ApsaraDB RDS for PostgreSQL is fully compatible with PostgreSQL. You can connect to an ApsaraDB RDS for PostgreSQL instance in a similar manner as you would connect to an open source PostgreSQL instance. In this topic, the pgAdmin 4 client is used to connect to an ApsaraDB RDS instance.

Use DMS to connect to an ApsaraDB RDS instance

For more information about how to use DMS to connect to an ApsaraDB RDS instance, see [Log on to an ApsaraDB for RDS instance by using DMS](#).

Use the pgAdmin 4 client to connect to an ApsaraDB RDS instance

1. Add the IP address of the pgAdmin client to an IP address whitelist of the ApsaraDB RDS instance. For more information about how to configure a whitelist, see [Configure an IP address whitelist](#).
2. Start the pgAdmin 4 client.

Note For information about how to download the pgAdmin 4 client, visit [pgAdmin 4 \(Windows\)](#).

3. Right-click **Servers** and choose **Create > Server**, as shown in the following figure.

- On the **General** tab of the **Create - Server** dialog box, enter the name of the server, as shown in the following figure.
- Click the **Connection** tab and enter the information of the instance, as shown in the following figure.

| Parameter | Description |
|-------------------|---|
| Host name/address | The internal endpoint of the ApsaraDB RDS instance. For more information about how to view the internal endpoint, see View and modify the internal endpoint and port number . |
| Port | The internal port number that is used to connect to the ApsaraDB RDS instance. For more information about how to view the internal port number, see View and modify the internal endpoint and port number . |
| Username | The name of the privileged account on the ApsaraDB RDS instance. For more information about how to obtain a privileged account, see Create a database and an account . |
| Password | The password of the privileged account of the ApsaraDB RDS instance. |

- Click **Save**.
- If the connection information is correct, choose **Servers > Server Name > Databases > postgres**. If the following page appears, the connection is established.

 **Notice** The postgres database is the default system database of the ApsaraDB RDS instance. Do not perform operations on this database.

10.6.2. Use DMS to log on to an ApsaraDB RDS instance

This topic describes how to use Data Management (DMS) to log on to an ApsaraDB RDS instance.

Prerequisites

The IP address whitelist is configured. For more information about how to configure an IP address whitelist, see [Configure an IP address whitelist](#).

Procedure

- [Log on to the ApsaraDB RDS console](#).
- On the **Instances** page, find the target instance.
- Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
- Click **Log On to DB** in the upper-right corner of the page.
- In the **Login instance** dialog box of the **DMS** console, check values of **Database type**, **Instance Area**, and **Connection string address**. If the information is correct, enter **Database account** and **Database password**, as shown in the following figure.

| Parameter | Description |
|---------------------------|---|
| Database type | The engine of the database. By default, the engine of the database to be connected is displayed. |
| Instance Area | The region where the instance is deployed. By default, the region of the current instance is displayed. |
| Connection string address | The endpoint of the instance. By default, the endpoint of the current instance is displayed. |
| Database account | The account of the database to be connected. |
| Database password | The password of the account used to connect to the database. |

6. Click **Login**.

Note If you want the browser to remember the password, select **Remember password** before you click **Login**.

10.6.3. View and modify the internal endpoint and port number

You must use the internal endpoint and port number to access an ApsaraDB RDS instance. This topic describes how to view and modify the internal endpoint and port number of an ApsaraDB RDS instance in the ApsaraDB RDS console.

View the internal endpoint and port number

1. [Log on to the ApsaraDB RDS console](#).

2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the **Basic Information** section, view the internal endpoint and port number of the instance.

Modify the internal endpoint and port number

1. [Log on to the ApsaraDB RDS console.](#)
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Connection**.
5. On the right side of the page, click **Change Endpoint**.
6. In the dialog box that appears, set **Connection Type** to **Internal Endpoint**.
7. Modify the endpoint prefix and port number and click **OK**.

FAQ

- Q: Do I need to modify the endpoint or port number in my application after I modify the endpoint or port number of an instance?

A: Yes, you must modify the endpoint or port number in the application after you modify them. Otherwise, the application cannot connect to databases of the instance.

- Q: Does the modification of the endpoint take effect immediately? Do I need to restart the instance?

A: No, you do not need to restart the instance. The modification takes effect immediately.

10.7. Accounts

10.7.1. Create an account

Before you start to use ApsaraDB RDS, you must create an account on an ApsaraDB RDS instance. This topic describes how to create an account on an ApsaraDB RDS for PostgreSQL instance.

Account types

ApsaraDB RDS for PostgreSQL instances support two types of accounts: privileged accounts and standard accounts. The following table describes these account types.

| Account type | Description |
|---------------------------|---|
| Privileged account | <ul style="list-style-type: none">• You can create and manage privileged accounts only by using the ApsaraDB RDS console or the API.• If your ApsaraDB RDS instance uses local SSDs, you can create only a single privileged account. If your ApsaraDB RDS instance uses standard or enhanced SSDs, you can create more than one privileged account. A privileged account allows you to manage all the standard accounts and databases that are created on your ApsaraDB RDS instance.• A privileged account has more permissions that allow you to manage your ApsaraDB RDS instance at more fine-grained levels. For example, you can grant the query permissions on different tables to different users.• A privileged account has the permissions to disconnect accounts that are created on your ApsaraDB RDS instance. |

| Account type | Description |
|-------------------------|---|
| Standard account | <ul style="list-style-type: none"> You can create and manage standard accounts by using the ApsaraDB RDS console, API operations, or SQL statements. You can create more than one standard account on your ApsaraDB RDS instance. You must grant the permissions on specific databases to a standard account. A standard account does not have the permissions to create, manage, or disconnect other accounts on your ApsaraDB RDS instance. |

Precautions

- If your ApsaraDB RDS instance uses local SSDs, you can create one privileged account in the ApsaraDB RDS console. After the privileged account is created, it cannot be deleted. You can also create and manage more than one standard account by using SQL statements.
- If your ApsaraDB RDS instance uses standard or enhanced SSDs, you can create more than one privileged account and standard account in the ApsaraDB RDS console. You can also create and manage more than one standard account by using SQL statements.
- To migrate data from an on-premises database to your ApsaraDB RDS instance, you must create a database and an account on the ApsaraDB RDS instance. Make sure that the created database has the same properties as the on-premises database. Also make sure that the created account has the same permissions on the created database as the account that is authorized to manage the on-premises database.
- Follow the least privilege principle to create accounts and grant them read-only permissions or read and write permissions on databases. If necessary, you can create more than one account and grant them only the permissions on specific databases. If an account does not need to write data to a database, grant only the read-only permissions on that database to the account.
- For security purposes, we recommend that you specify strong passwords for the accounts on your ApsaraDB RDS instance and change the passwords on a regular basis.

Create a privileged account on an instance that uses local SSDs

- Log on to the [ApsaraDB RDS console](#).
- On the **Instances** page, find the target instance.
- Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
- In the left-side navigation pane, click **Accounts**.
- On the Accounts page, click **Create Privileged Account** and configure the following parameters.

| Parameter | Description |
|--------------------------|---|
| Database Account | <ul style="list-style-type: none"> The name of the account must be 2 to 16 characters in length. The name can contain lowercase letters, digits, and underscores (_). The name must start with a letter and end with a letter or digit. |
| Password | <ul style="list-style-type: none"> The password of the account must be 8 to 32 characters in length. The password must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters.characters. Special characters include !@#%&^&*()_+ -= |
| Re-enter Password | Enter the password of the account again. |

6. Click **Create**.

Create a privileged or standard account on an instance that uses standard or enhanced SSDs

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Accounts**.
5. On the Accounts page, click **Create Account** and configure the following parameters.

| Parameter | Description |
|--------------------------|---|
| Database Account | <ul style="list-style-type: none"> ◦ The name of the account must be 2 to 16 characters in length. ◦ The name can contain lowercase letters, digits, and underscores (_). ◦ The name must start with a letter and end with a letter or digit. |
| Account Type | Select Privileged Account or Standard Account . |
| Password | <ul style="list-style-type: none"> ◦ The password of the account must be 8 to 32 characters in length. ◦ The password must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters.characters. ◦ Special characters include !@#\$\$%^&*()_+ -= |
| Re-enter Password | Enter the password of the account again. |
| Description | This parameter is optional. You can enter relevant description to make the instance identifiable. The description can be up to 256 characters in length. |

6. Click **Create**.

Create a standard account on an instance that uses local SSDs

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. Click **Log On to DB** in the upper-right corner of the page.
5. In the **Login instance** dialog box of the **DMS** console, check values of **Database type**, **Instance Area**, and **Connection string address**. If the information is correct, enter **Database account** and **Database password**, as shown in the following figure.

| Parameter | Description |
|----------------------------------|---|
| Database type | The engine of the database. By default, the engine of the database to be connected is displayed. |
| Instance Area | The region where the instance is deployed. By default, the region of the current instance is displayed. |
| Connection string address | The endpoint of the instance. By default, the endpoint of the current instance is displayed. |
| Database account | The account of the database to be connected. |
| Database password | The password of the account used to connect to the database. |

- Click **Login**. If you want the browser to remember the password, select **Remember password** before you click **Login**.

Note If a message prompts you that the connection fails, the problem may be caused by an improperly configured whitelist. Reconfigure the whitelist in the console. For more information, see [Configure an IP address whitelist](#).

- The **SQLConsole** page appears after you log on to the instance. Execute a statement in the following format to create a standard account:

```
CREATE USER name [ [ WITH ] option [ ... ] ]
where option can be:
  SUPERUSER | NOSUPERUSER
  | CREATEDB | NOCREATEDB
  | CREATEROLE | NOCREATEROLE
  | CREATEUSER | NOCREATEUSER
  | INHERIT | NOINHERIT
  | LOGIN | NOLOGIN
  | REPLICATION | NOREPLICATION
  | CONNECTION LIMIT connlimit
  | [ ENCRYPTED | UNENCRYPTED ] PASSWORD 'password'
  | VALID UNTIL 'timestamp'
  | IN ROLE role_name [, ...]
  | IN GROUP role_name [, ...]
  | ROLE role_name [, ...]
  | ADMIN role_name [, ...]
  | USER role_name [, ...]
  | SYSID uid
```

For example, if you want to create a user account named test2 whose password is 123456, execute the following statement:

```
create user test2 password '123456';
```

8. Click **execute**.

10.7.2. Reset the password

This topic describes how to use the ApsaraDB RDS console to reset the password of your database account if you forget the password.

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Accounts**.
5. In the **Actions** column corresponding to the account, click **Reset Password**.
6. In the dialog box that appears, enter a new password and click **OK**.

-  **Note** The password must meet the following requirements:
- The password must be 8 to 32 characters in length.
 - The password must contain at least three of the following characters: uppercase letters, lowercase letters, digits, and special characters.
 - Special characters include ! @ # \$ % ^ & * () _ + - =

10.8. Databases

10.8.1. Create a database

Before you start to use ApsaraDB RDS, you must create a database on an ApsaraDB RDS instance. This topic describes how to create a database on an ApsaraDB RDS for PostgreSQL instance.

Prerequisites

- An ApsaraDB RDS for PostgreSQL instance is created. For more information, see [Create an instance](#).
- An account is created. For more information, see [Create an account](#).

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. Click **Log On to DB** in the upper-right corner of the page.
5. In the **Login instance** dialog box of the DMS console, check values of **Database type**, **Instance Area**, and **Connection string address**. If the information is correct, enter **Database account** and **Database password**, as shown in the following figure.

| Parameter | Description |
|----------------------------------|---|
| Database type | The engine of the database. By default, the engine of the database to be connected is displayed. |
| Instance Area | The region where the instance is deployed. By default, the region of the current instance is displayed. |
| Connection string address | The endpoint of the instance. By default, the endpoint of the current instance is displayed. |
| Database account | The account of the database to be connected. |
| Database password | The password of the account used to connect to the database. |

6. Click **Login**. If you want the browser to remember the password, select **Remember password** before you click

Login.

Note If a message prompts you that the connection fails, the problem may be caused by an improperly configured whitelist. Reconfigure the whitelist in the console. For more information, see [Configure an IP address whitelist](#).

7. The **SQLConsole** page appears after you log on to the instance. Execute a statement in the following format to create a database:

```
CREATE DATABASE name;
```

For example, if you want to create a database named test, execute the following statement:

```
create database test;
```

8. Click **execute**.

10.8.2. Delete a database

This topic describes how to delete a database in the ApsaraDB RDS for PostgreSQL console.

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. Click **Log On to DB** in the upper-right corner of the page.
5. In the **Login instance** dialog box of the **DMS** console, check values of **Database type**, **Instance Area**, and **Connection string address**. If the information is correct, enter **Database account** and **Database password**, as shown in the following figure.

The screenshot shows a 'Login instance' dialog box with the following fields and controls:

- Database type**: A dropdown menu with a red box around it.
- Instance Area**: A dropdown menu with a red box around it.
- Connection string address**: A text input field with a red box around it.
- Database account**: A text input field with the placeholder text 'Please enter a database account'.
- Database password**: A text input field.
- Remember password**: A checkbox with a help icon.
- Buttons**: 'Test connection', 'Login', and 'Cancel'.

| Parameter | Description |
|---------------------------|---|
| Database type | The engine of the database. By default, the engine of the database to be connected is displayed. |
| Instance Area | The region where the instance is deployed. By default, the region of the current instance is displayed. |
| Connection string address | The endpoint of the instance. By default, the endpoint of the current instance is displayed. |
| Database account | The account of the database to be connected. |
| Database password | The password of the account used to connect to the database. |

- Click **Login**. If you want the browser to remember the password, select **Remember password** before you click **Login**.

Note If a message prompts you that the connection fails, the problem may be caused by an improperly configured whitelist. Reconfigure the whitelist in the console. For more information, see [Configure an IP address whitelist](#).

- The **SQLConsole** page appears after you log on to the instance. Execute the following statement to delete a database:

```
drop database <database name>;
```

For example, if you want to delete a database named test2, execute the following statement:

```
drop database test2;
```

- Click **execute**.

10.9. Networks, VPCs, and vSwitches

10.9.1. Change the network type of an ApsaraDB RDS for PostgreSQL instance

This topic describes how to change the network type of an ApsaraDB RDS for PostgreSQL instance between classic network and Virtual Private Cloud (VPC).

Prerequisites

ApsaraDB RDS instances use local SSDs.

Context

- Classic network:** ApsaraDB RDS instances in the classic network are not isolated. You can block unauthorized access only by configuring IP address whitelists on these instances.
- VPC:** Each VPC is an isolated network. We recommend that you use the VPC network type because it provides a higher security level.

You can configure route tables, CIDR blocks, and gateways in a VPC. To smoothly migrate applications to the cloud, you can use leased lines or VPNs to connect your self-managed data center to a VPC to create a virtual data center.

Change the network type from VPC to classic network

Precautions

- The ApsaraDB RDS instance must be in a VPC.
- After you change the network type from VPC to classic network, the internal endpoint of the ApsaraDB RDS instance remains unchanged. However, the IP address that is associated with the internal endpoint changes.
- After you change the network type from VPC to classic network, you cannot connect Elastic Compute Service (ECS) instances deployed in VPCs to the ApsaraDB RDS instance by using the internal endpoint. You must update the endpoint for the applications deployed on the ECS instances.
- When you change the network type, a 30-second network interruption may occur. To avoid business interruption, change the network type during off-peak hours or make sure that your applications are configured with automatic reconnection policies.

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Connection**.
5. In the upper-right corner of the Database Connection section, click **Switch to Classic Network**.
6. In the dialog box that appears, click **OK**.

 **Note** After the network type is changed to classic network, only ECS instances within the classic network can connect to the ApsaraDB RDS instance by using the internal endpoint. You must configure the internal endpoint for the ECS instances.

7. Configure a whitelist to allow ECS instances within the classic network to connect to the ApsaraDB RDS instance by using the internal endpoint.

 **Note**

- If the network isolation mode of the ApsaraDB RDS instance is standard whitelist, add the internal IP addresses of the ECS instances to a whitelist of your ApsaraDB RDS instance.
- If the network isolation mode of the ApsaraDB RDS instance is enhanced whitelist, add the internal IP addresses of the ECS instances to a classic network whitelist. If no classic network whitelists are available, create a whitelist. For more information about the enhanced whitelist mode, see [Switch to the enhanced whitelist mode](#).

Change the network type from classic network to VPC

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Connection**.
5. In the upper-right corner of the Database Connection section, click **Switch to VPC**.
6. In the Switch to VPC dialog box, select a VPC and vSwitch and specify whether to retain the classic network endpoint.

Note

- Select a VPC. We recommend that you select the VPC where your ECS instances are deployed. Otherwise, the ECS instances cannot communicate with the ApsaraDB RDS instance over the internal network.
- Select a vSwitch. If no vSwitches are available in the selected VPC, create one in the same zone where the ApsaraDB RDS instance is deployed. For more information, see [Create a vSwitch in Quick Start of VPC User Guide](#).
- Determine whether to select the **Reserve Original Classic Network Endpoint** option. The following table describes the details.

■ Not selected

The classic network endpoint is not retained, and the endpoint of the instance changes to a VPC endpoint.

When you change the network type, a 30-second network interruption may occur, and connections between ECS instances in the classic network and the ApsaraDB RDS instance are interrupted.

■ Selected

The classic network endpoint is retained, and a new VPC endpoint is generated. In such cases, the ApsaraDB RDS instance runs in hybrid access mode. ECS instances in both the classic network and a VPC can connect to the ApsaraDB RDS instance over the internal network. You must set **Expiration Time (Important)** to **14 Days Later**, **30 Days Later**, **60 Days Later**, or **120 Days Later** for the classic network. You can also modify the expiration time after the network type is changed. For more information, see [Hybrid network access mode](#).

When you change the network type, no network interruptions occur. Connections between ECS instances in the classic network and the ApsaraDB RDS instance remain available until the classic network endpoint expires.

To migrate your business to the VPC without interruption, you must add the VPC endpoint to access the ECS instances before the classic network endpoint expires. Seven days before the classic network endpoint expires, the system sends a text message to the phone number bound to your Apsara Stack account every day.

For more information, see [Hybrid access from both the classic network and VPCs](#).

7. Add the internal IP addresses of ECS instances in the selected VPC to a VPC whitelist. This allows the ECS instances to access the ApsaraDB RDS instance over the internal network. If no VPC whitelists are available, create a whitelist.

Note

- If you retain the classic network endpoint, add the VPC endpoint to the ECS instances before the classic network endpoint expires.
- If you do not retain the classic network endpoint, connections between ECS instances in the classic network and the ApsaraDB RDS instance over the internal network are interrupted. You must add the new endpoint to ECS instances in the VPC immediately after the network type is changed.

10.9.2. Configure hybrid access from both the classic network and VPCs

This topic describes how to use the hybrid access solution of ApsaraDB RDS to change the network type of an instance from classic network to Virtual Private Cloud (VPC) without network interruptions.

Prerequisites

- The ApsaraDB RDS instance uses local SSDs.
- The ApsaraDB RDS instance is deployed in the classic network.
- Available VPCs and vSwitches exist in the zone where the ApsaraDB RDS instance is deployed.

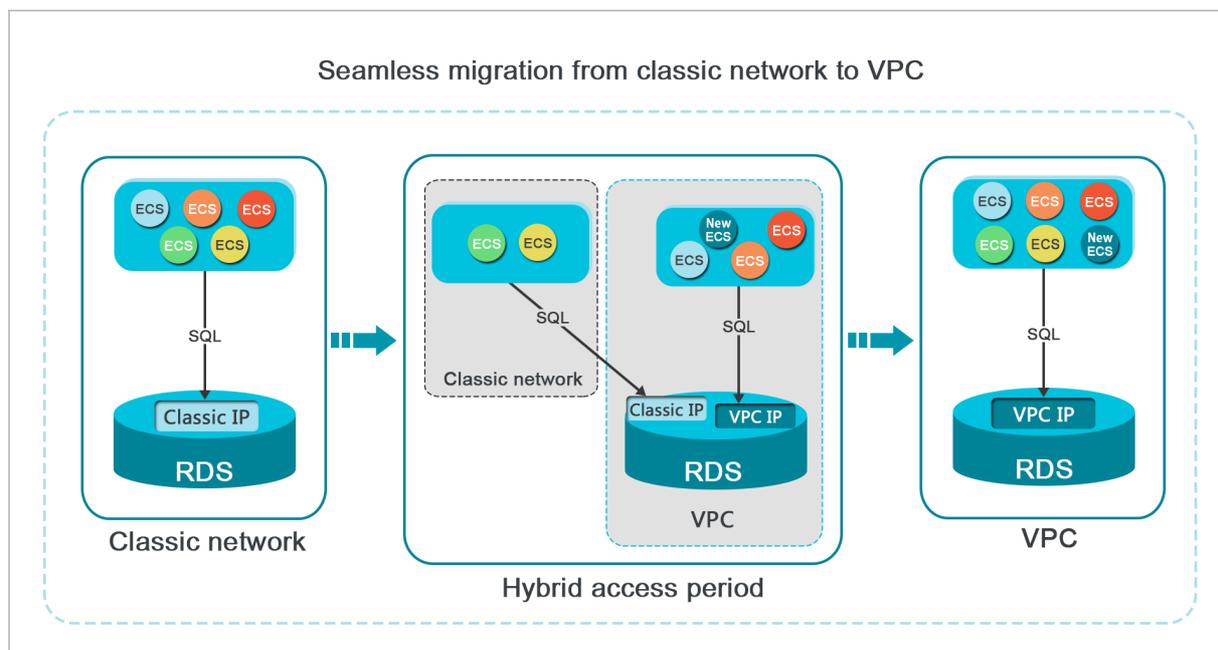
Context

In the past, when you change the network type of an ApsaraDB RDS instance from classic network to VPC, the internal endpoint of the ApsaraDB RDS instance would remain the same but the IP address bound to the endpoint would change to the corresponding IP address in the VPC. This change would cause a 30-second network interruption, and ECS instances within the classic network would not be able to access the ApsaraDB RDS instance by using the internal endpoint within this period. To smoothly change the network type, ApsaraDB RDS provides the hybrid access solution.

Hybrid access refers to the ability of an ApsaraDB RDS instance to be accessed by ECS instances in both the classic network and VPCs. During the hybrid access period, the ApsaraDB RDS instance reserves the original internal endpoint of the classic network and adds the internal endpoint of VPCs. This prevents network interruptions during the network type switchover.

For better security and performance, we recommend that you use the internal endpoint of VPCs. Hybrid access is available for a limited period of time. The internal endpoint of the classic network is released when the hybrid access period expires. In that case, your applications cannot access the ApsaraDB RDS database by using the internal endpoint of the classic network. You must configure the internal endpoint of VPCs in all your applications during the hybrid access period. This ensures smooth network switchover and minimizes the impact on your services.

For example, your company wants to use the hybrid access solution to change the network type from classic network to VPC. During the hybrid access period, some applications can access the database by using the internal endpoint of VPCs, and the other applications can access the database by using the original internal endpoint of the classic network. When all the applications access the database by using the internal endpoint of VPCs, the internal endpoint of the classic network can be released. The following figure illustrates the scenario.



Limits

During the hybrid access period, the instance has the following limits:

- The network type of your instance cannot be changed to classic network.
- Your instance cannot be migrated to another zone.

Change the network type from classic network to VPC

For more information, see [Change the network type from classic network to VPC](#).

Change the expiration time for the original internal endpoint of the classic network

During the period in which your instance can be accessed over the classic network or VPCs, you can specify the expiration time for the endpoint of the classic network. The setting takes effect immediately. For example, if the endpoint of the classic network is about to expire on August 18, 2017 and you change the expiration time to 14 days later on August 15, 2017, the endpoint of the classic network is released on August 29, 2017.

To change the expiration time, perform the following steps:

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Database Connection**.
5. On the **Instance Connection** tab, click **Change Expiration Time**.
6. In the **Change Expiration Time** dialog box, select an expiration time and click **OK**.

10.10. Monitoring

10.10.1. View monitored resources

ApsaraDB RDS provides a wide range of performance metrics. This topic describes how to view resource monitoring data in the ApsaraDB RDS console.

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Monitoring and Alerts**.
5. On the **Monitoring** tab, select a time range to query the corresponding monitoring data. The following table lists the specific metrics.

| Metric | Description |
|-------------------|--|
| Disk Space | The used disk space of the instance. Unit: MB. |
| IOPS | The number of I/O requests of the data and log disks per second. |
| Memory Usage | The memory usage of the instance. Unit: %. |
| CPU Utilization | The CPU utilization of the instance. Unit: %. |
| Total Connections | The total number of current connections of the instance. |

 **Note** You can click **Refresh** in the upper-right corner of the **Monitoring** tab to refresh the monitoring information.

10.10.2. Set a monitoring frequency

This topic describes how to set the monitoring frequency of an ApsaraDB RDS for PostgreSQL instance.

Context

ApsaraDB RDS for PostgreSQL provides the following monitoring frequencies:

- Every 5 seconds
- Every 60 seconds
- Every 300 seconds

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Monitoring and Alerts**.
5. On the **Monitoring** tab, click **Set Monitoring Frequency**.
6. In the **Set Monitoring Frequency** dialog box, select a monitoring frequency and click **OK**.

10.11. Data security

10.11.1. Switch to the enhanced whitelist mode

This topic describes how to switch from the standard whitelist mode to the enhanced whitelist mode for an ApsaraDB RDS instance. The enhanced whitelist mode provides higher security.

Network isolation modes

ApsaraDB RDS instances support the following network isolation modes:

- **Standard whitelist mode**
IP addresses from both the classic network and VPCs are added to the same IP address whitelist. However, the standard whitelist mode may incur security risks. Therefore, we recommend that you switch the network isolation mode to enhanced whitelist.
- **Enhanced whitelist mode**
IP addresses from the classic network and VPCs are added to different IP address whitelists. When you create an enhanced IP address whitelist, you must specify its network type.

Changes after you switch to the enhanced whitelist mode

- If the network type of the instance is VPC, the system generates a new whitelist that contains the same IP addresses as the original whitelists. The new IP address whitelist applies only to access from VPCs.
- If the network type of the instance is classic network, the system generates a new whitelist that contains the same IP addresses as the original whitelists. The new IP whitelist applies only to access from the classic network.
- If the instance supports [access from both the classic network and VPCs](#), two new IP address whitelists are created, and each contains the same IP addresses as the original whitelists. One whitelist applies to access from VPCs, and the other applies to access from the classic network.

 **Note** After you switch to the enhanced whitelist mode, the configured ECS instance groups remain unchanged.

Precautions

- You can switch from the standard whitelist mode to the enhanced whitelist mode, but not the other way around.
- In enhanced whitelist mode, a classic network whitelist also allows access from the Internet. If you want to access the instance from a host over the Internet, you can add the public IP address of the host to a classic network whitelist.

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. On the **Whitelist Settings** tab, click **Switch to Enhanced Whitelist (Recommended)**.
6. In the message that appears, click **Confirm**.

10.11.2. Configure an IP address whitelist

This topic describes how to configure a whitelist for an ApsaraDB RDS instance. Only entities that are listed in a whitelist can access your ApsaraDB RDS instance.

Context

Whitelists make your ApsaraDB RDS instance more secure and do not interrupt the operations of your ApsaraDB RDS instance when you configure whitelists. We recommend that you perform maintenance on your whitelists on a regular basis.

To configure a whitelist, perform the following operations:

- Configure a whitelist: Add IP addresses to allow them to connect to the ApsaraDB RDS instance.

 **Note** The IP address whitelist labeled **default** contains only the default IP address 0.0.0.0/0, which allows all entities to access your ApsaraDB RDS instance.

- Configure an ECS security group: Add an ECS security group for the ApsaraDB RDS instance to allow ECS instances in the group to connect to the ApsaraDB RDS instance.

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. On the **Whitelist Settings** tab, click **Edit** corresponding to the **default** whitelist.

 **Note** You can also click **Create Whitelist** to create a whitelist.

6. In the **Edit Whitelist** dialog box, enter the IP addresses or CIDR blocks used to access the instance and click

OK. The following section describes the rules:

- If you enter the CIDR block 10.10.10.0/24 in the IP Addresses field, all IP addresses in the 10.10.10.X format can access your ApsaraDB RDS instance.
- If you enter more than one IP address or CIDR block, you must separate them with commas (.). Do not add spaces before or after the commas. Example: 192.168.0.1,172.16.213.9.
- If you click **Add Internal IP Addresses of ECS Instances**, the IP addresses of all ECS instances created within your Alibaba Cloud account are displayed. You can select the required IP addresses to add them to the IP address whitelist.

10.11.3. Configure SSL encryption

This topic describes how to enable SSL encryption for an ApsaraDB RDS instance.

Prerequisites

The ApsaraDB RDS instance uses standard SSDs.

Precautions

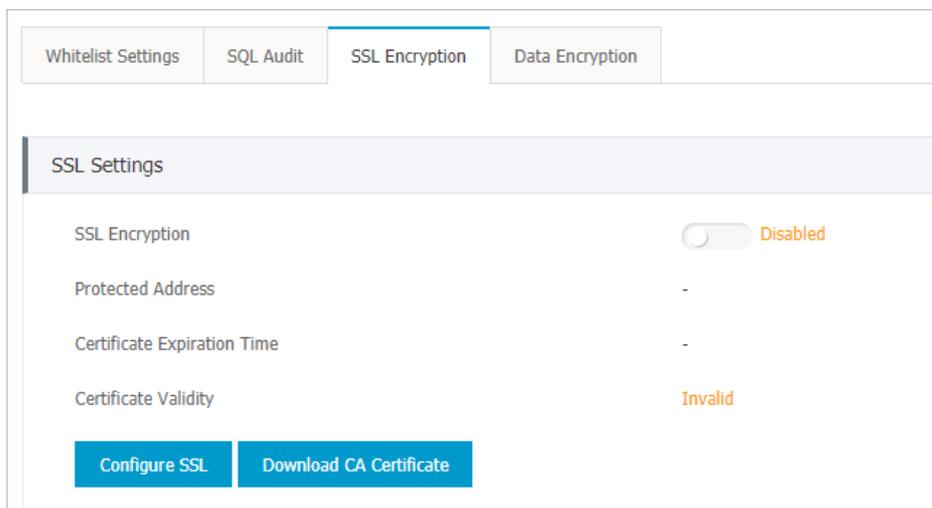
- After SSL encryption is enabled, data transmitted over an internal network or the Internet is encrypted by using SSL. SSL encryption protects data from leaks.
- After SSL encryption is enabled, you must close the existing connection and establish a new one to bring SSL encryption into effect.

Enable SSL encryption

SSL 3.0 has been upgraded by the Internet Engineering Task Force (IETF) to Transport Layer Security (TLS), but the term SSL encryption is still commonly used in the communications industry. Therefore, SSL encryption is used in this topic to refer to TLS encryption.

 **Note** ApsaraDB RDS supports TLS 1.0, 1.1, and 1.2.

1. [Log on to the ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. Click the **SSL Encryption** tab.

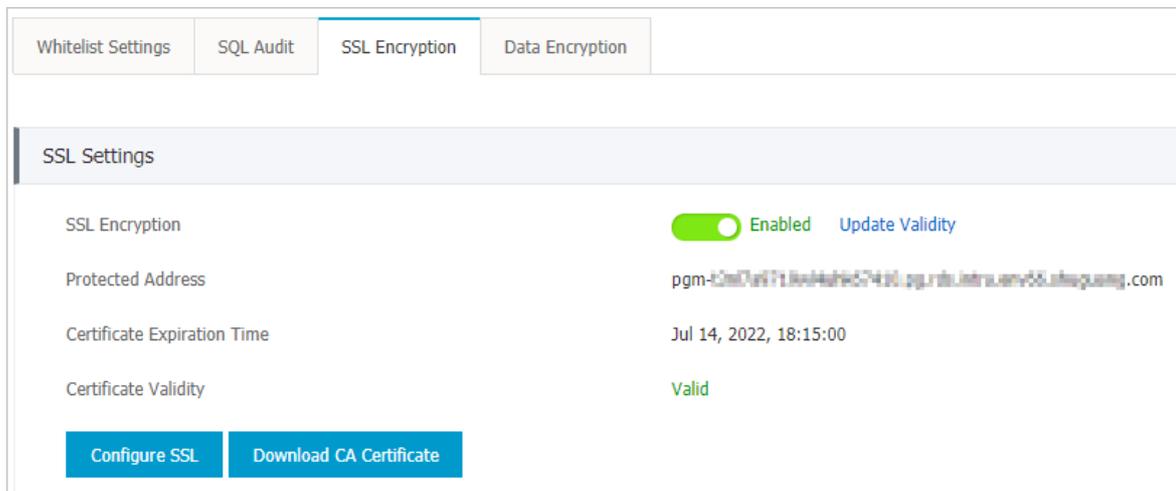


6. Click **Enable SSL**.

Note After SSL encryption is enabled, you must set the SSL mode to **Prefer** when you log on from your client.

Disable SSL encryption

1. Log on to the [ApsaraDB for RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the **Actions** column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. Click the **SSL Encryption** tab.



6. Click **Disable SSL**.

10.11.4. Configure data encryption

This topic describes how to configure data encryption for an ApsaraDB RDS instance that uses standard or enhanced SSDs. The disk encryption feature maximizes the protection for your data and eliminates the need to modify business or application configurations. ApsaraDB RDS automatically applies disk encryption to both the snapshots that are generated from the encrypted SSDs and the SSDs that are created from those snapshots.

Prerequisites

The storage type of the instance is standard SSD.

Configure disk encryption

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

Note When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top navigation bar, choose **Products > Security > Key Management Service**.
5. On the Keys page, click **Create Key**.
6. Configure the following parameters.

| Section | Parameter | Description |
|-----------------------|-------------------------|--|
| Region | Organization | The organization to which the key belongs. |
| | Resource Set | The resource set to which the key belongs. |
| | Region | The region to which the key belongs. |
| Basic Settings | Key Type | KMS supports the following key types: <ul style="list-style-type: none"> ○ Symmetric keys: <ul style="list-style-type: none"> ■ Aliyun_AES_256 ■ Aliyun_SM4 ○ Asymmetric keys: <ul style="list-style-type: none"> ■ RSA_2048 ■ EC_P256 ■ EC_P256K ■ EC_SM2 |
| | Key Purpose | ENCRYPT/DECRYPT: The purpose of the CMK is to encrypt or decrypt data. |
| | Protection Level | <ul style="list-style-type: none"> ○ SOFTWARE: Use a software module to protect the CMK. ○ HSM: Host the CMK in a hardware security module (HSM). Managed HSM uses the HSM as dedicated hardware to safeguard the CMK. |
| | Alias | The identifier of the CMK. For more information, see <i>Use aliases in KMS User Guide</i> . |
| | Description | The description of the CMK. |

| Section | Parameter | Description |
|-------------------|---------------------|---|
| Advanced Settings | Rotation Period | <p>Specifies whether to enable automatic rotation. If you choose to enable automatic rotation, you must select a rotation period. For more information about rotation, see <i>Key rotation in KMS User Guide</i>. Valid values:</p> <ul style="list-style-type: none"> 30 Days 90 Days 180 Days 365 Days Custom: Customize a period that ranges from 7 to 730 days. <p>Note You can specify this parameter only if Key Type is set to Aliyun_AES_256 or Aliyun_SM4.</p> |
| | Key Material Source | <p>The source of key material.</p> <ul style="list-style-type: none"> Key Management Service: Use KMS to generate key material. External: Manually import external key material. <p>Note If Rotation Period is set to Enable, the External option is unavailable.</p> |

- Click **Submit**.
- Create an ApsaraDB RDS instance with disk encryption enabled. For more information, see [Create an ApsaraDB RDS for PostgreSQL instance that uses standard or enhanced SSDs](#).

10.12. Logs and audit

10.12.1. Configure SQL audit

This topic describes how to configure the SQL audit feature to audit SQL executions and check the details. SQL audit does not affect instance performance.

Precautions

- SQL audit does not affect instance performance.
- SQL audit logs are retained for 30 days.
- Log files exported from SQL audit are retained for two days. The system deletes files that are retained for longer than two days.
- SQL audit is disabled by default. You must manually enable it.
- You cannot view logs that are generated before SQL audit is enabled.

Enable SQL audit

- [Log on to the ApsaraDB RDS console](#).
- On the **Instances** page, find the target instance.
- Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
- In the left-side navigation pane, click **Data Security**.
- Click the **SQL Audit** tab.

6. Click **Enable SQL Audit** or **Enable now**.
7. In the message that appears, click **Confirm**.

 **Note** After SQL audit is enabled, you can query SQL information based on conditions such as the time, database, user, and keyword.

Disable SQL audit

You can disable SQL audit when it is no longer needed. To disable SQL audit, perform the following steps:

 **Notice** After SQL audit is disabled, all SQL audit logs are deleted. We recommend that you export and store audit logs to your computer before you disable SQL audit.

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Data Security**.
5. Click the **SQL Audit** tab. Click **Export File**.

 **Note** If more than 1 million SQL audit logs meet the filter conditions you specify, only 1 million logs can be exported. SQL audit logs are exported at a speed of 900 entries per second. It takes about 20 minutes to export 1 million SQL audit logs.

6. Click **Files**. Find a file and click **Download** in the **Action** column to download the file to your computer.
7. Click **Disable SQL Audit**.
8. In the message that appears, click **Confirm**.

10.12.2. Manage logs

You can view logs for errors, slow queries, and primary/secondary instance switching for ApsaraDB RDS for PostgreSQL instances in the ApsaraDB RDS console or by executing SQL statements. These logs help you troubleshoot errors. This topic describes how to manage logs in the console.

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Logs**.
5. On the **Logs** page, click the **Error Logs**, **Slow Query Logs**, or **Primary/Secondary Switching Logs** tab, select a time range, and then click **Search**.

| Tab | Description |
|------------------------|--|
| Error Logs | Records database running errors that occurred within the last month. |
| Slow Query Logs | Records SQL statements within the last month that took longer than one second to execute. Duplicated SQL statements are removed. |

| Tab | Description |
|----------------------------------|--|
| Primary/Secondary Switching Logs | Records switchovers between the primary and secondary instances within the last month. |

10.13. Backup

10.13.1. Back up an ApsaraDB RDS for PostgreSQL instance

This topic describes how to back up an ApsaraDB RDS for PostgreSQL instance. You can configure a backup policy that is used to automatically back up your ApsaraDB RDS instance. If you do not configure a backup policy, the default backup policy is used. You can also manually back up your ApsaraDB RDS instance.

Precautions

- Do not perform data definition language (DDL) operations during a backup. If you do so, the backup may fail due to table locks.
- We recommend that you back up your ApsaraDB RDS instance during off-peak hours.
- If the amount of data is large, it may take a long time to back up your ApsaraDB RDS instance.
- Backup files are retained for a specified retention period. We recommend that you download the required backup files to your computer before they are deleted.

Backup description

ApsaraDB RDS for PostgreSQL allows you to perform full physical backup and back up archived redo log files of databases.

Configure a backup policy to automatically back up your ApsaraDB RDS instance

ApsaraDB RDS automatically backs up your instance based on the specified backup policy.

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Backup and Restoration**.
5. On the **Backup and Restoration** page, click the **Backup Settings** tab and click **Edit**.
6. In the dialog box that appears, configure the following parameters and click **OK**. The following table lists the parameters.

| Parameter | Description |
|-----------------------|--|
| Data Retention Period | The number of days for which you want to retain data backup files. Valid values: 7 to 730. Unit: days. Default value: 7. |
| Backup Cycle | The cycle to create backups. You can select one or more days of the week. <div style="border: 1px solid #ccc; background-color: #e0f2f1; padding: 5px;"> <p> Note To ensure data security, we recommend that you back up your ApsaraDB RDS instance at least twice a week.</p> </div> |

| Parameter | Description |
|----------------------|--|
| Backup Time | The period of time for which you want to back up data. Unit: hours. |
| Log Backup | Specifies whether to enable the log backup feature. <div style="border: 1px solid #ccc; background-color: #e0f2f1; padding: 5px; margin-top: 5px;">  Notice If you disable this feature, all log backup files are deleted and your instance cannot be restored to previous points in time. </div> |
| Log Retention Period | <ul style="list-style-type: none"> ◦ The period of time for which you want to retain log backup files. Valid values: 7 to 730. Unit: days. Default value: 7. ◦ The log retention period must be less than or equal to the data retention period. |

Manually back up your ApsaraDB RDS instance

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the upper-right corner of the page, click **Back Up Instance**. The **Back Up Instance** dialog box appears.
5. Select the backup mode and backup policy, and click **OK**.

 **Note** The backup mode is **Full Backup** and the backup policy is **Instance Backup**.

What's next

You can click the  icon in the upper-right corner of the page to view the task progress displayed in the **Task Progress** list.

10.13.2. Download data and log backup files

This topic describes how to download unencrypted data and log backup files in the ApsaraDB RDS console to archive the files and restore data to an on-premises database.

Procedure

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Backup and Restoration** to go to the **Backup and Restoration** page.
5. Click the **Data Backup** or **Archived Logs** tab.
 - To download data backup files, click the **Data Backup** tab.
 - To download log files, click the **Archived Logs** tab.
6. Select a time range to which you want to restore the instance.
7. Find the data backup or log file that you want to download and click **Download** in the **Actions** column.

Note

- If you want to use a data backup file to restore data, select the backup file that is the closest to the time for restoration.
- If you want to use a log file to restore data to an on-premises database, take note of the following items:
 - The instance No. of the log file must be the same as that of the data backup file.
 - The start time of the log file must be later than the data backup time and earlier than the time for restoration.

8. In the message that appears, select a download method.

| Download method | Description |
|------------------------|---|
| Download | Download the file by using the public endpoint. |
| Copy Internal Endpoint | Copy the internal endpoint to download the file. If your ECS and ApsaraDB RDS instances are deployed within the same region, you can log on to the ECS instance and use the internal endpoint to download the file. This method is fast and secure. |
| Copy Public Endpoint | Copy the public endpoint to download the file. If you want to use other tools to download the file, use the public endpoint. |

Note If you use a Linux operating system, you can run the following command to download the file:

```
wget -c '<Public endpoint of the backup file, which is the download URL>' -O <File name>
```

- The `-c` option enables resumable download.
- The `-O` option saves the downloaded file by using a specified name. We recommend that you use the file name contained in the download URL.
- If the URL contains more than one parameter, enclose the download URL in a pair of single quotation marks (').

```
root@iZbc...:~# wget -c 'http://rdslog-hz-...-cn-hangzhou.aliyuncs.com/.../hostins...mysql-bin.000457?OSSAccessKeyId=...' -O mysql-bin.000457
```

10.13.3. Create a logical backup for an ApsaraDB RDS for PostgreSQL instance

This topic describes how to use `pg_dump` to create a logical backup for an ApsaraDB RDS for PostgreSQL instance and export the backup file to your computer.

Context

The `pg_dump` utility provided with PostgreSQL is used to back up individual databases. For more information, visit [pg_dump](#).

In this example, an ApsaraDB RDS for PostgreSQL instance that runs Linux 7 and PostgreSQL 10 is used.

Prerequisites

- The IP address of your ECS instance or host is added to a whitelist of the ApsaraDB RDS for PostgreSQL instance.

For more information, see [Configure an IP address whitelist](#).

- Your ECS instance or host runs the same version of PostgreSQL as the ApsaraDB RDS for PostgreSQL instance.

Precautions

We recommend that you use the privileged account of the ApsaraDB RDS for PostgreSQL instance to ensure that you have all the required permissions.

Back up a database

- Log on to your ECS instance or host. Then, run the following command to back up a database from the ApsaraDB RDS for PostgreSQL instance:

```
pg_dump -h '<hostname>' -U <username> -p <port> -Fc <dbname> > <dumpdir>
```

| Parameter | Description |
|-----------|---|
| hostname | The endpoint of the ApsaraDB RDS for PostgreSQL instance. <div style="border: 1px solid #add8e6; padding: 5px;"> ? Note If your ECS instance connects to the ApsaraDB RDS for PostgreSQL instance by using an internal endpoint, make sure that the ECS and ApsaraDB RDS instances have the same network type. If both instances use the VPC network type, make sure that both instances are deployed within the same VPC. For more information about how to view the internal endpoint, see View and modify the internal endpoint and port number. </div> |
| username | The username of the privileged account of the ApsaraDB RDS for PostgreSQL instance. |
| port | The port number of the ApsaraDB RDS for PostgreSQL instance. |
| -Fc | The output file format. <code>-Fc</code> specifies the custom format, which is ideal when you use <code>pg_restore</code> to import logical backup files and restore databases. For more information, visit pg_dump . |
| dbname | The name of the database that you want to back up. |
| dumpdir | The directory and name of the logical backup file to export. |

Example:

```
pg_dump -h 'pgm-bpxxxxxx.pg.rds.aliyuncs.com' -U test123 -p 3433 -Fc testdb > /tmp/testdb.dump
```

- When `Password:` appears, enter the password of the privileged account of the ApsaraDB RDS for PostgreSQL instance and press the Enter key.

```

[root@iZbp... etc]# pg_dump -h 'pgm-... pg.rds.aliyuncs.com' -U test123 -p 3433 -Fc testdb > /tmp/testdb.dump
Password:
[root@iZbp... etc]# ll /tmp/testdb.dump
-rw-r--r-- 1 root root 2006 Nov  5 16:05 /tmp/testdb.dump
[root@iZbp... etc]#
    
```

Back up one or more tables

- Log on to your ECS instance or host. Then, run the following command to back up one or more tables from a database in the ApsaraDB RDS for PostgreSQL instance:

```
pg_dump -h '<hostname>' -U <username> -p <port> -t <table> -Fc <dbname> > <dumpdir>
```

| Parameter | Description |
|-----------|---|
| hostname | The endpoint of the ApsaraDB RDS for PostgreSQL instance. Note If your ECS instance connects to the ApsaraDB RDS for PostgreSQL instance by using an internal endpoint, make sure that the ECS and ApsaraDB RDS instances have the same network type. If both instances use the VPC network type, make sure that both instances are deployed within the same VPC. For more information about how to view the internal endpoint, see View and modify the internal endpoint and port number . |
| username | The username of the privileged account of the ApsaraDB RDS for PostgreSQL instance. |
| port | The port number of the ApsaraDB RDS for PostgreSQL instance. |
| table | The name of the table that you want to back up. You can use <code>-t <table></code> to specify more than one table. |
| -Fc | The output file format. <code>-Fc</code> specifies the custom format, which is ideal when you use <code>pg_restore</code> to import logical backup files and restore databases. For more information, visit pg_dump . |
| dbname | The name of the database that you want to back up. |
| dumpdir | The directory and name of the logical backup file to export. |

Example:

```
pg_dump -h 'pgm-bpxxxxx.pg.rds.aliyuncs.com' -U test123 -p 3433 -t products1 -Fc testdb2 > /tmp/testdb2.dump
```

- When `Password:` appears, enter the password of the privileged account of the ApsaraDB RDS for PostgreSQL instance and press the Enter key.

```
[root@iZ... ~]# pg_dump -h 'pgm-bp...pg.rds.aliyuncs.com' -U test123 -p 3433 -t products1 -Fc testdb2 > /tmp/testdb2.d
ump
Password:
[root@iZ... ~]#
```

Back up a database with one or more tables excluded

- Log on to your ECS instance or host. Then, run the following command to back up a database from the ApsaraDB RDS instance with one or more tables excluded:

```
pg_dump -h '<hostname>' -U <username> -p <port> -T <table> -Fc <dbname> > <dumpdir>
```

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|-----------|--|
| hostname | <p>The endpoint of the ApsaraDB RDS for PostgreSQL instance.</p> <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e0f0ff;"> <p>Note If your ECS instance connects to the ApsaraDB RDS for PostgreSQL instance by using an internal endpoint, make sure that the ECS and ApsaraDB RDS instances have the same network type. If both instances use the VPC network type, make sure that both instances are deployed within the same VPC. For more information about how to view the internal endpoint, see View and modify the internal endpoint and port number.</p> </div> |
| username | The username of the privileged account of the ApsaraDB RDS for PostgreSQL instance. |
| port | The port number of the ApsaraDB RDS for PostgreSQL instance. |
| table | The name of the table that you want to exclude. You can use <code>-T <table></code> to specify more than one table. |
| -Fc | The output file format. <code>-Fc</code> specifies the custom format, which is ideal when you use <code>pg_restore</code> to import logical backup files and restore databases. For more information, visit pg_dump . |
| dbname | The name of the database that you want to back up. |
| dumpdir | The directory and name of the logical backup file to export. |

Example:

```
pg_dump -h 'pgm-bpxxxx.pg.rds.aliyuncs.com' -U test123 -p 3433 -T products1 -Fc testdb2 > /tmp/testdb2.dump
```

- When `Password:` appears, enter the password of the privileged account of the ApsaraDB RDS for PostgreSQL instance and press the Enter key.



Back up the schema of a database with data excluded

- Log on to your ECS instance or host. Then, run the following command to back up the schema of a database from the ApsaraDB RDS for PostgreSQL instance:

```
pg_dump -h '<hostname>' -U <username> -p <port> -s -Fc <dbname> > <dumpdir>
```

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|-----------|---|
| hostname | The endpoint of the ApsaraDB RDS for PostgreSQL instance. Note If your ECS instance connects to the ApsaraDB RDS for PostgreSQL instance by using an internal endpoint, make sure that the ECS and ApsaraDB RDS instances have the same network type. If both instances use the VPC network type, make sure that both instances are deployed within the same VPC. For more information about how to view the internal endpoint, see View and modify the internal endpoint and port number . |
| username | The username of the privileged account of the ApsaraDB RDS for PostgreSQL instance. |
| port | The port number of the ApsaraDB RDS for PostgreSQL instance. |
| -s | Specifies whether to back up only the schema of the database. The data of the database is not backed up. For more information, visit pg_dump . |
| -Fc | The output file format. <code>-Fc</code> specifies the custom format, which is ideal when you use <code>pg_restore</code> to import logical backup files and restore databases. For more information, visit pg_dump . |
| dbname | The name of the database that you want to back up. |
| dumpdir | The directory and name of the logical backup file to export. |

Example:

```
pg_dump -h 'pgm-bpxxxx.pg.rds.aliyuncs.com' -U test123 -p 3433 -s -Fc testdb2 > /tmp/testdb2.dump
```

- When `Password:` appears, enter the password of the privileged account of the ApsaraDB RDS for PostgreSQL instance and press the Enter key.

```
[root@izb]# pg_dump -h 'pgm-bpxxxx.pg.rds.aliyuncs.com' -U test123 -p 3433 -s -Fc testdb2 > /tmp/testdb2.dump
Password:
[root@izb]# ll /tmp/
total 16
-rwxr-xr-x 1 root root 0 Nov 5 15:28 Aegis--
-rw-r--r-- 1 root root 4 Nov 5 15:27 CmsGoAgent.pid
drwx----- 3 root root 4096 Nov 5 15:27 systemd-private-
-rw-r--r-- 1 root root 2013 Nov 7 14:43 testdb2.dump
```

10.13.4. Create a full backup of an ApsaraDB RDS for PostgreSQL instance

This topic describes how to use the `pg_basebackup` utility provided by open source PostgreSQL to create a full backup of your ApsaraDB RDS for PostgreSQL instance and export the backup files to your computer.

Prerequisites

- The IP address of your ECS instance or host is added to a whitelist of your ApsaraDB RDS for PostgreSQL instance. For more information, see [Configure an IP address whitelist](#).
- Your ECS instance or host runs the same version of PostgreSQL as the ApsaraDB RDS for PostgreSQL instance.

Context

`pg_basebackup` backs up all data of a PostgreSQL instance. Backup files can be used for point-in-time recovery. For more information, visit [pg_basebackup](#).

In this example, CentOS 7 is used to create a full backup.

Precautions

We recommend that you use the privileged account of the ApsaraDB RDS for PostgreSQL instance to ensure that you have all the required permissions.

Procedure

Note `pg_basebackup` cannot back up a single database or database object. For more information about how to back up a single database or database object, see [Create a logical backup for an ApsaraDB RDS for PostgreSQL instance](#).

1. Log on to your ECS instance or host. Then, run the following command to back up a database from the ApsaraDB RDS for PostgreSQL instance:

```
pg_basebackup -Ft -Pv -Xf -z -D <backupdir> -Z5 -h '<hostname>' -p <port> -U <username> -W
```

The following table describes the parameters in this command. For more information, visit [pg_basebackup](#).

| Parameter | Description |
|-----------|--|
| backupdir | The directory of backup files that are exported. The system automatically creates this directory. However, if this directory already exists and is not empty, the system reports an error. |
| hostname | The internal endpoint of the ApsaraDB RDS for PostgreSQL instance. For more information about how to view the internal endpoint, see View and modify the internal endpoint and port number . |
| port | The port number of the ApsaraDB RDS for PostgreSQL instance. |
| username | A username of the ApsaraDB RDS for PostgreSQL instance. |

Example:

```
pg_basebackup -Ft -Pv -Xf -z -D /pg12/backup1/ -Z5 -h pgm-bpxxxxx.pg.rds.aliyuncs.com -p 1433 -U test1 -W
```

2. When **Password:** appears, enter the password of the username of the ApsaraDB RDS for PostgreSQL instance and press the Enter key.

```
[root@izbp-... ~]# pg_basebackup -Ft -Pv -Xf -z -D /pg12/backup/ -Z5 -h pgm-bpxxxxx.pg.rds.aliyuncs.com -p 1433 -U test1 -W
Password:
pg_basebackup: initiating base backup, waiting for checkpoint to complete
WARNING: skipping special file ".s.PGSQL.3002"
pg_basebackup: checkpoint completed
pg_basebackup: write-ahead log start point: 14/8F000028 on timeline 1
WARNING: skipping special file ".s.PGSQL.3002"/base.tar.gz
49965/49965 Kb (100%), 1/1 tablespaces
pg_basebackup: write-ahead log end point: 14/8F0003A0
pg_basebackup: syncing data to disk ...
pg_basebackup: base backup completed
[root@izbp-... jQZ ~]# ll /pg12/backup/
total 3956
-rw-r--r-- 1 root root 4047901 Apr 13 14:04 base.tar.gz
[root@izbp-... ~]#
```

10.14. Restoration

10.14.1. Restore data of an ApsaraDB RDS for PostgreSQL instance

This topic describes how to use the backup data of an ApsaraDB RDS for PostgreSQL instance to restore data.

Precautions

- The new instance must have the same whitelist, backup, and parameter settings as the original instance.
- The new instance must have the same data and account information as the backup set or instance at the time point.

Prerequisites

The original instance must meet the following requirements:

- The original instance is in the Running state and is not locked.
- The original instance does not have ongoing migration tasks.
- If you want to restore data to a point in time, the log backup feature is enabled for the original instance.
- If you want to restore an instance from a backup set, the original instance has at least one backup set.

Restore data of an ApsaraDB RDS for PostgreSQL instance

1. [Log on to the ApsaraDB RDS console](#).
2. On the **Instances** page, find the target instance.
3. Click the instance ID or click **Manage** in the Actions column corresponding to the instance to go to the **Basic Information** page.
4. In the left-side navigation pane, click **Backup and Restoration**.
5. In the upper-right corner of the page, click **Restore Database (Previously Clone Database)**.
6. Configure the following parameters.

| Section | Parameter | Description |
|-------------------------|------------------------|--|
| Region | Region | The region where the instance is deployed. |
| Restore Database | Restore Mode | <ul style="list-style-type: none"> ◦ By Time: You can restore data to a point in time within the retention period of the log backup. For more information about how to view or change the retention period of log backups, see Back up an ApsaraDB RDS for PostgreSQL instance. ◦ By Backup Set <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 5px;"> <p> Note The By Time option appears only when the log backup feature is enabled.</p> </div> |
| | Restore Time | The time to which the database is restored. This parameter is displayed when you set Restore Mode to By Time . |
| | Backup Set | The backup set used to restore the database. This parameter is displayed when you set Restore Mode to By Backup Set . |
| Specifications | Instance Name | The name of the instance. |
| | Database Engine | The engine of the database. The value of this parameter is set to PostgreSQL and cannot be changed. |
| | Engine Version | The version of the database engine. The value of this parameter is set to the engine version of the current instance and cannot be changed. |
| | Edition | The edition of the instance. |
| | Storage Type | The storage type of the instance. The value of this parameter is set to the storage type of the current instance and cannot be changed. |

| Section | Parameter | Description |
|---------------------|-------------------------|--|
| | Instance Type | The instance type of the instance. Memory size determines the maximum number of connections and IOPS. The actual values are displayed in the console. For more information, see Instance types in <i>Instance types of ApsaraDB RDS Product Introduction</i> . |
| | Storage Capacity | The storage capacity of the instance, including the space to store data, system files, binlog files, and transaction files. Valid values: 20 to 600. Unit: GB. The value must be in 1 GB increments. |
| Network Type | Network Type | <p>The network type of the instance. ApsaraDB RDS instances support the following network types:</p> <ul style="list-style-type: none"> ◦ Classic Network: Cloud services on a classic network are not isolated from each other. Unauthorized access to a cloud service is blocked only by the security group or whitelist policy of the service. ◦ VPC: A virtual private cloud (VPC) helps you build an isolated network environment on Apsara Stack. You can customize route tables, CIDR blocks, and gateways in a VPC. We recommend that you select VPC for improved security. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note If you set the network type to VPC, you must also select a VPC and a vSwitch.</p> </div> |

7. Click **Submit**.

10.14.2. Restore data from a logical backup file

This topic describes how to restore data from a logical backup file to an ApsaraDB RDS for PostgreSQL instance or an on-premises PostgreSQL database.

Context

A logical backup file is used to restore a small amount of data, such as data in a table. For a large amount of data, we recommend that you restore it from a full physical backup file to a new ApsaraDB RDS instance and then use Data Transmission Service (DTS) to migrate data to the original ApsaraDB RDS instance.

Prerequisites

Data in the ApsaraDB RDS for PostgreSQL instance is logically backed up. For more information, see [Create a logical backup for an ApsaraDB RDS for PostgreSQL instance](#).

Precautions

- We recommend that you do not restore data to the default postgres database.
- When you restore the data of a table, the system does not restore the database objects on which the table depends. The restoration may fail.

Restore the data of a database

1. Log on to the ECS instance or on-premises host that houses the logical backup file and run the following command to restore the data of a database:

```
pg_restore -h '<hostname>' -U <username> -p <port> -d <dbname> <dumpdir>
```

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|-----------|--|
| hostname | The endpoint of the ApsaraDB RDS for PostgreSQL instance. Note If your ECS instance connects to the ApsaraDB RDS for PostgreSQL instance by using an internal endpoint, make sure that the ECS and ApsaraDB RDS instances have the same network type. If both instances are of the VPC network type, make sure that both instances reside within the same VPC. For more information about how to view the internal endpoint, see View and modify the internal endpoint and port number . |
| username | The username of the privileged account of the ApsaraDB RDS for PostgreSQL instance. |
| port | The port number of the ApsaraDB RDS for PostgreSQL instance. |
| dbname | The name of the database whose data you want to restore. |
| dumpdir | The directory and name of the logical backup file to use. |

Example:

```
pg_restore -h 'pgm-bpxxxxx.pg.rds.aliyuncs.com' -U test123 -p 3433 -d testdb2 /tmp/testdb2.dump
```

- When **Password:** appears, enter the password of the privileged account of your ApsaraDB RDS instance and press the Enter key.

Note You can ignore alerts generated by the embedded plpgsql plug-in.

```
[root@izbj... ~]# pg_restore -h 'pgm-bp...pg.rds.aliyuncs.com' -U ... -p 3433 -d testdb4 /tmp/testdb2.dump
Password:
pg_restore: [archiver (db)] Error while PROCESSING TOC:
pg_restore: [archiver (db)] Error from TOC entry 3076; 0 0 COMMENT EXTENSION plpgsql
pg_restore: [archiver (db)] could not execute query: ERROR: must be owner of extension plpgsql
Command was: COMMENT ON EXTENSION plpgsql IS 'PL/pgSQL procedural language';
WARNING: errors ignored on restore: 1
```

Restore the data of a table

- Log on to the ECS instance or on-premises host that houses the logical backup file and run the following command to restore the data of a table:

```
pg_restore -h '<hostname>' -U <username> -p <port> -d <dbname> -t <table> -c <dumpdir>
```

| Parameter | Description |
|-----------|--|
| hostname | The endpoint of the ApsaraDB RDS for PostgreSQL instance. Note If your ECS instance connects to the ApsaraDB RDS for PostgreSQL instance by using an internal endpoint, make sure that the ECS and ApsaraDB RDS instances have the same network type. If both instances are of the VPC network type, make sure that both instances reside within the same VPC. For more information about how to view the internal endpoint, see View and modify the internal endpoint and port number . |

| Parameter | Description |
|-----------|---|
| username | The username of the privileged account of the ApsaraDB RDS for PostgreSQL instance. |
| port | The port number of the ApsaraDB RDS for PostgreSQL instance. |
| dbname | The name of the database whose data you want to restore. |
| table | The name of the table whose data you want to restore. |
| -c | <code>-c</code> : specifies to delete the database objects on which the table depends before data restoration. For more information, visit pg_restore . |
| dumpdir | The directory and name of the logical backup file to use. |

Example:

```
pg_restore -h 'pgm-bpxxxx.pg.rds.aliyuncs.com' -U test123 -p 3433 -d testdb2 -t products -c /tmp/testdb.dump
```

- When `Password:` appears, enter the password of the privileged account of your ApsaraDB RDS instance and press the Enter key.



Restore the schema of a database with data excluded

- Log on to the ECS instance or on-premises host that houses the logical backup file and run the following command to restore only the schema of a database:

```
pg_restore -h '<hostname>' -U <username> -p <port> -d <dbname> -s <dumpdir>
```

| Parameter | Description |
|-----------|--|
| hostname | The endpoint of the ApsaraDB RDS for PostgreSQL instance. <div style="border: 1px solid #add8e6; padding: 5px;"> <p> Note If your ECS instance connects to the ApsaraDB RDS for PostgreSQL instance by using an internal endpoint, make sure that the ECS and ApsaraDB RDS instances have the same network type. If both instances are of the VPC network type, make sure that both instances reside within the same VPC. For more information about how to view the internal endpoint, see View and modify the internal endpoint and port number.</p> </div> |
| username | The username of the privileged account of the ApsaraDB RDS for PostgreSQL instance. |
| port | The port number of the ApsaraDB RDS for PostgreSQL instance. |
| dbname | The name of the database whose schema you want to restore. |
| -s | <code>-s</code> : specifies to restore only the schema of the database. The data of the database is not restored. For more information, visit pg_restore . |
| dumpdir | The directory and name of the logical backup file to use. |

Example:

```
pg_restore -h 'pgm-bpxxxxx.pg.rds.aliyuncs.com' -U test123 -p 3433 -d testdb4 -s /tmp/testdb2.dump
```

- When **Password:** appears, enter the password of the privileged account of your ApsaraDB RDS instance and press the Enter key.

Note You can ignore alerts generated by the embedded plpgsql plug-in.

```
[root@iZbp... ~]# pg_restore -h 'pgm-bp...pg.rds.aliyuncs.com' -U ... -p 3433 -d testdb4 -s /tmp/testdb2.dump
Password:
pg_restore: [archiver (db)] Error while PROCESSING TOC:
pg_restore: [archiver (db)] Error from TOC entry 3075; 0 0 COMMENT EXTENSION plpgsql
pg_restore: [archiver (db)] could not execute query: ERROR: must be owner of extension plpgsql
Command was: COMMENT ON EXTENSION plpgsql IS 'PL/pgSQL procedural language';

WARNING: errors ignored on restore: 1
```

10.15. Plug-ins

10.15.1. Plug-ins supported

This topic describes the plug-ins that are supported by ApsaraDB RDS for PostgreSQL and their available versions.

PostgreSQL 12

| Plug-in | Version |
|--------------------|---------|
| btree_gin | 1.3 |
| btree_gist | 1.5 |
| citext | 1.6 |
| cube | 1.4 |
| dblink | 1.2 |
| dict_int | 1 |
| earthdistance | 1.1 |
| fuzzystrmatch | 1.1 |
| hstore | 1.6 |
| intagg | 1.1 |
| intarray | 1.2 |
| isn | 1.2 |
| ltree | 1.1 |
| pg_buffercache | 1.3 |
| pg_prewarm | 1.2 |
| pg_stat_statements | 1.7 |
| pg_trgm | 1.4 |
| pgcrypto | 1.3 |

| Plug-in | Version |
|------------------------------------|---------|
| pgrowlocks | 1.2 |
| pgstattuple | 1.5 |
| postgres_fdw | 1 |
| sslinfo | 1.2 |
| tablefunc | 1 |
| unaccent | 1.1 |
| plpgsql | 1 |
| plperl | 1 |
| pg_roaringbitmap | 0.5.0 |
| rdkit | 3.8 |
| mysql_fdw | 1.1 |
| ganos_geometry_sfcgal | 3.0 |
| ganos_geometry_topology | 3.0 |
| ganos_geometry | 3.0 |
| ganos_networking | 3.0 |
| ganos_pointcloud_geometry | 3.0 |
| ganos_pointcloud | 3.0 |
| ganos_raster | 3.0 |
| ganos_spatialref | 3.0 |
| ganos_trajectory | 3.0 |
| ganos_tiger_geocoder | 3.0 |
| ganos_address_standardizer | 3.0 |
| ganos_address_standardizer_data_us | 3.0 |
| wal2json | 2.0 |
| hll | 2.14 |
| plproxy | 2.9.0 |
| tsm_system_rows | 1.0 |
| tsm_system_time | 1.0 |

| Plug-in | Version |
|-------------|---------|
| smlar | 1.0 |
| tds_fdw | 1.0 |
| bigm | 1.2 |
| timescaledb | 1.7.1 |

PostgreSQL 11

| Plug-in | Version |
|--------------------|---------|
| plpgsql | 1 |
| pg_stat_statements | 1.6 |
| btree_gin | 1.3 |
| btree_gist | 1.5 |
| citext | 1.5 |
| cube | 1.4 |
| rum | 1.3 |
| dblink | 1.2 |
| dict_int | 1 |
| earthdistance | 1.1 |
| hstore | 1.5 |
| intagg | 1.1 |
| intarray | 1.2 |
| isn | 1.2 |
| ltree | 1.1 |
| pgcrypto | 1.3 |
| pgrowlocks | 1.2 |
| pg_prewarm | 1.2 |
| pg_trgm | 1.4 |
| postgres_fdw | 1 |
| sslinfo | 1.2 |
| tablefunc | 1 |
| timescaledb | 1.7.1 |

| Plug-in | Version |
|------------------------------|---------|
| unaccent | 1.1 |
| fuzzystrmatch | 1.1 |
| pgstattuple | 1.5 |
| pg_buffercache | 1.3 |
| zhparser | 1 |
| pg_pathman | 1.5 |
| plperl | 1 |
| orafce | 3.8 |
| pg_concurrency_control | 1 |
| varbitx | 1 |
| postgis | 2.5.1 |
| pgrouting | 2.6.2 |
| postgis_sfcgal | 2.5.1 |
| postgis_topology | 2.5.1 |
| address_standardizer | 2.5.1 |
| address_standardizer_data_us | 2.5.1 |
| ogr_fdw | 1 |
| ganos_pointcloud | 3.0 |
| ganos_spatialref | 3.0 |
| log_fdw | 1.0 |
| wal2json | 2.2 |
| PL/v8 | 2.3.13 |
| pg_cron | 1.1 |
| pase | 0.0.1 |
| hll | 2.14 |
| oss_fdw | 1.1 |
| tds_fdw | 2.0.1 |
| plproxy | 2.9.0 |
| tsm_system_rows | 1.0 |

| Plug-in | Version |
|-----------------|---------|
| tsm_system_time | 1.0 |
| smlar | 1.0 |
| zombodb | 4.0 |
| bigm | 1.2 |

PostgreSQL 10

| Plug-in | Version |
|--------------------|---------|
| pg_stat_statements | 1.6 |
| btree_gin | 1.2 |
| btree_gist | 1.5 |
| chkpass | 1 |
| citext | 1.4 |
| cube | 1.2 |
| dblink | 1.2 |
| dict_int | 1 |
| earthdistance | 1.1 |
| hstore | 1.4 |
| intagg | 1.1 |
| intarray | 1.2 |
| isn | 1.1 |
| ltree | 1.1 |
| pgcrypto | 1.3 |
| pgrowlocks | 1.2 |
| pg_prewarm | 1.1 |
| pg_trgm | 1.3 |
| postgres_fdw | 1 |
| sslinfo | 1.2 |
| tablefunc | 1 |
| unaccent | 1.1 |
| postgis_sfcgal | 2.5.1 |

| Plug-in | Version |
|------------------------------|---------|
| postgis_topology | 2.5.1 |
| fuzzystrmatch | 1.1 |
| postgis_tiger_geocoder | 2.5.1 |
| address_standardizer | 2.5.1 |
| address_standardizer_data_us | 2.5.1 |
| ogr_fdw | 1 |
| plperl | 1 |
| plv8 | 1.4.2 |
| plls | 1.4.2 |
| plcoffee | 1.4.2 |
| uuid-oss | 1.1 |
| zhparser | 1 |
| pgrouting | 2.6.2 |
| pg_hint_plan | 1.3.0 |
| pgstattuple | 1.5 |
| oss_fdw | 1.1 |
| ali_decoding | 0.0.1 |
| varbitx | 1 |
| pg_buffercache | 1.3 |
| q3c | 1.5.0 |
| pg_sphere | 1 |
| smlar | 1 |
| rum | 1.3 |
| pg_pathman | 1.5 |
| aggs_for_arrays | 1.3.1 |
| mysql_fdw | 1 |
| orafce | 3.6 |
| plproxy | 2.8.0 |
| pg_concurrency_control | 1 |

| Plug-in | Version |
|------------------------------------|---------|
| postgis | 2.5.1 |
| ganos_geometry_sfcgal | 2.2 |
| ganos_geometry_topology | 2.2 |
| ganos_geometry | 2.2 |
| ganos_networking | 2.2 |
| ganos_pointcloud_geometry | 2.2 |
| ganos_pointcloud | 2.2 |
| ganos_raster | 2.2 |
| ganos_spatialref | 2.2 |
| ganos_trajectory | 2.2 |
| ganos_tiger_geocoder | 2.2 |
| ganos_address_standardizer | 2.2 |
| ganos_address_standardizer_data_us | 2.2 |

PostgreSQL 9.4

| Plug-in | Version |
|--------------------|---------|
| plpgsql | 1 |
| pg_stat_statements | 1.2 |
| btree_gin | 1 |
| btree_gist | 1 |
| chkpass | 1 |
| citext | 1 |
| cube | 1 |
| dblink | 1.1 |
| dict_int | 1 |
| earthdistance | 1 |
| hstore | 1.3 |
| intagg | 1 |
| intarray | 1 |
| isn | 1 |

| Plug-in | Version |
|------------------------|---------|
| ltree | 1 |
| pgcrypto | 1.1 |
| pgrowlocks | 1.1 |
| pg_prewarm | 1 |
| pg_trgm | 1.1 |
| postgres_fdw | 1 |
| sslinfo | 1 |
| tablefunc | 1 |
| tsearch2 | 1 |
| unaccent | 1 |
| postgis | 2.2.8 |
| postgis_topology | 2.2.8 |
| fuzzystrmatch | 1 |
| postgis_tiger_geocoder | 2.2.8 |
| plperl | 1 |
| pltcl | 1 |
| plv8 | 1.4.2 |
| plls | 1.4.2 |
| plcoffee | 1.4.2 |
| uuid-oss | 1 |
| zhparser | 1 |
| pgrouting | 2.0.0 |
| rdkit | 3.4 |
| pg_hint_plan | 1.1.3 |
| pgstattuple | 1.2 |
| oss_fdw | 1.1 |
| jsonbx | 1 |
| ali_decoding | 0.0.1 |
| varbitx | 1 |

| Plug-in | Version |
|------------------------|---------|
| pg_buffercache | 1 |
| smlar | 1 |
| pg_sphere | 1 |
| q3c | 1.5.0 |
| pg_awr | 1 |
| imgsmr | 1 |
| orafce | 3.6 |
| pg_concurrency_control | 1 |

10.15.2. Use mysql_fdw to read data from and write data to a MySQL database

This topic describes how to use the mysql_fdw plug-in of ApsaraDB RDS for PostgreSQL to read data from and write data to a database on an ApsaraDB RDS for MySQL instance or a self-managed MySQL database.

Prerequisites

- The instance runs PostgreSQL 10.
- Communication between your ApsaraDB RDS for PostgreSQL instance and the MySQL database is normal.

Context

PostgreSQL 9.6 and later support parallel computing. PostgreSQL 11 can use joins on up to a billion data records to complete queries in seconds. A number of users prefer to use PostgreSQL to build small-sized data warehouses and process highly concurrent access requests. PostgreSQL 13 is under development. It will support columnar storage engines that further improve analysis capabilities.

The mysql_fdw plug-in establishes a connection to synchronize data from a MySQL database to your ApsaraDB RDS for PostgreSQL instance.

Procedure

1. Log on to a database of your ApsaraDB RDS for PostgreSQL instance. For more information, see [Connect to an ApsaraDB RDS for PostgreSQL instance](#).
2. Create the mysql_fdw plug-in.

```
create extension mysql_fdw;
```

3. Define a MySQL server.

```
CREATE SERVER <Name of the MySQL server>
FOREIGN DATA WRAPPER mysql_fdw
OPTIONS (host '<Endpoint used to connect to the MySQL server>', port '<Port used to connect to the MySQL server>');
```

Example:

```
CREATE SERVER mysql_server  
FOREIGN DATA WRAPPER mysql_fdw  
OPTIONS (host 'rm-xxx.mysql.rds.aliyuncs.com', port '3306');
```

4. Map the MySQL server to an account created on your ApsaraDB RDS for PostgreSQL instance. Then, the account can be used to access data in the MySQL database on the MySQL server.

```
CREATE USER MAPPING FOR <Username of the account to which the MySQL server is mapped>  
SERVER <Name of the MySQL server>  
OPTIONS (username '<Username used to log on to the MySQL database>', password '<Password used to log on to the MySQL database>');
```

Example:

```
CREATE USER MAPPING FOR pgtest  
SERVER mysql_server  
OPTIONS (username 'mysqltest', password 'Test1234!');
```

5. Create a foreign MySQL table by using the account that you mapped to the MySQL server in the previous step.

 **Note** The field names in the foreign MySQL table must be the same as those in the table of the MySQL database. You can choose to create only the fields you want to query. For example, if the table in the MySQL database contains the ID, NAME, and AGE fields, you can create only the ID and NAME fields in the foreign MySQL table.

```
CREATE FOREIGN TABLE <Name of the foreign MySQL table> (<Name of Field 1> <Data type of Field 1>, <Name of Field 2> <Data type of Field 2>...) server <Name of the MySQL server> options (dbname '<Name of the MySQL database>', table_name '<Name of the table in the MySQL database>');
```

Example:

```
CREATE FOREIGN TABLE ft_test (id1 int, name1 text) server mysql_server options (dbname 'test123', table_name 'test');
```

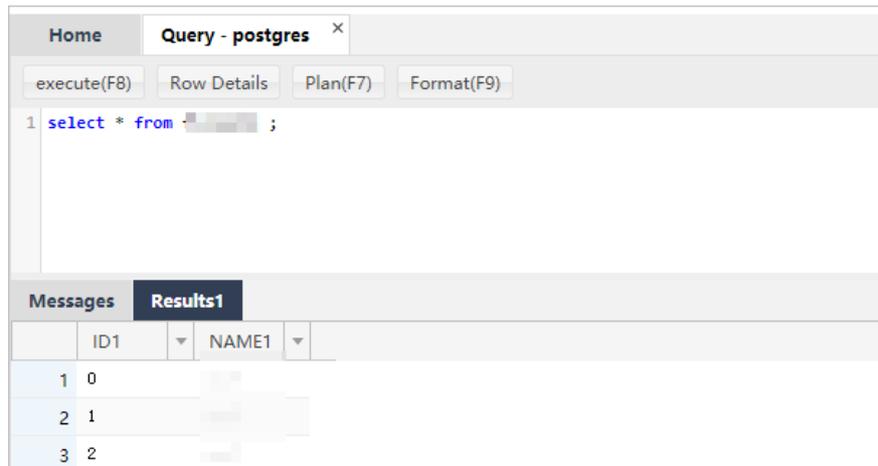
What to do next

You can use the foreign MySQL table to test the performance of read and write operations on the MySQL database.

 **Note** Data can be written to the table in the MySQL database only when the table is assigned a primary key. If the table is not assigned a primary key, the following error is reported:

```
ERROR: first column of remote table must be unique for INSERT/UPDATE/DELETE operation.
```

```
select * from ft_test ;  
insert into ft_test values (2,'abc');  
insert into ft_test select generate_series(3,100),'abc';  
select count(*) from ft_test ;
```



Run `postgres=> explain verbose select count(*) from ft_test;` to find out how the requests sent from your ApsaraDB RDS for PostgreSQL instance are executed to query data from the MySQL database. Command output:

```

QUERY PLAN
-----
Aggregate (cost=1027.50..1027.51 rows=1 width=8)
  Output: count(*)
  -> Foreign Scan on public.ft_test (cost=25.00..1025.00 rows=1000 width=0)
    Output: id, info
    Remote server startup cost: 25
    Remote query: SELECT NULL FROM `test123`.`test`
(6 rows)

```

10.15.3. Use oss_fdw to read and write foreign data files

This topic describes how to use the `oss_fdw` plug-in to load data between Object Storage Service (OSS) and PostgreSQL or PPAS databases.

oss_fdw parameters

The `oss_fdw` plug-in uses a method similar to other Foreign Data Wrapper (FDW) interfaces to encapsulate foreign data stored in OSS. You can use `oss_fdw` to read data stored in OSS. This process is similar to reading data tables. `oss_fdw` provides unique parameters to connect and parse file data in OSS.

Note

- `oss_fdw` can read and write files of the following types in OSS: TXT and CSV files as well as GZIP-compressed TXT and CSV files.
- The value of each parameter must be enclosed in double quotation marks (") and cannot contain unnecessary spaces.

CREATE SERVER parameters

- `ossendpoint`: the endpoint used to access OSS over the internal network, also known as the host.
- `id oss`: the AccessKey ID of the OSS account.
- `key oss`: the AccessKey secret of the OSS account.
- `bucket`: the bucket where the data you want to access is stored. You must create an OSS account before you specify this parameter.

The following fault tolerance parameters can be used for data import and export. If network connectivity is poor, you can adjust these parameters to ensure successful import and export.

- `oss_connect_timeout`: the connection timeout period. Default value: 10. Unit: seconds.
- `oss_dns_cache_timeout`: the DNS timeout period. Default value: 60. Unit: seconds.
- `oss_speed_limit`: the minimum data transmission rate. Default value: 1024. Unit: bytes/s.
- `oss_speed_time`: the maximum waiting period during which the data transmission rate is lower than the minimum value. Default value: 15. Unit: seconds.

If the default values of `oss_speed_limit` and `oss_speed_time` are used, a timeout error occurs when the transmission rate is lower than 1,024 bytes/s for 15 consecutive seconds.

CREATE FOREIGN TABLE parameters

- `filepath`: a file name that contains a path in OSS.
 - The file name specified by this parameter contains the directory name but not the bucket name.
 - This parameter matches multiple files in the corresponding path in OSS. You can load multiple files to a database.
 - You can import files that adhere to the `filepath` or `filepath.x` format to a database. The values of `x` must be consecutive numbers starting from 1.

For example, among the files named `filepath`, `filepath.1`, `filepath.2`, `filepath.3`, and `filepath.5`, the first four files are matched and imported. The `filepath.5` file is not imported.
- `dir`: the virtual file directory in OSS.
 - The specified directory must end with a forward slash (/).
 - All files (excluding subfolders and files in subfolders) in the virtual file directory specified by `dir` are matched and imported to a database.
- `prefix`: the prefix of the path name corresponding to the data file. The prefix does not support regular expressions. The `prefix`, `filepath`, and `dir` parameters are mutually exclusive. Therefore, only one of them can be specified at a time.
- `format`: the file format, which can only be CSV.
- `encoding`: the file data encoding format. It supports common PostgreSQL encoding formats, such as UTF-8.
- `parse_errors`: the fault-tolerant parsing mode. If an error occurs during the parsing process, the entire row of data is ignored.
- `delimiter`: the string used to delimit columns.
- `quote`: the quote character for files.
- `escape`: the escape character for files.
- `null`: sets the column matching the specified string to null. For example, `null 'test'` is used to set the value of the 'test' column to null.
- `force_not_null`: sets the value of a column to a non-null value. For example, `force_not_null 'id'` is used to set the value of the 'id' column to empty strings.
- `compressiontype`: the format of the files to be read or written in OSS.
 - `none`: The files are uncompressed. This is the default value.
 - `gzip`: The files are compressed in the GZIP format.
- `compressionlevel`: the degree to which data files written to OSS are compressed. Valid values: 1 to 9. Default value: 6.

Note

- You must specify filepath and dir in the OPTIONS parameter.
- You must specify filepath or dir.
- The export mode can only be dir.

Export mode parameters for CREATE FOREIGN TABLE

- `oss_flush_block_size`: the buffer size for the data written to OSS at a time. Default value: 32. Valid values: 1 to 128. Unit: MB.
- `oss_file_max_size`: the maximum size of a data file allowed to be written to OSS. If a data file reaches the maximum size, the remaining data is written to another data file. Default value: 1024. Valid values: 8 to 4000. Unit: MB.
- `num_parallel_worker`: the maximum number of threads that are allowed to run in parallel to compress the data written to OSS. Valid values: 1 to 8. Default value: 3.

Auxiliary functions

FUNCTION `oss_fdw_list_file` (relname text, schema text DEFAULT 'public')

- This function obtains the name and size of the OSS file that a foreign table matches.
- The file size is measured in bytes.

The following result is returned after `select * from oss_fdw_list_file('t_oss');` is executed:

```

name      | size
-----+-----
oss_test/test.gz.1 | 739698350
oss_test/test.gz.2 | 739413041
oss_test/test.gz.3 | 739562048
(3 rows)

```

Auxiliary features

`oss_fdw.rds_read_one_file`: In read mode, this feature is used to specify a file to match the foreign table. The foreign table matches only the specified file during data import.

Example: `set oss_fdw.rds_read_one_file = 'oss_test/example16.csv.1';`

The following result is returned after `set oss_fdw.rds_read_one_file = 'oss_test/test.gz.2';` and `select * from oss_fdw_list_file('t_oss');` are executed:

```

name      | size
-----+-----
oss_test/test.gz.2 | 739413041
(1 rows)

```

oss_fdw example

```
# Create the plug-in for a PostgreSQL database.
create extension oss_fdw; -- For a PPAS database, execute select rds_manage_extension('create','oss_fdw');
# Create a server.
CREATE SERVER ossserver FOREIGN DATA WRAPPER oss_fdw OPTIONS
  (host 'oss-cn-hangzhou.aliyuncs.com', id 'xxx', key 'xxx', bucket 'mybucket');
# Create an OSS foreign table.
CREATE FOREIGN TABLE ossexample
  (date text, time text, open float,
  high float, low float, volume int)
  SERVER ossserver
  OPTIONS ( filepath 'osstest/example.csv', delimiter ',',
  format 'csv', encoding 'utf8', PARSE_ERRORS '100');
# Create a table named example to which to import data.
create table example
  (date text, time text, open float,
  high float, low float, volume int);
# Load data from ossexample to example.
insert into example select * from ossexample;
# Result
# oss_fdw estimates the file size in OSS and formulates a query plan.
explain insert into example select * from ossexample;
      QUERY PLAN
-----
Insert on example (cost=0.00..1.60 rows=6 width=92)
-> Foreign Scan on ossexample (cost=0.00..1.60 rows=6 width=92)
    Foreign OssFile: osstest/example.csv.0
    Foreign OssFile Size: 728
(4 rows)
# Write the data in the example table to OSS.
insert into ossexample select * from example;
explain insert into ossexample select * from example;
      QUERY PLAN
-----
Insert on ossexample (cost=0.00..16.60 rows=660 width=92)
-> Seq Scan on example (cost=0.00..16.60 rows=660 width=92)
(2 rows)
```

Additional considerations

- `oss_fdw` is a foreign table plug-in developed based on the PostgreSQL FOREIGN TABLE framework.
- The data import performance varies based on the PostgreSQL cluster resources (CPU, I/O, and memory) and OSS.
- To ensure data import performance, the ApsaraDB RDS for PostgreSQL instance must be in the same region as the OSS bucket.

ID and key encryption

If the `id` and `key` parameters for `CREATE SERVER` are not encrypted, the `select * from pg_foreign_server` statement execution result displays the information. Your AccessKey ID and AccessKey secret are exposed. You can use symmetric encryption to hide your AccessKey ID and AccessKey secret. Use different AccessKey pairs for different instances to further protect your information. However, to avoid incompatibility with earlier versions, do not add data types as you would in Greenplum.

Encrypted information:

```
postgres=# select * from pg_foreign_server ;
  srvname | srvowner | srvid | srvttype | srsversion | srvacl |
-----+-----+-----+-----+-----+-----+-----
  ossserver | 10 | 16390 | | | | {host=oss-cn-hangzhou-zmf.aliyuncs.com,id=MD5xxxxxxxx,key=MD5xxxxxxxx,bu
cket=067862}
```

The encrypted information is preceded by the MD5 hash value. The remainder of the total length divided by 8 is 3. Therefore, encryption is not performed again when the exported data is imported. You cannot create an AccessKey pair that is preceded by MD5.

10.16. Use Pgpool for read/write splitting in ApsaraDB RDS for PostgreSQL

This topic describes how to use the Pgpool tool of PostgreSQL installed on an ECS instance to implement read/write splitting for your primary and read-only ApsaraDB RDS for PostgreSQL instances.

Context

If you do not use Pgpool to ensure high availability, Pgpool is stateless. The decrease in performance can be ignored. Additionally, Pgpool supports horizontal scaling of your database system. You can use Pgpool and the high availability architecture of ApsaraDB RDS for PostgreSQL to implement read/write splitting.

Set up a test environment

If you have purchased a primary ApsaraDB RDS instance that runs PostgreSQL 10 and have attached read-only instances to the primary instance, you need only to [install Pgpool](#). For more information, see [Create an instance](#) and [Create a read-only ApsaraDB RDS for PostgreSQL instance](#). After you install Pgpool, go to [Configure Pgpool](#).

1. Run the `vi /etc/sysctl.conf` command to open the `sysctl.conf` file. Modify the following configurations:

```
# add by digoal.zhou
fs.aio-max-nr = 1048576
fs.file-max = 76724600
# Optional. Set the kernel.core_pattern parameter to /data01/corefiles/core_%e_%u_%t_%s.%p.
# The /data01/corefiles directory that is used to store core dumps is created with the 777 permission before testing.
If a symbolic link is used, change the directory to 777.
kernel.sem = 4096 2147483647 2147483646 512000
# Specify the semaphore. You can run the ipcs -l or -u command to obtain the semaphore count. Each group of 16 pr
ocesses requires a semaphore with a count of 17.
kernel.shmall = 107374182
# Specify the total size of shared memory segments. Recommended value: 80% of the memory capacity. Unit: pages.

kernel.shmmax = 274877906944
# Specify the maximum size of a single shared memory segment. Recommended value: 50% of the memory capacity.
Unit: bytes. In PostgreSQL versions later than 9.2, the use of shared memory significantly drops.
kernel.shmuni = 819200
# Specify the total number of shared memory segments that can be generated. At least two shared memory segmen
ts must be generated within each PostgreSQL cluster.
net.core.netdev_max_backlog = 10000
net.core.rmem_default = 262144
# The default setting of the socket receive buffer in bytes.
net.core.rmem_max = 4194304
# The maximum receive socket buffer size in bytes
net.core.wmem_default = 262144
# The default setting (in bytes) of the socket send buffer.
net.core.wmem_max = 4194304
```

```
# The maximum send socket buffer size in bytes.
net.core.somaxconn = 4096
net.ipv4.tcp_max_syn_backlog = 4096
net.ipv4.tcp_keepalive_intvl = 20
net.ipv4.tcp_keepalive_probes = 3
net.ipv4.tcp_keepalive_time = 60
net.ipv4.tcp_mem = 8388608 12582912 16777216
net.ipv4.tcp_fin_timeout = 5
net.ipv4.tcp_synack_retries = 2
net.ipv4.tcp_syncookies = 1
# Enable SYN cookies. If an SYN waiting queue overflows, you can enable SYN cookies to defend against a small number of SYN attacks.
net.ipv4.tcp_timestamps = 1
# Reduce the time after which a network socket enters the TIME-WAIT state.
net.ipv4.tcp_tw_recycle = 0
# If you set this parameter to 1 to enable the recycle function, network sockets in the TIME-WAIT state over TCP connections are recycled. However, if network address translation (NAT) is used, TCP connections may fail. We recommend that you set this parameter to 0 on the database server.
net.ipv4.tcp_tw_reuse = 1
# Enable the reuse function. This function enables network sockets in the TIME-WAIT state to be reused over new TCP connections.
net.ipv4.tcp_max_tw_buckets = 262144
net.ipv4.tcp_rmem = 8192 87380 16777216
net.ipv4.tcp_wmem = 8192 65536 16777216
net.nf_conntrack_max = 1200000
net.netfilter.nf_conntrack_max = 1200000
vm.dirty_background_bytes = 409600000
# If the size of dirty pages reaches the specified limit, a background scheduling process (for example, pdflush) is invoked to flush the dirty pages to disks. These are the pages that are generated n seconds earlier. The value of n is calculated by using the following formula: n = Value of the dirty_expire_centisecs parameter/100.
# The default limit is 10% of the memory capacity. If the memory capacity is large, we recommend that you specify the limit in bytes.
vm.dirty_expire_centisecs = 3000
# Specify the maximum period to retain dirty pages. Dirty pages are flushed to disks after the time period specified by this parameter elapses. The value 3000 indicates 30 seconds.
vm.dirty_ratio = 95
# The processes that users call to write data onto disks must actively flush dirty pages to disks. This applies when the background scheduling process to flush dirty pages is slow and the size of dirty pages exceeds 95% of the memory capacity. These processes include fsync and fdatasync.
# Set this parameter properly to prevent user-called processes from flushing dirty pages to disks. This allows you to create multiple ApsaraDB RDS instances on a single server and use control groups to limit the input/output operations per second (IOPS) per instance.
vm.dirty_writeback_centisecs = 100
# Specify the time interval at which the background scheduling process (such as pdflush) flushes dirty pages to disks. The value 100 indicates 1 second.
vm.swappiness = 0
# Disable the swap function.
vm.mmap_min_addr = 65536
vm.overcommit_memory = 0
# Specify whether you can allocate more memory space than the physical host has available. If you set this parameter to 1, the system always considers the available memory space sufficient. If the memory capacity provided in the test environment is low, we recommend that you set this parameter to 1.
vm.overcommit_ratio = 90
# Specify the memory capacity that can be allocated when the overcommit_memory parameter is set to 2.
vm.swappiness = 0
# Disable the swap function.
vm.zone_reclaim_mode = 0
# Disable non-uniform memory access (NUMA). You can also disable NUMA in the vmlinux file.
net.ipv4.ip_local_port_range = 40000 65535
# Specify the range of TCP or UDP port numbers for the physical host to allocate.
```

```
fs.nr_open=2048000
# Specify the maximum number of file handles that a single process can open.
# Take note of the following parameters:
#vm.extra_free_kbytes = 4096000 # If the physical host provides a low memory capacity, do not specify a large value such as 4096000. If you specify a large value, the physical host may not start.
#vm.min_free_kbytes = 6291456 # We recommend that you increase the value of the vm.min_free_kbytes parameter by 1 GB for every 32 GB of memory.
# If the physical host does not provide much memory, we recommend that you do not configure vm.extra_free_kbytes and vm.min_free_kbytes.
# vm.nr_hugepages = 66536
# If the size of the shared buffer exceeds 64 GB, we recommend that you use huge pages. You can specify the page size by setting the Hugepagesize parameter in the /proc/meminfo file.
#vm.lowmem_reserve_ratio = 1 1 1
# If the memory capacity exceeds 64 GB, we recommend that you set this parameter. Otherwise, we recommend that you retain the default value 256 256 32.
```

2. Run the `vi /etc/security/limits.conf` command to open the `limits.conf` file. Modify the following configurations:

```
* soft nfile 1024000
* hard nfile 1024000
* soft nproc unlimited
* hard nproc unlimited
* soft core unlimited
* hard core unlimited
* soft memlock unlimited
* hard memlock unlimited
# Comment out the other parameters in the limits.conf file.
# Comment out the /etc/security/limits.d/20-nproc.conf file.
```

3. Run the following commands to open the `rc.local` file:

```
chmod +x /etc/rc.local
vi /etc/rc.local
```

Modify the following configurations to disable transparent huge pages, configure huge pages, and start PostgreSQL:

```
# Disable transparent huge pages.
if test -f /sys/kernel/mm/transparent_hugepage/enabled; then
  echo never > /sys/kernel/mm/transparent_hugepage/enabled
fi
# Configure huge pages for two instances. Each instance has a shared buffer of 16 GB.
sysctl -w vm.nr_hugepages=17000
# Start the two instances.
su - postgres -c "pg_ctl start -D /data01/pg12_3389/pg_root"
su - postgres -c "pg_ctl start -D /data01/pg12_8002/pg_root"
```

4. Create a file system.

 **Warning** If you use a new disk, you must verify that the new disk belongs to the `vdb` partition instead of the `vda` partition. If the new disk belongs to the `vda` partition, data may be deleted from the new disk.

```
parted -a optimal -s /dev/vdb mklabel gpt mkpart primary 1MiB 100%FREE
mkfs.ext4 /dev/vdb1 -m 0 -O extent,uninit_bg -E lazy_init=1 -b 4096 -T largefile -L vdb1
vi /etc/fstab
LABEL=vdb1 /data01 ext4 defaults,noatime,nodiratime,nodelalloc,barrier=0,data=writeback 0 0
mkdir /data01
mount -a
```

5. Start the irqbalance command line tool.

```
systemctl status irqbalance
systemctl enable irqbalance
systemctl start irqbalance
systemctl status irqbalance
```

6. Install PostgreSQL 10 and Pgpool.

```
yum install -y https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm
yum install -y https://download.postgresql.org/pub/repos/yum/repopms/EL-7-x86_64/pgdg-redhat-repo-latest.noarch.rpm
yum search all postgresql
yum search all pgpool
yum install -y postgresql12*
yum install -y pgpool-II-12-extensions
```

7. Initialize the data directory of your database system.

```
mkdir /data01/pg12_3389
chown postgres:postgres /data01/pg12_3389
```

8. Configure environment variables for the postgres user.

```
su - postgres
vi .bash_profile
```

Append the following parameters to the environment variables:

```
export PS1="$USER@`/bin/hostname -s`-> "
export PGPORT=3389
export PGDATA=/data01/pg12_3389/pg_root
export LANG=en_US.utf8
export PGHOME=/usr/pgsql-12
export LD_LIBRARY_PATH=$PGHOME/lib:/lib64:/usr/lib64:/usr/local/lib64:/lib:/usr/lib:/usr/local/lib:$LD_LIBRARY_PATH
export DATE=`date +"%Y%m%d%H%M"`
export PATH=$PGHOME/bin:$PATH:
export MANPATH=$PGHOME/share/man:$MANPATH
export PGHOST=$PGDATA
export PGUSER=postgres
export PGDATABASE=db1
alias rm='rm -i'
alias ll='ls -lh'
unalias vi
```

9. Initialize your primary ApsaraDB RDS instance.

```
initdb -D $PGDATA -U postgres -E UTF8 --lc-collate=C --lc-ctype=en_US.utf8
```

10. Modify the postgresql.conf file.

```
listen_addresses = '0.0.0.0'
port = 3389
max_connections = 1500
superuser_reserved_connections = 13
unix_socket_directories = '., /var/run/postgresql, /tmp'
tcp_keepalives_idle = 60
tcp_keepalives_interval = 10
tcp_keepalives_count = 10
shared_buffers = 16GB
huge_pages = on
work_mem = 8MB
```

```
work_mem = 1MB
maintenance_work_mem = 1GB
dynamic_shared_memory_type = posix
vacuum_cost_delay = 0
bgwriter_delay = 10ms
bgwriter_lru_maxpages = 1000
bgwriter_lru_multiplier = 10.0
bgwriter_flush_after = 512kB
effective_io_concurrency = 0
max_worker_processes = 128
max_parallel_maintenance_workers = 3
max_parallel_workers_per_gather = 4
parallel_leader_participation = off
max_parallel_workers = 8
backend_flush_after = 256
wal_level = replica
synchronous_commit = off
full_page_writes = on
wal_compression = on
wal_buffers = 16MB
wal_writer_delay = 10ms
wal_writer_flush_after = 1MB
checkpoint_timeout = 15min
max_wal_size = 64GB
min_wal_size = 8GB
checkpoint_completion_target = 0.2
checkpoint_flush_after = 256kB
random_page_cost = 1.1
effective_cache_size = 48GB
log_destination = 'csvlog'
logging_collector = on
log_directory = 'log'
log_filename = 'postgresql-%a.log'
log_truncate_on_rotation = on
log_rotation_age = 1d
log_rotation_size = 0
log_min_duration_statement = 1s
log_checkpoints = on
log_connections = on
log_disconnections = on
log_line_prefix = '%m [%p] '
log_statement = 'ddl'
log_timezone = 'Asia/Shanghai'
autovacuum = on
log_autovacuum_min_duration = 0
autovacuum_vacuum_scale_factor = 0.1
autovacuum_analyze_scale_factor = 0.05
autovacuum_freeze_max_age = 800000000
autovacuum_multixact_freeze_max_age = 900000000
autovacuum_vacuum_cost_delay = 0
vacuum_freeze_table_age = 750000000
vacuum_multixact_freeze_table_age = 750000000
datestyle = 'iso, mdy'
timezone = 'Asia/Shanghai'
lc_messages = 'en_US.utf8'
lc_monetary = 'en_US.utf8'
lc_numeric = 'en_US.utf8'
lc_time = 'en_US.utf8'
default_text_search_config = 'pg_catalog.english'
```

11. Modify the `pg_hba.conf` file.

Note Pgpool-II is installed on the same ECS instance as the database server where PostgreSQL resides. If you specify the 127.0.0.1 IP address in the `pg_hba.conf` file, you must enter the correct password to ensure a successful logon.

```
# "local" is for Unix domain socket connections only
local all all trust
# IPv4 local connections:
host all all 127.0.0.1/32 md5
# IPv6 local connections:
host all all ::1/128 trust
# Allow replication connections from localhost, by a user with the
# replication privilege.
local replication all trust
host replication all 127.0.0.1/32 trust
host replication all ::1/128 trust
host db123 digoal 0.0.0.0/0 md5
```

12. Execute a statement in the database to create a user authorized with streaming replication permissions.
Example:

```
create role rep123 login replication encrypted password 'xxxxxxx';
```

13. Execute statements in the database to create a user and authorize it to manage your ApsaraDB RDS instances.
Example:

```
create role digoal login encrypted password 'xxxxxxx';
create database db123 owner digoal;
```

14. Create a user who is authorized to check the health heartbeats between Pgpool and your read-only ApsaraDB RDS instances. With the parameters of Pgpool properly configured, this user can check the write-ahead logging (WAL) replay latency on each read-only ApsaraDB RDS instance. Example:

```
create role nobody login encrypted password 'xxxxxxx';
```

Create a secondary ApsaraDB RDS instance

To simplify the test procedure, perform the following steps to create a secondary ApsaraDB RDS instance on the same ECS instance as your primary ApsaraDB RDS instance:

1. Use the `pg_basebackup` tool to create a secondary ApsaraDB RDS instance.

```
pg_basebackup -D /data01/pg12_8002/pg_root -F p --checkpoint=fast -P -h 127.0.0.1 -p 3389 -U rep123
```

2. Run the following commands to open the `postgresql.conf` file of the secondary ApsaraDB RDS instance:

```
cd /data01/pg12_8002/pg_root
vi postgresql.conf
```

Modify the following configurations:

```
# The secondary ApsaraDB RDS instance has the following configurations different from the primary ApsaraDB RDS instance:
port = 8002
primary_conninfo = 'hostaddr=127.0.0.1 port=3389 user=rep123' # You do not need to set the password. This is because trust relationships are configured on the primary ApsaraDB RDS instance.
hot_standby = on
wal_receiver_status_interval = 1s
wal_receiver_timeout = 10s
recovery_target_timeline = 'latest'
```

3. Configure the standby.signal file of the secondary ApsaraDB RDS instance.

```
cd /data01/pg12_8002/pg_root
touch standby.signal
```

4. Execute the `SELECT * FROM pg_stat_replication;` statement in the database to check whether data is properly synchronized between the primary and secondary ApsaraDB RDS instances. The following output is returned:

```
-[ RECORD 1 ]-----+-----
pid          | 21065
usesysid     | 10
username     | postgres
application_name | walreceiver
client_addr  | 127.0.0.1
client_hostname |
client_port  | 47064
backend_start | 2020-02-29 00:26:28.485427+08
backend_xmin  |
state        | streaming
sent_lsn     | 0/52000060
write_lsn    | 0/52000060
flush_lsn    | 0/52000060
replay_lsn   | 0/52000060
write_lag    |
flush_lag    |
replay_lag   |
sync_priority | 0
sync_state   | async
reply_time   | 2020-02-29 01:32:40.635183+08
```

Configure Pgpool

1. Query the location where Pgpool is installed.

```
rpm -qa | grep pgpool
pgpool-II-12-extensions-4.1.1-1.rhel7.x86_64
pgpool-II-12-4.1.1-1.rhel7.x86_64
rpm -ql pgpool-II-12-4.1.1
```

2. Run the following commands to open the pgpool.conf file:

```
cd /etc/pgpool-II-12/
cp pgpool.conf.sample-stream pgpool.conf
vi pgpool.conf
```

Modify the following configurations:

```
listen_addresses = '0.0.0.0'
port = 8001
socket_dir = '/tmp'
reserved_connections = 0
pcp_listen_addresses = ''
pcp_port = 9898
pcp_socket_dir = '/tmp'
# - Backend Connection Settings -
backend_hostname0 = '127.0.0.1'
# Host name or IP address to connect to for backend 0
backend_port0 = 3389
# Port number for backend 0
backend_weight0 = 1
# Weight for backend 0 (only in load balancing mode)
```

```
backend_data_directory0 = '/data01/pg12_3389/pg_root'
    # Data directory for backend 0
backend_flag0 = 'ALWAYS_MASTER'
    # Controls various backend behavior
    # ALLOW_TO_FAILOVER, DISALLOW_TO_FAILOVER
    # or ALWAYS_MASTER
backend_application_name0 = 'server0'
    # walsender's application_name, used for "show pool_nodes" command
backend_hostname1 = '127.0.0.1'
backend_port1 = 8002
backend_weight1 = 1
backend_data_directory1 = '/data01/pg12_8002/pg_root'
backend_flag1 = 'DISALLOW_TO_FAILOVER'
backend_application_name1 = 'server1'
# - Authentication -
enable_pool_hba = on
    # Use pool_hba.conf for client authentication
pool_passwd = 'pool_passwd'
    # File name of pool_passwd for md5 authentication.
    # "" disables pool_passwd.
    # (change requires restart)
allow_clear_text_frontend_auth = off
    # Allow Pgpool-II to use clear text password authentication
    # with clients, when pool_passwd does not
    # contain the user password
# - Concurrent session and pool size -
num_init_children = 128
    # Number of concurrent sessions allowed
    # (change requires restart)
max_pool = 4
    # Number of connection pool caches per connection
    # (change requires restart)
# - Life time -
child_life_time = 300
    # Pool exits after being idle for this many seconds
child_max_connections = 0
    # Pool exits after receiving that many connections
    # 0 means no exit
connection_life_time = 0
    # Connection to backend closes after being idle for this many seconds
    # 0 means no close
client_idle_limit = 0
    # Client is disconnected after being idle for that many seconds
    # (even inside an explicit transactions!)
    # 0 means no disconnection
#-----
# LOGS
#-----
# - Where to log -
log_destination = 'syslog'
    # Where to log
    # Valid values are combinations of stderr,
    # and syslog. Default to stderr.
log_connections = on
    # Log connections
log_standby_delay = 'if_over_threshold'
    # Log standby delay
    # Valid values are combinations of always,
    # if_over_threshold, none
#-----
# CONNECTIONS
```

```

# FILE LOCATIONS
#-----
pid_file_name = '/var/run/pgpool-II-12/pgpool.pid'
    # PID file name
    # Can be specified as relative to the"
    # location of pgpool.conf file or
    # as an absolute path
    # (change requires restart)
logdir = '/tmp'
    # Directory of pgPool status file
    # (change requires restart)
#-----
# CONNECTION POOLING
#-----
connection_cache = on
    # Activate connection pools
    # (change requires restart)
    # Semicolon separated list of queries
    # to be issued at the end of a session
    # The default is for 8.3 and later
reset_query_list = 'ABORT; DISCARD ALL'
#-----
# LOAD BALANCING MODE
#-----
load_balance_mode = on
    # Activate load balancing mode
    # (change requires restart)
ignore_leading_white_space = on
    # Ignore leading white spaces of each query
white_function_list = ''
    # Comma separated list of function names
    # that don't write to database
    # Regexp are accepted
black_function_list = 'currval,lastval,nextval,setval'
    # Comma separated list of function names
    # that write to database
    # Regexp are accepted
black_query_pattern_list = ''
    # Semicolon separated list of query patterns
    # that should be sent to primary node
    # Regexp are accepted
    # valid for streaming replicaton mode only.
database_redirect_preference_list = ''
    # comma separated list of pairs of database and node id.
    # example: postgres:primary,mydb[0-4]:1,mydb[5-9]:2'
    # valid for streaming replicaton mode only.
app_name_redirect_preference_list = ''
    # comma separated list of pairs of app name and node id.
    # example: 'psql:primary,myapp[0-4]:1,myapp[5-9]:standby'
    # valid for streaming replicaton mode only.
allow_sql_comments = off
    # if on, ignore SQL comments when judging if load balance or
    # query cache is possible.
    # If off, SQL comments effectively prevent the judgment
    # (pre 3.4 behavior).
disable_load_balance_on_write = 'transaction'
    # Load balance behavior when write query is issued
    # in an explicit transaction.
    # Note that any query not in an explicit transaction
    # is not affected by the parameter.

```

```
# 'transaction' (the default): if a write query is issued,
# subsequent read queries will not be load balanced
# until the transaction ends.
# 'trans_transaction': if a write query is issued,
# subsequent read queries in an explicit transaction
# will not be load balanced until the session ends.
# 'always': if a write query is issued, read queries will
# not be load balanced until the session ends.
statement_level_load_balance = off
# Enables statement level load balancing
#-----
# MASTER/SLAVE MODE
#-----
master_slave_mode = on
# Activate master/slave mode
# (change requires restart)
master_slave_sub_mode = 'stream'
# Master/slave sub mode
# Valid values are combinations stream, slony
# or logical. Default is stream.
# (change requires restart)
# - Streaming -
sr_check_period = 3
# Streaming replication check period
# Disabled (0) by default
sr_check_user = 'nobody'
# Streaming replication check user
# This is necessary even if you disable streaming
# replication delay check by sr_check_period = 0
sr_check_password = ''
# Password for streaming replication check user
# Leaving it empty will make Pgpool-II to first look for the
# Password in pool_passwd file before using the empty password
sr_check_database = 'postgres'
# Database name for streaming replication check
delay_threshold = 512000
# Threshold before not dispatching query to standby node
# Unit is in bytes
# Disabled (0) by default
#-----
# HEALTH CHECK GLOBAL PARAMETERS
#-----
health_check_period = 5
# Health check period
# Disabled (0) by default
health_check_timeout = 10
# Health check timeout
# 0 means no timeout
health_check_user = 'nobody'
# Health check user
health_check_password = ''
# Password for health check user
# Leaving it empty will make Pgpool-II to first look for the
# Password in pool_passwd file before using the empty password
health_check_database = ''
# Database name for health check. If '', tries 'postgres' first,
health_check_max_retries = 60
# Maximum number of times to retry a failed health check before giving up.
health_check_retry_delay = 1
# Amount of time to wait (in seconds) between retries.
```

```

connect_timeout = 10000
    # Timeout value in milliseconds before giving up to connect to backend.
    # Default is 10000 ms (10 second). Flaky network user may want to increase
    # the value. 0 means no timeout.
    # Note that this value is not only used for health check,
    # but also for ordinary connection to backend.
#-----
# FAILOVER AND FAILBACK
#-----
failover_on_backend_error = off
    # Initiates failover when reading/writing to the
    # backend communication socket fails
    # If set to off, pgpool will report an
    # error and disconnect the session.
relcache_expire = 0 # After the configuration file is restructured, we recommend that you set this parameter to 1, re
load the configuration file, and then set this parameter to 0 again. You can also set this parameter to a specific point
in time.
    # Life time of relation cache in seconds.
    # 0 means no cache expiration(the default).
    # The relation cache is used for cache the
    # query result against PostgreSQL system
    # catalog to obtain various information
    # including table structures or if it's a
    # temporary table or not. The cache is
    # maintained in a pgpool child local memory
    # and being kept as long as it survives.
    # If someone modify the table by using
    # ALTER TABLE or some such, the relcache is
    # not consistent anymore.
    # For this purpose, cache_expiration
    # controls the life time of the cache.
relcache_size = 8192
    # Number of relation cache
    # entry. If you see frequently:
    # "pool_search_relcache: cache replacement happend"
    # in the pgpool log, you might want to increate this number.

```

3. Run the `cd /etc/pgpool-II-12` command to configure the `pool_passwd` file.

 **Note** If you connect to your ApsaraDB RDS instances by using Pgpool, you must configure the `pool_passwd` file. This is because Pgpool supports the authentication protocol of PostgreSQL.

```

# Run the following command:
#pg_md5 --md5auth --username=username password
# Generate the passwords of the digoal and nobody users. The passwords are automatically written into the pool_p
asswd file.
pg_md5 --md5auth --username=digoal "xxxxxxx"
pg_md5 --md5auth --username=nobody "xxxxxxx"

```

4. Use the system to automatically generate the `pool_passwd` file.

```

cd /etc/pgpool-II-12
cat pool_passwd

```

5. Run the following commands to configure the `pgpool_hba` file:

```

cd /etc/pgpool-II-12
cp pool_hba.conf.sample pool_hba.conf
vi pool_hba.conf

```

Configure the following parameters:

```
host all all 0.0.0.0/0 md5
```

6. Configure the `pcp.conf` file.

Note The `pcp.conf` file is used to manage the users and passwords of Pgpool. It is not related to the users and passwords of your ApsaraDB RDS instances.

```
cd /etc/pgpool-II-12
# pg_md5 abc # In this command, you set the password to abc and encrypt it by using the MD5 encryption algorithm.
900150983cd24fb0d6963f7d28e17f72
cp pcp.conf.sample pcp.conf
vi pcp.conf
# USERID:MD5PASSWD
manage:900150983cd24fb0d6963f7d28e17f72 # In this command, the manage user is used to manage PCP.
```

7. Start Pgpool.

```
cd /etc/pgpool-II-12
pgpool -f ./pgpool.conf -a ./pool_hba.conf -F ./pcp.conf
```

Note If you want to view the logs of Pgpool, run the following command:

```
less /var/log/messages
```

8. Use Pgpool to connect to your ApsaraDB RDS instances.

```
psql -h 127.0.0.1 -p 8001 -U digoal postgres
```



```
[root@izbp... pgpool-II-12]# psql -h pgm-bp... .pg.rds.aliyuncs.com -p 3433 -U digoal postgres
Password for user digoal:
psql (12.2, server 10.10)
Type "help" for help.

postgres=>
```

FAQ

• Q: How do I test whether read/write splitting is enabled?

A: You can connect to your ApsaraDB RDS instances by using Pgpool and call the `pg_is_in_recovery()` function. Then, close the connection, establish a connection again, and call the `pg_is_in_recovery()` function again. If you receive a value of `false` and then a value of `true`, Pgpool routes requests to your primary ApsaraDB RDS instance and then to your read-only ApsaraDB RDS instances, and read/write splitting is enabled.

• Q: Does Pgpool increase the latency?

A: Pgpool increases the latency slightly. In the test environment you set up in this topic, the latency increases by about 0.12 milliseconds.

• Q: How does Pgpool check the latency and health on my read-only ApsaraDB RDS instances?

- A: If the WAL replay latency on a read-only ApsaraDB RDS instance exceeds the specified limit, Pgpool stops routing SQL requests to the read-only instance. Pgpool resumes routing SQL requests to the read-only instance only after it detects that the WAL replay latency on the read-only instance falls below the specified limit.

Note Connect to your primary ApsaraDB RDS instance and query the location where the current WAL data record is written. This location is referred to as log sequence number (LSN) 1. Then, connect to a read-only ApsaraDB RDS instance and query the location where the current WAL data record is replayed. This location is referred to as LSN 2. You can obtain the number of bytes between LSN 1 and LSN 2. This number indicates the latency.

- Pgpool monitors the health of your read-only ApsaraDB RDS instances. If a read-only instance is unhealthy, Pgpool stops routing requests to the read-only instance.

- Q: How do I stop Pgpool and reload the configuration of Pgpool?

A: Run the `pgpool --help` command to obtain more information about the commands used in Pgpool. Example:

```
cd /etc/pgpool-II-12
pgpool -f ./pgpool.conf -m fast stop
```

- Q: How do I configure Pgpool if more than one read-only ApsaraDB RDS instance is attached to my primary ApsaraDB RDS instance?

A: Add the configurations of all the attached read-only ApsaraDB RDS instances to the `pgpool.conf` file.

Example:

```
backend_hostname1 = 'xx.xx.xxx.xx'
backend_port1 = 8002
backend_weight1 = 1
backend_data_directory1 = '/data01/pg12_8002/pg_root'
backend_flag1 = 'DISALLOW_TO_FAILOVER'
backend_application_name1 = 'server1'
backend_hostname2 = 'xx.xx.xx.xx'
backend_port1 = 8002
backend_weight1 = 1
backend_data_directory1 = '/data01/pg12_8002/pg_root'
backend_flag1 = 'DISALLOW_TO_FAILOVER'
backend_application_name1 = 'server1'
```

- Q: How do I use `pcp` commands to view the status of my read-only ApsaraDB RDS instances?

A: To obtain the status of your read-only ApsaraDB RDS instances by using `pcp` commands, run the following command:

```
# pcp_node_info -U manage -h /tmp -p 9898 -n 1 -v
Password: # Enter the password.
Hostname      : 127.0.0.1
Port         : 8002
Status       : 2
Weight       : 0.500000
Status Name   : up
Role         : standby
Replication Delay : 0
Replication State :
Replication Sync State :
Last Status Change : 2020-02-29 00:20:29
```

- Q: Which listening ports are used by Pgpool for read/write splitting?

A: The following listening ports are used by Pgpool for read/write splitting:

- Primary ApsaraDB RDS instance: Port 3389
- Secondary ApsaraDB RDS instance: Port 8002
- Pgpool: Port 8001
- PCP: Port 9898

10.17. Use ShardingSphere to develop ApsaraDB RDS for PostgreSQL

ShardingSphere is an open source ecosystem that consists of a set of distributed database middleware solutions.

Prerequisites

All PostgreSQL versions used with ApsaraDB RDS support ShardingSphere.

Context

ApsaraDB RDS for PostgreSQL supports database-integrated sharding plug-ins (such as Citus, Postgres-XC, and AntDB) and massively parallel processing (MPP) products. It also supports sharding middleware products that are similar to those widely used in MySQL, such as ShardingSphere.

ShardingSphere is suitable for services that run in databases with thorough, well-organized logical sharding. It offers the following features:

- Data sharding
 - Database and table sharding
 - Read/write splitting
 - Sharding strategy customization
 - Decentralized distributed primary key
- Distributed transaction
 - Unified transaction API
 - XA transaction
 - BASE transaction
- Database orchestration
 - Dynamic configuration
 - Orchestration and governance
 - Data encryption
 - Tracing and observability
 - Elastic scaling out (planning)

For more information, visit the [ShardingSphere documentation](#).

ShardingSphere products

ShardingSphere includes three independent products. You can choose the product that best suits your business requirements. The following table describes these products.

| Parameter | Sharding-JDBC | Sharding-Proxy | Sharding-Sidecar |
|---------------------------|--|----------------------|----------------------|
| Supported database engine | All JDBC-compatible database engines such as MySQL, PostgreSQL, Oracle, and SQL Server | MySQL and PostgreSQL | MySQL and PostgreSQL |

| Parameter | Sharding-JDBC | Sharding-Proxy | Sharding-Sidecar |
|----------------------------------|-----------------|----------------------|------------------|
| Connections consumed | High | Low | High |
| Supported heterogeneous language | Java | All | All |
| Performance | Low consumption | Moderate consumption | Low consumption |
| Decentralized | Yes | No | Yes |
| Stateless API | No | Yes | No |

Prepare configuration templates

1. On your ECS instance, run the following command to go to the directory where configuration templates are stored. The directory is under the root directory in this example.

```
cd /root/apache-shardingsphere-incubating-4.0.0-sharding-proxy-bin/conf
```

2. Run the `ll` command to view all files stored in the directory: Command output:

```
total 24
-rw-r--r-- 1 501 games 3019 Jul 30 2019 config-encrypt.yaml
-rw-r--r-- 1 501 games 3582 Apr 22 2019 config-master_slave.yaml
-rw-r--r-- 1 501 games 4278 Apr 22 2019 config-sharding.yaml
-rw-r--r-- 1 501 games 1918 Jul 30 2019 server.yaml
```

Note

- config-encrypt.yaml: the data encryption configuration file.
- config-master_slave.yaml: the read/write splitting configuration file.
- config-sharding.yaml: the data sharding configuration file.
- server.yaml: the common configuration file.

3. Modify the configuration files.

Note For more information about the configuration files, visit the [ShardingSphere documentation](#). In this example, the data sharding and common configuration files are used.

- Example of a data sharding configuration file:

```
schemaName: # The name of the logical data source.
dataSources: # The configuration of the data source. You can configure more than one data source by using the data_source_name variable.
<data_source_name>: # You do not need to configure a database connection pool. This is different in Sharding-JDBC.
url: # The URL used to connect to your database.
username: # The username used to log on to the database.
password: # The password used to log on to the database.
connectionTimeoutMilliseconds: 30000 # The connection timeout period in milliseconds.
idleTimeoutMilliseconds: 60000 # The idle-connection reclaiming timeout period in milliseconds.
maxLifetimeMilliseconds: 1800000 # The maximum connection time to live (TTL) in milliseconds.
maxPoolSize: 65 # The maximum number of connections allowed.
shardingRule: # You do not need to configure a sharding rule, because it is the same in Sharding-JDBC.
```

- Example of a common configuration file:

```
Proxy properties
# You do not need to configure proxy properties that are the same in Sharding-JDBC
props:
  acceptor.size: # The number of worker threads that receive requests from the client. The default number is equal to the number of CPU cores multiplied by 2.
  proxy.transaction.type: # The type of transaction processed by the proxy. Valid values: LOCAL | XA | BASE. Default value: LOCAL. Value XA specifies to use Atomikos as the transaction manager. Value BASE specifies to copy the .jar package that implements the ShardingTransactionManager operation to the library.
  proxy.opentracing.enabled: # Specifies whether to enable link tracing. Link tracing is disabled by default.
  check.table.metadata.enabled: # Specifies whether to check the consistency of metadata among sharding tables during startup. Default value: false.
  proxy.frontend.flush.threshold: # The number of packets returned in a batch during a complex query.
Permission verification
This part of the configuration is used to verify your permissions when you attempt to log on to Sharding-Proxy. After you configure the username, password, and authorized databases, you must use the correct username and password to log on to Sharding-Proxy from the authorized databases.
authentication:
  users:
    root: # The username of the root user.
    password: root# The password of the root user.
    sharding: # The username of the sharding user.
    password: sharding# The password of the sharding user.
    authorizedSchemas: sharding_db, masterslave_db # The databases in which the specified user is authorized. If you want to specify more than one database, separate them with commas (,). You are granted the permissions of the root user by default. This way, you can access all databases.
```

Set up a test environment

- On your ECS instance, install Java.

```
yum install -y java
```

- Configure an ApsaraDB RDS instance that runs PostgreSQL 10.
 - Create an account with username r1.
 - Set the password of the account to "PW123321!".

- Create the following databases whose owners are user r1: db0, db1, db2, and db3.
- Add the IP address of your ECS instance to an IP address whitelist of the ApsaraDB RDS for PostgreSQL instance.

 Note

- For more information about how to create an ApsaraDB RDS for PostgreSQL instance, database, and account, see [Create an instance](#) and [Create a database and an account](#).
- For more information about how to configure an IP address whitelist, see [Configure an IP address whitelist](#).

- Run `vi /root/apache-shardingsphere-incubating-4.0.0-sharding-proxy-bin/conf/server.yaml` to configure the following common configuration file:

```
authentication:
users:
  r1:
    password: PW123321!
    authorizedSchemas: db0,db1,db2,db3
props:
  executor.size: 16
  sql.show: false
```

Test horizontal sharding

1. Run `vi /root/apache-shardingsphere-incubating-4.0.0-sharding-proxy-bin/conf/config-sharding.yaml` to modify the following data sharding configuration file:

```
schemaName: sdb
dataSources:
  db0:
    url: jdbc:postgresql://pgm-bpxxxxx.pg.rds.aliyuncs.com:1433/db0
    username: r1
    password: PW123321!
    connectionTimeoutMilliseconds: 30000
    idleTimeoutMilliseconds: 60000
    maxLifetimeMilliseconds: 1800000
    maxPoolSize: 65
  db1:
    url: jdbc:postgresql://pgm-bpxxxxx.pg.rds.aliyuncs.com:1433/db1
    username: r1
    password: PW123321!
    connectionTimeoutMilliseconds: 30000
    idleTimeoutMilliseconds: 60000
    maxLifetimeMilliseconds: 1800000
    maxPoolSize: 65
  db2:
    url: jdbc:postgresql://pgm-bpxxxxx.pg.rds.aliyuncs.com:1433/db2
    username: r1
    password: PW123321!
    connectionTimeoutMilliseconds: 30000
    idleTimeoutMilliseconds: 60000
    maxLifetimeMilliseconds: 1800000
    maxPoolSize: 65
  db3:
    url: jdbc:postgresql://pgm-bpxxxxx.pg.rds.aliyuncs.com:1433/db3
    username: r1
    password: PW123321!
    connectionTimeoutMilliseconds: 30000
```

```
idleTimeoutMilliseconds: 60000
maxLifetimeMilliseconds: 1800000
maxPoolSize: 65
shardingRule:
  tables:
    t_order:
      actualDataNodes: db${0..3}.t_order${0..7}
      databaseStrategy:
        inline:
          shardingColumn: user_id
          algorithmExpression: db${user_id % 4}
      tableStrategy:
        inline:
          shardingColumn: order_id
          algorithmExpression: t_order${order_id % 8}
      keyGenerator:
        type: SNOWFLAKE
        column: order_id
    t_order_item:
      actualDataNodes: db${0..3}.t_order_item${0..7}
      databaseStrategy:
        inline:
          shardingColumn: user_id
          algorithmExpression: db${user_id % 4}
      tableStrategy:
        inline:
          shardingColumn: order_id
          algorithmExpression: t_order_item${order_id % 8}
      keyGenerator:
        type: SNOWFLAKE
        column: order_item_id
  bindingTables:
    - t_order,t_order_item
  defaultTableStrategy:
    none:
```

2. Start ShardingSphere and listen to Port 8001.

```
cd /root/apache-shardingsphere-incubating-4.0.0-sharding-proxy-bin/bin/
./start.sh 8001
```

3. Connect to the destination database.

```
psql -h 127.0.0.1 -p 8001 -U r1 sdb
```

4. Create a table.

```
create table t_order(order_id int8 primary key, user_id int8, info text, c1 int, crt_time timestamp);
create table t_order_item(order_item_id int8 primary key, order_id int8, user_id int8, info text, c1 int, c2 int, c3 int,
c4 int, c5 int, crt_time timestamp);
```

 **Note** When you create a table, the system creates horizontal shards in the destination database based on the sharding strategy that you specify.

FAQ

- If you want to view the SQL parsing and routing statements used in ShardingSphere, run `vi /root/apache-shardingsphere-incubating-4.0.0-sharding-proxy-bin/conf/server.yaml`.

```

authentication:
users:
  r1:
    password: PW123321!
    authorizedSchemas: db0,db1,db2,db3
props:
  executor.size: 16
  sql.show: true # Specifies to log parsed SQL statements.

```

- If you want to test writes and queries, run the following commands:

```

insert into t_order (user_id, info, c1, crt_time) values (0,'a',1,now());
insert into t_order (user_id, info, c1, crt_time) values (1,'b',2,now());
insert into t_order (user_id, info, c1, crt_time) values (2,'c',3,now());
insert into t_order (user_id, info, c1, crt_time) values (3,'c',4,now());
select * from t_order;

```

The following result is returned in this example:

| order_id | user_id | info | c1 | crt_time |
|--------------------|---------|------|----|----------------------------|
| 433352561047633921 | 0 | a | 1 | 2020-02-09 19:48:21.856555 |
| 433352585668198400 | 1 | b | 2 | 2020-02-09 19:48:27.726815 |
| 433352610813050881 | 2 | c | 3 | 2020-02-09 19:48:33.721754 |
| 433352628370407424 | 3 | c | 4 | 2020-02-09 19:48:37.907683 |

(4 rows)

- If you want to view ShardingSphere logs, run the following command:

```
/root/apache-shardingsphere-incubating-4.0.0-sharding-proxy-bin/logs/stdout.log
```

- If you want to use pgbench for stress testing, run the following commands:

```

vi test.sql
\set user_id random(1,100000000)
\set order_id random(1,2000000000)
\set order_item_id random(1,2000000000)
insert into t_order (user_id, order_id, info, c1, crt_time) values (:user_id, :order_id,random()::text, random()*1000, now()) on conflict (order_id) do update set info=excluded.info,c1=excluded.c1,crt_time=excluded.crt_time;
insert into t_order_item (order_item_id, user_id, order_id, info, c1,c2,c3,c4,c5,crt_time) values (:order_item_id, :user_id,:order_id,random()::text, random()*1000,random()*1000,random()*1000,random()*1000,random()*1000, now()) on conflict(order_item_id) do nothing;
pgbench -M simple -n -r -P 1 -f ./test.sql -c 24 -j 24 -h 127.0.0.1 -p 8001 -U r1 sdb -T 120
progress: 1.0 s, 1100.9 tps, lat 21.266 ms stddev 6.349
progress: 2.0 s, 1253.0 tps, lat 18.779 ms stddev 7.913
progress: 3.0 s, 1219.0 tps, lat 20.083 ms stddev 13.212

```

11. Cloud Native Distributed Database PolarDB-X

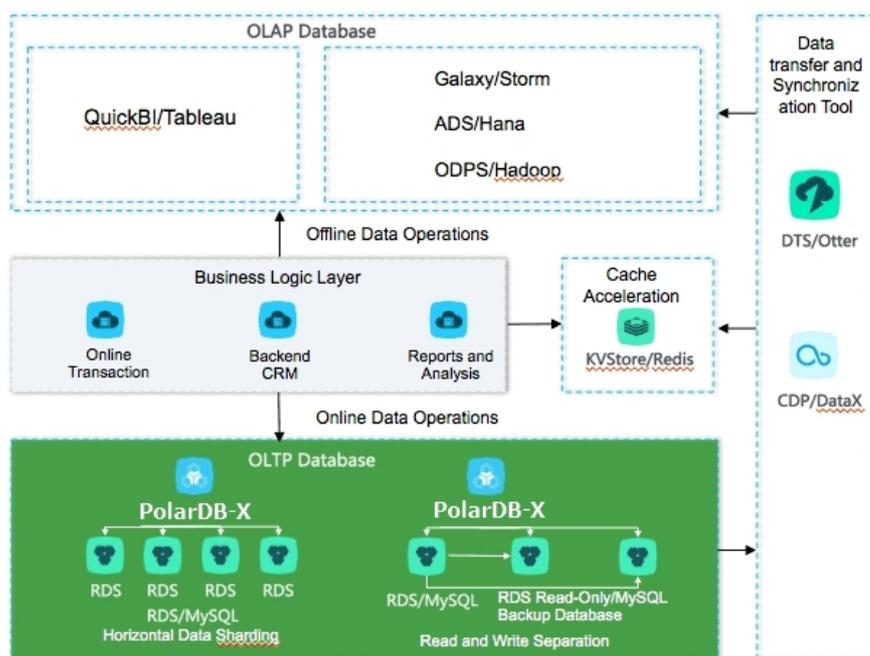
11.1. What is PolarDB-X?

Cloud Native Distributed Database PolarDB-X is a database product that is developed by Alibaba Group and focuses on scaling out single-instance relational databases. This service is compatible with former Distributed Relational Database Service (DRDS).

Compatible with the MySQL communication protocols, PolarDB-X supports most MySQL data manipulation language (DML) and data definition language (DDL) syntax. It provides the core capabilities and features of distributed databases, such as sharding, smooth scale-out, service upgrade and downgrade, and transparent read/write splitting. PolarDB-X is lightweight (stateless), flexible, stable, and efficient, and provides you with O&M capabilities throughout the lifecycle of distributed databases.

PolarDB-X is used for operations on large-scale online data. PolarDB-X maximizes the operation efficiency by partitioning data in specific business scenarios. This meets the requirements of online business on relational databases in an effective way.

Figure of the PolarDB-X architecture



Fixed issues

- Capacity bottlenecks of single-instance databases: As the data volume and access increase, traditional single-instance databases encounter great challenges that cannot be solved by hardware upgrades in a complete way. In distributed database solutions, multiple instances work in a joint way. This resolves the bottlenecks of data storage capacity and access volumes in an effective way.
- Difficult scale-out of relational databases: Due to the inherent attributes of distributed databases, you can change the shards where data is stored through smooth data migration. This way, the dynamic scale-out of relational databases is achieved.

11.2. Quick start

This topic describes how to get started with Cloud Native Distributed Database PolarDB-X.

A PolarDB-X instance is physically a distributed cluster that consists of multiple PolarDB-X server nodes and underlying storage instances. A PolarDB-X database is a logical concept and only contains metadata. Specific data is stored in the physical database of the underlying storage instance. To get started with PolarDB-X, follow these steps:

1. [Create a PolarDB-X instance.](#)
2. [Create a database.](#)

To create a database in a PolarDB-X instance, you must select one or more ApsaraDB RDS for MySQL instances as the data storage nodes. If no RDS instance exists, create one first. For more information about how to create and manage ApsaraDB RDS for MySQL instances, see *User Guide of RDS*.

3. After a PolarDB-X database is created, you also need to create tables in the PolarDB-X database like in a single-instance database. However, the syntax is different, mainly in the expression of data partitioning information in the PolarDB-X table creation statement. For more information about how to create a table, see [Table creation syntax](#).

11.3. Log on to the PolarDB-X console

This topic describes how to log on to the Cloud Native Distributed Database PolarDB-X console by using Google Chrome.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

 **Note** When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top navigation bar, choose **Products > Database Services > Distributed Relational Database Service**.

11.4. Instance management

11.4.1. Create an instance

Before you use PolarDB-X, create an instance. This topic describes how to create a PolarDB-X instance.

1. [Log on to the PolarDB-X console.](#)
2. In the upper-right corner of the page, click **Create Instance**.

3. On the **Create PolarDB-X Instance** page, configure the parameters as needed.

[Parameters for creating an instance](#) describes the parameters.

Parameters for creating an instance

| Category | Parameter | Description |
|----------------|----------------|--|
| Region | Organization | The organization to which the instance belongs. |
| | Resource Set | The resource set to which the instance belongs. |
| | Region | The region in which the instance is deployed. The instances in different regions cannot communicate with each other over an internal network. After an instance is created, you cannot change the region for the instance. |
| | Zone | The zone in which the instance is deployed. |
| Basic Settings | Instance Type | The type of the instance. Select an instance type from the options that are available on the page. |
| | Edition | The edition of the instance. Valid values: <ul style="list-style-type: none"> ◦ Standard ◦ Enterprise ◦ Starter |
| | Specifications | The specifications of the instance. The specifications vary based on instance editions. Select the instance specifications from the options that are available on the page. |
| Network Type | Network Type | The network type of the PolarDB-X instance. Valid values: <ul style="list-style-type: none"> ◦ Classic Network: Cloud services on the classic network are not isolated from each other. Unauthorized access to a cloud service is blocked only by the security group or based on the whitelist policy of the service. ◦ VPC: You can create a virtual private cloud (VPC) to build an isolated network environment on Apsara Stack. You can customize route tables, IP address ranges, and gateways in a VPC. We recommend that you select a VPC for improved security. If you set Network Type to VPC, you must set VPC and vSwitch. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note The PolarDB-X instance and the Elastic Compute Service (ECS) instance to connect must use the same network type. Otherwise, they cannot communicate with each other over an internal network. </div> |
| | VPC | Select a VPC. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note If you set Network Type to VPC, specify this parameter. </div> |
| | vSwitch | Select a vSwitch in the VPC. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note If you set Network Type to VPC, specify this parameter. </div> |

4. Click **Submit**.

After the instance is created, the instance appears in the instance list and the status of the instance changes to **Running**. The instance name is a unique identifier of an instance. You can identify a PolarDB-X instance based on this unique identifier.

11.4.2. Change instance specifications

When you use PolarDB-X, you can change the specifications of a PolarDB-X instance as needed.

Procedure

1. [Log on to the PolarDB-X console](#).
2. Find the target instance in the instance list.
3. Click the target instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the **Common Actions** section or the **Configuration Information** section, click **Upgrade** or **Downgrade** to go to the Change Specifications page.

 **Note** Alternatively, on the **DRDS Instance Management** page, choose **More > Downgrade** in the **Actions** column of the instance.

5. On the Change Specifications page, set **Instance Edition** and **Instance Specifications**, and click **Submit**. After a few minutes, you can view the new specifications of the PolarDB-X instance in the instance list.

 **Note** Specifications downgrade leads to transient disconnections between applications and PolarDB-X within a short period of time. Make sure that your applications can be automatically reconnected.

11.4.3. Read-only PolarDB-X instances

11.4.3.1. Overview

Read-only PolarDB-X instances are extension and supplement to primary PolarDB-X instances and are compatible with SQL query syntax of primary PolarDB-X instances.

Features

Read-only and primary PolarDB-X instances can share the same replica of data. You can perform complex data query and analysis directly on read-only or primary ApsaraDB RDS for MySQL instances. Multiple instance types are provided to handle highly concurrent access requests and reduce the response time (RT) for complex queries. Resource isolation alleviates the load pressure on the primary instances and reduces the link complexity of the business architecture. It reduces the O&M and budget costs, eliminating the need for additional data synchronization.

Instance type

Concurrent read-only instances: For high-concurrency and high-traffic simple queries or offline data extraction, resource isolation protects you against highly concurrent queries, ensuring the stability of online business links.

Note For the businesses with primary PolarDB-X instances, concurrent read-only instances can be used in the following scenarios:

- High-concurrency and high-traffic simple queries are performed.
- Data is extracted offline.

Limits

- Primary and read-only PolarDB-X instances must be in the same region, but they can be in different zones.
- A read-only PolarDB-X instance must belong to a primary PolarDB-X instance. Before creating a read-only instance, you must create a primary instance. After you create a database on the primary instance, the database is replicated to the read-only instance. If you delete the database from the primary instance, the corresponding database on the read-only instance is also deleted.
- You are not allowed to migrate data to read-only PolarDB-X instances.
- You are not allowed to create or delete databases in PolarDB-X read-only instances.
- PolarDB-X read-only instances cannot be cloned.
- PolarDB-X read-only instances support data definition language (DDL) statements but do not support data manipulation language (DML) statements for data modification.

11.4.3.2. Create a read-only PolarDB-X instance

This topic describes how to create a read-only Cloud Native Distributed Database PolarDB-X (PolarDB-X) instance.

Procedure

1. [Log on to the PolarDB-X console.](#)
2. Find the target instance in the instance list.
3. Click the target instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. On the Basic Information page, click **Create DRDS Read-only Instance** in the Related Instances section.
5. Set Region, Basic Settings, and Network Type, and then click **Submit**.

Parameters for creating a read-only PolarDB-X instance

| Type | Parameter | Description |
|----------------|-------------------------|---|
| Region | Region | The region where the read-only instance resides. Services in different regions are not interconnected over the internal network. After the instance is created, the region cannot be changed. |
| | Zone | The zone where the read-only instance resides. |
| Basic Settings | Instance Type | The type of the read-only instance. Select an instance type from the options available on the page. |
| | Instance Edition | The edition of the read-only instance. Valid values: <ul style="list-style-type: none"> ◦ Starter ◦ Standard ◦ Enterprise |
| | Instance Specifications | The specifications of the read-only instance. The rules vary with instance editions. Select the instance specifications from the options available on the page. |

| Type | Parameter | Description |
|---------------------|---------------------|---|
| | Description | The description of the read-only instance. We recommend that you provide an informative description to simplify future management operations. |
| Network Type | Network Type | <p>The network type of the read-only instance. PolarDB-X instances support the following network types:</p> <ul style="list-style-type: none"> ◦ Classic Network: Cloud services on a classic network are not isolated from each other. Unauthorized access to a cloud service is blocked only by the security group or whitelist policy of the service. ◦ VPC: A virtual private cloud (VPC) helps you build an isolated network environment on Apsara Stack. You can customize routing tables, IP address ranges, and gateways in a VPC. We recommend that you select VPC for higher security. Select VPC for Network Type, and then set VPC and VSwitch. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note Make sure that the PolarDB-X instance has the same network type as the Elastic Compute Service (ECS) instance to which you want to connect. If the PolarDB-X and ECS instances have different network types, they cannot communicate over an internal network.</p> </div> |

6. It takes several minutes to create the instance. Please wait. After the instance is created, it appears in the instance list in the PolarDB-X console.

11.4.3.3. Manage a read-only PolarDB-X instance

Read-only PolarDB-X instances are managed in a similar way as primary instances. However, databases cannot be created or deleted on the read-only instance management page. Databases on read-only instances are created or deleted with databases on primary instances. In the PolarDB-X console, you can go to the read-only instance management page in two ways.

Manage a read-only PolarDB-X instance by its ID

1. [Log on to the PolarDB-X console.](#)
2. On the **DRDS Instance Management** page, find the target read-only instance.
3. Click the target instance ID or choose **More > Manage** from the Actions column of the target read-only instance to access the **Basic Information** page.

Manage a read-only PolarDB-X instance by the ID of its primary instance

1. [Log on to the PolarDB-X console.](#)
2. On the **DRDS Instance Management** page, find the target primary instance.
3. Click the target instance ID or choose **More > Manage** from the Actions column of the target primary instance to access the **Basic Information** page.
4. On the **Basic Information** page, move the pointer over the number of read-only instances in the **Related Instances** section to view the ID of the read-only PolarDB-X instance.
5. Click the ID of the target read-only PolarDB-X instance. The **Basic Information** page of the read-only instance appears.

11.4.3.4. Release a read-only PolarDB-X instance

If you no longer need a read-only PolarDB-X instance, you can release it.

Prerequisites

The read-only instance must be in the **Running** state.

Procedure

1. [Log on to the PolarDB-X console.](#)
2. Find the target instance in the instance list.
3. In the PolarDB-X instance list, find the target instance, and choose **More > Release** from the Actions column.

 **Notice** You cannot recover the PolarDB-X instances that have been released. Exercise caution when you perform this operation.

4. In the **Release DRDS Instance** dialog box, click **OK**.

11.4.4. Restart a PolarDB-X instance

This topic describes how to restart a PolarDB-X instance.

Prerequisites

The PolarDB-X instance must be in the **Running** state.

Procedure

1. [Log on to the PolarDB-X console.](#)
2. Find the target instance in the instance list.
3. Click the target instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. Click **Restart Instance** in the upper-right corner.
5. In the **Restart Instance** dialog box, click **OK**.

 **Notice** Restarting a PolarDB-X instance terminates all its connections. Make appropriate service arrangements before you restart a PolarDB-X instance. Exercise caution when you perform this operation.

11.4.5. Release a PolarDB-X instance

This topic describes how to release a running PolarDB-X instance in the PolarDB-X console.

Prerequisites

- All databases on the PolarDB-X instance have been deleted.
- The PolarDB-X instance must be in the **Running** state.

Procedure

1. [Log on to the PolarDB-X console.](#)
2. Find the target instance in the instance list.
3. In the PolarDB-X instance list, find the target instance, and choose **More > Release** from the Actions column.
4. In the **Release DRDS Instance** dialog box, click **OK**.

 **Warning** After the PolarDB-X instance is released, data is not deleted from its attached ApsaraDB RDS for MySQL instances. However, a released PolarDB-X instance cannot be restored. Exercise caution when you perform this operation.

11.4.6. Recover data

11.4.6.1. Backup and restoration

PolarDB-X allows you to back up data of instances and databases and restore them by using the backup data. Instances can be automatically or manually backed up. PolarDB-X provides fast backup and consistent backup. Existing backup sets are used to restore data in instances. Data is restored to the new PolarDB-X and ApsaraDB RDS for MySQL instances by using the existing backup sets.

Considerations

- By default, the automatic backup policy of PolarDB-X is disabled. You must manually enable it.
- The log backup capability of PolarDB-X relies on underlying ApsaraDB RDS for MySQL instances. Therefore, the log backup policy configured in the PolarDB-X console is automatically synchronized to all underlying ApsaraDB RDS for MySQL instances. After the policy is configured, do not modify it in the ApsaraDB RDS console. Otherwise, related data backup sets may be invalid.
- The backup and restoration feature of PolarDB-X relies on log backup. We recommend that you enable the log backup policy by default to prevent backup sets from becoming invalid.
- Data definition language (DDL) operations cannot be performed during the backup process. Otherwise, instance backup and restoration may fail.
- During data backup, ensure that the underlying ApsaraDB RDS for MySQL instances for the PolarDB-X instance are normal. Otherwise, data backup may fail.
- Consistent backup and restoration is supported only by PolarDB-X 5.3.8 and later.
- Ensure that all tables have primary keys. Otherwise, data accuracy may be affected during consistent backup and restoration.
- During consistent backup, distributed transactions on PolarDB-X instances are locked for seconds. During the locking period, the execution of non-transactional SQL statements and non-distributed transactions is not affected. However, the commitment of distributed transactions is blocked and the response time (RT) for executing SQL statements may have millisecond-level jitters. We recommend that you perform consistent backup during off-peak hours.
- Due to changes in the inventory of PolarDB-X and ApsaraDB RDS for MySQL resources, PolarDB-X automatically adjusts the instance type and zone during instance restoration. We recommend that you confirm and adjust the instance type and zone after the instance restoration to avoid business disruption.

Backup methods

For different scenarios, PolarDB-X provides fast backup and consistent backup and the related data restoration capabilities. The following table compares the two backup methods.

| Backup method | Scenario | Benefit | Disadvantage |
|---------------|--|---|---|
| Fast backup | Applies to routine backup and restoration scenarios. | <ul style="list-style-type: none"> • Fast data backup and restoration is enabled. • Data can be restored by backup set or by time. • All PolarDB-X instance versions support this feature. | In sharding scenarios, data consistency can be ensured only within a single ApsaraDB RDS for MySQL instance. Global data consistency cannot be ensured. |

| Backup method | Scenario | Benefit | Disadvantage |
|-------------------|---|--|--|
| Consistent backup | Applies to backup and restoration for the financial industry and online core transactions that require high data consistency. | In sharding scenarios, global data consistency is ensured. | <ul style="list-style-type: none"> Backup and restoration is slow. Data can be restored by backup set, but cannot be restored by time. Only PolarDB-X 5.3.8 and later support this feature. During data backup, distributed transactions on PolarDB-X instances are locked for seconds. During the locking period, the RT for executing SQL statements may have millisecond-level jitters. Therefore, we recommend that you perform consistent backup during off-peak hours. |

11.4.6.2. Configure an automatic backup policy

PolarDB-X provides the automatic backup feature. This topic describes how to configure an automatic backup policy.

Procedure

1. [Log on to the PolarDB-X console.](#)
2. Find the target instance in the instance list.
3. Click the target instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Data Recovery > Backup and Recovery**.
5. On the page that appears, choose **Backup Policy > Edit**.
6. In the **Backup Policy** dialog box, set parameters as needed, and click **OK**.

11.4.6.3. Configure local logs

You can use local logs and the backup and restoration feature or the SQL flashback feature of PolarDB-X to accurately restore an instance or a database to the desired point in time. This topic describes how to configure local logs.

Procedure

1. [Log on to the PolarDB-X console.](#)
2. Find the target instance in the instance list.
3. Click the target instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Data Recovery > Backup and Recovery**.

5. Click the **Local Log Settings** tab and then click **Edit**.
6. In the **Local Binlog Settings** dialog box, specify the parameters based on the business requirements and click **OK**.

 **Notice** The local log settings are applied to all the underlying ApsaraDB RDS for MySQL instances.

11.4.6.4. Manual backup

PolarDB-X also provides the manual backup capability, so that you can back up data at any time. This topic describes how to manually back up instances and databases.

Procedure

1. [Log on to the PolarDB-X console](#).
2. Find the target instance in the instance list.
3. Click the target instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Data Recovery > Backup and Recovery**.
5. On the page that appears, click **Data Backup** on the right.
6. In the dialog box that appears, set Backup Method and Backup Level.
 - Backup Method can be set to **Fast Backup** and **Consistent Backup**. For more information about differences between the two methods, see [Backup methods](#).

 **Notice** If you select Consistent Backup, distributed transactions are locked within seconds and the response time (RT) may vary by sub-seconds. Therefore, we recommend that you perform this operation during off-peak hours.

- Backup Level can be set to **Instance Backup** or **Database Backup**. You can select **Instance Backup** to back up the entire instance, or select **Database Backup** to back up a database as needed.
7. Click **OK**.

11.4.6.5. Recover data

You can use the data recovery feature of PolarDB-X to recover an instance or a database to the time when the backup is created. You can perform this operation at any time. This topic describes how to recover the data of an instance or a database to a specific point in time.

Procedure

1. [Log on to the PolarDB-X console](#).
2. Find the target instance in the instance list.
3. Click the target instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Data Recovery > Backup and Recovery**.
5. On the page that appears, click **Data Recovery (Original Clone Instance)** on the right.
6. Select a recovery method
 - **By Time**: Recover data to the selected point in time. You must set **Restoration Time** and **Recovery Level**.
 - **By Backup Set**: Recover data from the selected backup file.

 **Note** You can also click **Recover** in the Actions column of the target backup set to recover data by backup set

7. Click **Precheck** to check whether a valid backup set is available for data recovery. If the precheck fails, the data cannot be restored.
8. Click **Enable** to access the order confirmation page.
9. Confirm the order details and then click **Enable** to recover the data. You can view the data recovery progress in **Task Progress** in the upper-right corner of the page.

11.4.6.6. SQL flashback

11.4.6.6.1. Overview

PolarDB-X provides the SQL flashback feature to recover data of particular rows.

When you mistakenly run an SQL statement such as INSERT, UPDATE, or DELETE on PolarDB-X, provide the relevant SQL information to match the event in the binary log file and generate the corresponding recovery file. You can download the file and recover data as needed. SQL flashback automatically chooses **fuzzy match** or **exact match** to locate lost data caused by the error. For more information, see [Exact match and fuzzy match](#) and [Rollback SQL statements and original SQL statements](#).

Features

- Easy-to-use: SQL flashback allows you to retrieve the lost data by entering required information about the corresponding SQL statement.
- Fast and lightweight: Regardless of the backup policy of ApsaraDB RDS for MySQL instances, you only need to enable log backup before an SQL statement error occurs.
- Flexible recovery: Rollback SQL statements and original SQL statements are available for different scenarios.
- Exact match: SQL flashback supports exact match of data about the corresponding SQL statement, which improves precision of data recovery.

Limits

- SQL flashback depends on the binary log retention time and the log backup feature of ApsaraDB RDS for MySQL must be enabled. Binary log files can be retained only for a certain period. Use SQL flashback to generate files for recovery as soon as possible when an error occurs.
- The recovery files generated by SQL flashback are retained for seven days by default, and you need to download these files as soon as possible.
- The following conditions must be met for SQL flashback exact match:
 - The PolarDB-X instance version is 5.3.4-15378085 or later.
 - The version of the ApsaraDB RDS for MySQL instance used by the PolarDB-X database is 5.6 or later.
 - SQL flashback exact match is enabled before the error SQL statement is executed.
 - The TRACE_ID information for the error SQL statement is provided.
- To ensure the precision of data recovery, the exact match feature is enabled by default for the database created in a PolarDB-X instance of 5.3.4-15378085 or later. After this feature is enabled, SQL execution information is included in the binary log file by default, which requires more storage space for ApsaraDB RDS for MySQL instances. If you need to use the exact match feature, we recommend that you upgrade PolarDB-X before enabling the feature. For more information, see [Enable exact match](#).

11.4.6.6.2. Generate a recovery file

Procedure

1. [Log on to the PolarDB-X console.](#)
2. Find the target instance in the instance list.
3. Click the target instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Data Recovery > SQL Flashback**. The **SQL Flashback** page appears.
5. On the **SQL Flashback** page, enter the basic information about a mistaken SQL statement, including Database, Time Range, Table Name, TRACE_ID, and SQL Statement Type. The following table describes the parameters.

| Parameter | Description |
|--------------------|--|
| Database | The database where the mistaken SQL statement was executed. |
| Time Range | The time range during which the mistaken SQL statement was executed. The start time is earlier than the start time when the mistaken SQL statement was executed, whereas the end time is later than the time when the execution of the mistaken SQL statement ended. To ensure efficient recovery, we recommend that you limit the time range to five minutes. |
| Table Name | The name of the table on which the mistaken SQL statement was executed. This parameter is optional. |
| TRACE_ID | The unique TRACE_ID that PolarDB-X allocates for each executed SQL statement. You can obtain the TRACE_ID of the mistaken SQL statement by using the SQL audit feature of PolarDB-X. |
| SQL Statement Type | The type of the mistaken SQL statement. Valid values: <ul style="list-style-type: none"> ◦ INSERT ◦ UPDATE ◦ DELETE |

6. Click **Precheck**. The system checks whether a binary log file exists within the specified time range. For more information about binary log files, see [Configure local logs](#).

Note

- If no binary log file exists within specified the time range, the precheck fails and the system cannot recover the data for you.
- If a binary log file exists within the specified time range, the precheck is successful and you can go to the next step.

7. Set SQL Statement Type for Recovery to **Rollback SQL** or **Original SQL Statement**. For more information about differences between the two methods, see [Rollback SQL statements and original SQL statements](#).
8. Click **Generate SQL** to generate an SQL flashback task. The statuses of the SQL flashback tasks that are running on the current instance appear at the bottom of the page.

What's next

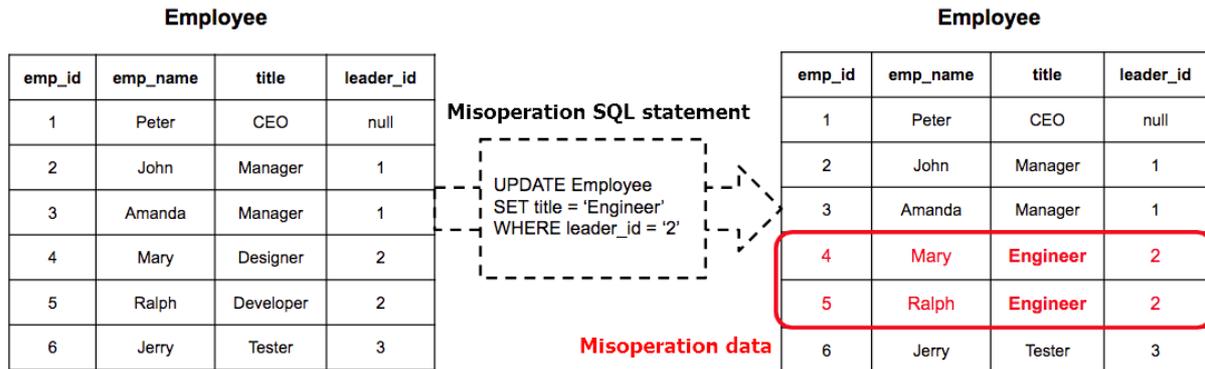
After an SQL flashback task is completed, the task information such as the exact match status and the number of recovered rows appears. You can click **Download** next to the target SQL flashback task to download the corresponding recovery file.

 **Notice** By default, the recovery file is retained for seven days. Download it as soon as possible.

11.4.6.6.3. Rollback SQL statements and original SQL statements

To support different business scenarios, PolarDB-X SQL flashback provides rollback SQL statements and original SQL statements. Before generating an SQL statement for recovering data, you must select a corresponding recovery method based on your scenario.

Recovery methods



| Recovery method | Description | Example |
|------------------------|--|---|
| Rollback SQL statement | Traverses the events in the binary log file in reverse order to reverse the INSERT, UPDATE, and DELETE events. <ul style="list-style-type: none"> The reverse of INSERT is equivalent to DELETE. The reverse of DELETE is equivalent to INSERT. The reverse of UPDATE is equivalent to the value before UPDATE. | <pre>UPDATE Employee SET title = 'Developer' WHERE emp_id = '5'</pre> <pre>UPDATE Employee SET title = 'Designer' WHERE emp_id = '4'</pre> |
| Original SQL statement | Traverses the events in the binary log file in order to mirror all records of the INSERT, UPDATE, and DELETE events. <ul style="list-style-type: none"> An INSERT mirror is equivalent to INSERT. A DELETE mirror is equivalent to INSERT. An UPDATE mirror is equivalent to the value before INSERT. | <pre>INSERT INTO Employee(emp_id,emp_name,title,leader_id) values('4','Mary','Designer','2')</pre> <pre>INSERT INTO Employee(emp_id,emp_name,title,leader_id) values('5','Ralph','Developer','2')</pre> |

11.4.6.6.4. Exact match and fuzzy match

SQL flashback supports **exact match** and **fuzzy match** for binary log events. You do not need to select a match policy. SQL flashback automatically detects and selects the optimal match policy, and notifies you when the flashback task is completed.

| Match mode | Description | Advantage | Disadvantage |
|-------------|---|--|--|
| Exact match | The system performs exact match on the event of a mistaken SQL statement in the binary log file and generates a recovery file. | The recovery file contains only data that is deleted or modified by the mistaken SQL statement. You can use the file directly to ensure the precision and efficiency of data recovery. | The following requirements must be met: <ul style="list-style-type: none"> The PolarDB-X instance is Version 5.3.4-15378085 or later. The version of the ApsaraDB RDS for MySQL instance used by the PolarDB-X database is Version 5.6 or later. You have enabled exact match of SQL flashback before the mistaken SQL statement is executed. You must provide the TRACE_ID of the mistaken SQL statement. |
| Fuzzy match | The system matches the information about the mistaken SQL statement in the binary log file, including the time range, table name, and SQL statement type. Then, the system generates a recovery file. | Fuzzy match is supported for all instances, regardless of the instance version or parameter settings. | Data that is deleted or modified by the mistaken SQL statement cannot be accurately matched. The recovery file contains data changes made by other business SQL operations. You must filter the required data. |

Enable exact match

 **Note** Fuzzy match is enabled by default.

1. Log on to PolarDB-Xconsole, and go to the parameter settings page of the specified instance. For more information, see [Set parameters](#).
2. Change the value of `ENABLE_SQL_FLASHBACK_EXACT_MATCH` to `ON`.

11.4.6.7. Table recycle bin

11.4.6.7.1. Overview

The table recycle bin of PolarDB-X allows you to recover mistakenly deleted tables.

After the table recycle bin is enabled for your PolarDB-X database, the tables that are deleted by using the DROP TABLE statement are moved to the recycle bin and are no longer visible to you. After the tables are moved to the recycle bin for two hours, they are automatically cleared and cannot be recovered. You can view, recover, and clear the deleted tables in the recycle bin.

Limits and notes

- The table recycle bin feature is only supported by PolarDB-X 5.3.3-1670435 and later. For more information, see [View the instance version](#).
- The table recycle bin is disabled for your PolarDB-X database by default. For more information about how to

enable it, see [Enable the table recycle bin](#).

- The table recycle bin of PolarDB-X does not support the recovery of tables deleted by the TRUNCATE TABLE command.
- Tables in the recycle bin still occupy the storage space of ApsaraDB RDS for MySQL before they are automatically cleared. To release the storage space as soon as possible, you can access the recycle bin to manually delete them.

11.4.6.7.2. Enable the table recycle bin

This topic describes how to enable the table recycle bin.

Procedure

1. [Log on to the PolarDB-X console](#).
2. On the **DRDS Instance Management** page, find the target instance.
3. Click the instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Data Recovery > Table Recycle Bin**. The **Table Recycle Bin** page appears.
5. On the top of the **Table Recycle Bin** page, click the tab of the database for which the table recycle bin needs to be enabled.
6. Click **Enable**.
7. In the dialog box that appears, click **OK**.

11.4.6.7.3. Recover tables

This topic describes how to recover your tables from the table recycle bin.

Procedure

1. [Log on to the PolarDB-X console](#).
2. On the **DRDS Instance Management** page, find the target instance.
3. Click the instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Data Recovery > Table Recycle Bin**. The **Table Recycle Bin** page appears.
5. On the top of the **Table Recycle Bin** page, click the tab of the database in which the tables need to be recovered.
6. Click **Restore** in the Actions column of the target table.

11.4.6.7.4. Delete tables from the recycle bin

This topic describes how to delete unnecessary tables from the table recycle bin.

Procedure

1. [Log on to the PolarDB-X console](#).
2. On the **DRDS Instance Management** page, find the target instance.
3. Click the instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Data Recovery > Table Recycle Bin**. The **Table Recycle Bin** page appears.

- On the top of the **Table Recycle Bin** page, click the tab of the database in which the tables need to be cleared.
- Click **Delete** in the **Actions** column of the target table.

11.4.6.7.5. Disable the table recycle bin

If you no longer need the table recycle bin, you can disable it.

Procedure

- Log on to the [PolarDB-X console](#).
- On the **DRDS Instance Management** page, find the target instance.
- Click the instance ID or choose **More > Manage** from the **Actions** column of the instance to access the **Basic Information** page.
- In the left-side navigation pane, choose **Data Recovery > Table Recycle Bin**. The **Table Recycle Bin** page appears.
- On the top of the **Table Recycle Bin** page, click the tab of the database for which the table recycle bin needs to be disabled.
- Click **Disable** to disable the table recycle bin for the database.

11.4.7. Set parameters

PolarDB-X allows you to set parameters for instances and databases. You can view and modify parameter values in the PolarDB-X console based on business needs.

 **Note** Parameters cannot be set for read-only PolarDB-X instances.

Procedure

- Log on to the [PolarDB-X console](#).
- On the **DRDS Instance Management** page, find the target instance.
- Click the instance ID or choose **More > Manage** from the **Actions** column of the instance to access the **Basic Information** page.
- In the left-side navigation pane, choose **Diagnostics and Optimization > Parameter Settings**. Click the **Instance** or **Database** tab to view parameters that you can modify for instances and databases, respectively. For more information about the parameters, see [Parameter description](#).
- Click  next to the parameter you want to modify, enter the target value, and click **OK**.
- Click **Submit** in the upper-right corner to commit the modification.

 **Note** To undo parameter modification, click **Cancel** in the upper-right corner.

Parameter description

| Parameter | Level | Description |
|---------------------------|----------|---|
| Slow SQL threshold | Instance | The threshold for slow SQL statements. SQL statements whose thresholds exceed this threshold are recorded in logical slow SQL logs. |
| Logical idle link timeout | Instance | The logical timeout period of the idle connection between user applications and PolarDB-X (unit: ms). |

| Parameter | Level | Description |
|--|----------|---|
| Maximum package size | Instance | The maximum network packet for the interaction between user applications and PolarDB-X (unit: byte). |
| Instance memory pool size limit | Instance | The maximum size of the memory pool for an instance. If the memory usage on an instance exceeds the value, an error is reported and the query ends. |
| Whether to prohibit all table deletion/update | Database | Specifies whether to disable full table deletion or update. |
| Whether to open the recycle bin | Database | Specifies whether to enable the recycle bin for storing deleted PolarDB-X logical tables. |
| Temporary table size | Database | The size of the temporary table used during distributed queries in PolarDB-X (unit: row). |
| Number of join tables | Database | The maximum number of table shards that can be combined through JOIN when you query multiple table shards in a database. |
| Physical SQL timeout | Database | The timeout period of SQL statements for interaction between PolarDB-X and ApsaraDB RDS for MySQL (unit: ms). The value 0 indicates the timeout period is not limited. |
| SQL exact flashback switch | Database | Specifies whether to support SQL flashback exact match. It is disabled by default. After it is enabled, information about the queries is added to the binary log file used by the PolarDB-X database. |
| Whether to enable logical INFORMATION_SCHEMA A query | Database | Specifies whether to enable logical INFORMATION_SCHEMA query (not relying on the shadow database but returning the aggregation results of logical databases and tables). When it is disabled, the original status is restored (relying on the shadow database and returning the physical database and table information). |
| Transaction log cleanup start time period | Database | The period during which transaction log cleanup starts at a random time. |
| Library-level memory pool size limit | Database | The maximum size of the database-level memory pool. When the memory usage of a PolarDB-X database exceeds this value, an error is reported and the query terminates. The value -1 indicates no limit. |
| Query-level memory pool size limit | Database | The maximum size of the query-level memory pool. When the memory usage of a query exceeds this value, an error is reported and the query terminates. The value -1 indicates no limit. |
| Whether CBO is enabled | Database | Specifies whether to enable the cost-based optimizer (CBO), including features such as Join Reorder and Hash Join. |
| Whether to enable the asynchronous DDL engine | Database | Specifies whether to enable the data definition language (DDL) engine. If you disable it, the execution logic of the original DDL engine remains. |

| Parameter | Level | Description |
|---|----------|---|
| Whether to enable asynchronous-only mode under asynchronous DDL engine | Database | Specifies whether to enable the asynchronous-only mode when the asynchronous DDL engine is enabled. <ul style="list-style-type: none"> Enabled: The status is returned immediately after the client connects to PolarDB-X and executes the DDL statement. The execution status can be viewed only through asynchronous DDL management statements. Disable: The synchronous mode remains. That is, the status is returned only after the client completes executing the DDL statement. |
| Maximum number of physical tables allowed to be created in a single physical database | Database | The maximum number of table shards that can be created in a database shard. |
| INFORMATION_SCHEMA.TABLES queries whether statistics are aggregated | Database | Specifies whether to aggregate statistics of INFORMATION_SCHEMA.TABLES queries. To ensure the performance, it is not aggregated by default. |
| Maximum number of physical sharding links | Database | The maximum number of connections between PolarDB-X and a single ApsaraDB RDS for MySQL shard. |
| Minimum number of physical sharding links | Database | The minimum number of connections between PolarDB-X and a single ApsaraDB RDS for MySQL shard. |
| Physical idle link timeout | Database | The idle time of the connection between PolarDB-X and ApsaraDB RDS for MySQL (unit: minute). |

11.4.8. SQL audit and analysis

11.4.8.1. Description

Cloud Native Distributed Database PolarDB-X (PolarDB-X) combines the SQL audit and analysis feature with Log Service (SLS). This feature not only audits historical SQL records, but also provides real-time diagnosis and analysis of SQL execution status, performance metrics, and security risks. You can enable SQL audit and analysis in the PolarDB-X console.

Benefits

- **Easy operation:** SQL audit and analysis can be enabled with easy configuration to help you audit and analyze SQL logs in real time.
- **Lossless performance:** Pulling SQL log files from PolarDB-X nodes and uploading these logs to SLS in real time does not affect instance performance.
- **Trace to historical issues:** This feature supports importing historical SQL logs to trace issues.
- **Real-time analysis:** This feature provides real-time SQL analysis and an out-of-the-box report center based on SLS. This feature also supports custom reports and drill-down analysis, and helps you understand the execution status, performance, and security risks of databases.
- **Real-time alerts:** This feature supports real-time monitoring and alerts based on customized metrics to ensure timely response to critical business exceptions.

Limits and instructions

- You must activate Alibaba Cloud SLS to use the SQL audit and analysis feature.
- SQL audit logs are saved for 30 days by default. You can modify the log storage time as needed.
- Do not delete or modify the default settings for the project, Logstore, index, or dashboard that are created by SLS. SLS updates and upgrades the SQL log audit feature from time to time. The indexes and default reports of the exclusive Logstore are also automatically updated.
- The maximum length of a single SQL statement is 5 MB.

Scenarios

- Troubleshoot SQL problems

After the SQL audit and analysis feature is enabled, you can quickly search SQL logs to locate and troubleshoot problems. For example, to check whether a DROP operation is performed, you can perform the following query:

```
sql_type: Drop
```

The query result contains information such as the SQL execution time, user, and IP address of the client that runs the SQL statement.

- Analyze costly SQL templates

In most applications, SQL statements are dynamically generated based on several templates, with different parameters. The real-time analysis feature of SLS allows you to obtain the list of costly SQL statements in the current database.

For example, execute the following query:

```
| SELECT sql_code as "SQL template ID",  
round(total_time * 1.0 / sum(total_time) over() * 100,2) as "execution time share (%)",  
execute_times as "number of execution times",  
round(avg_time) as "average execution time",  
round(avg_rows) as "average number of affected rows",  
CASE WHEN length(sql) > 200 THEN concat(substr(sql, 1, 200), '.....') ELSE trim(lpad(sql, 200,$) end as "sample SQL" FR  
OM (SELECT sql_code, count(1) as execute_times,  
sum(response_time) as total_time,  
avg(response_time) as avg_time,  
avg(affect_rows) as avg_rows,  
arbitrary(sql) as sql FROM log GROUP BY sql_code) ORDER BY "execution time share (%)" desc limit 10
```

The search result contains the SQL template ID, ratio of response time of the statement generated from the template in the total response time of SQL statements, number of executions, average execution time, average number of affected rows, and sample SQL statement. You can find and optimize the most costly SQL templates in the application based on the analysis result.

- Collect log statistics

To help you analyze issues, PolarDB-X combines the SQL audit and analysis feature with SLS and provides out-of-the-box reports. You can diagnose and analyze the running status, performance, and potential security risks of databases in real time.

11.4.8.2. Enable SQL audit and analysis

The SQL audit and analysis feature is disabled by default. You can manually enable it in the PolarDB-X console. By default, you can perform only audit and analysis on the log data generated after the SQL audit and analysis feature is enabled. You can also import a portion of historical data.

Prerequisites

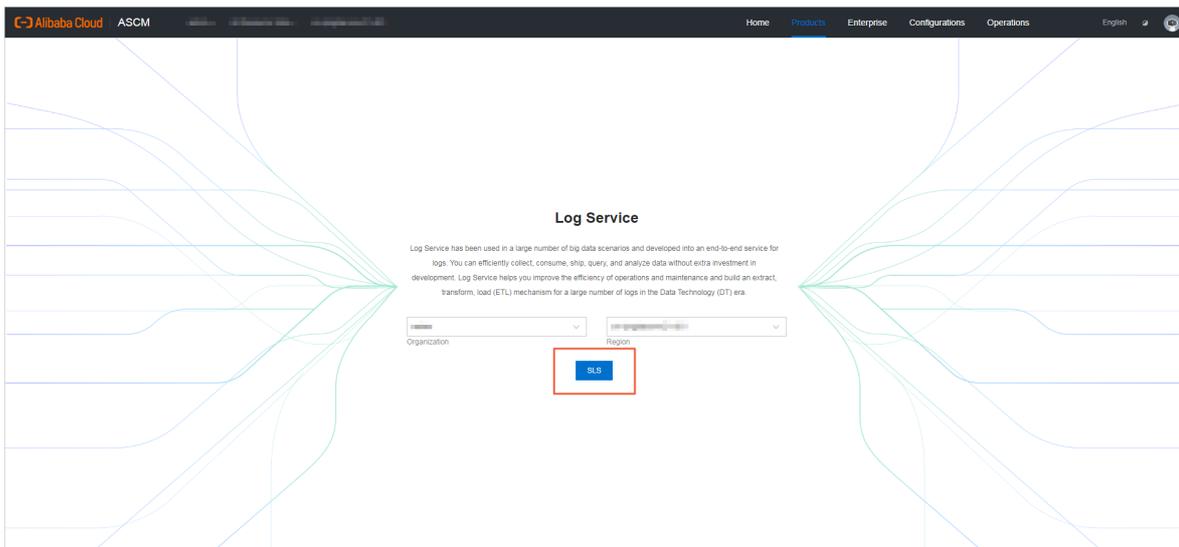
The SQL audit and analysis feature depends on Log Service (SLS). You must activate SLS before you use this feature.

Procedure

1. Log on to the SLS console. For more information, see *Log Service User Guide > Quick Start > Log on to the Log Service console*.
2. Select the organization to which the PolarDB-X instance belongs.

Note The logon account must be consistent with the logon account of PolarDB-X.

3. Click SLS to go to the Log Service page.



4. Log on to the PolarDB-X console.
5. Find the target instance in the instance list.
6. Click the target instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
7. In the left-side navigation pane, choose **Diagnostics and Optimization > SQL Audit and Analysis**.
8. In the left-side section, select the database for which you want to enable the SQL audit and analysis feature.
9. On the SQL Audit and Analysis page, turn on the switch next to **SQL Audit Log Status of Current Database** on the right.

Note On the SQL Audit and Analysis page, you can also turn on the switch of **SQL Audit and Analysis** next to the target database in the left-side section.

10. Confirm whether to import historical data.

Note By default, you can analyze and audit only the logs that are generated after the SQL audit and analysis feature is enabled. If you find the historical data of the PolarDB-X database is modified but the SQL audit and analysis feature is not enabled, you can import historical data and include historical logs in the audit and analysis scope to trace data tampering. PolarDB-X dynamically checks the scope of historical data that can be imported based on the log storage on the PolarDB-X instance. Logs within seven days can be imported.

- o If you need to import historical data, enable **Import Historical Data or Not**, specify the backtrace start time and end time, and then click **Enable**.
- o If you do not need to import historical data, click **Enable**.

What's next

Every time you use the SQL audit and analysis feature, you must repeat the preceding steps.

11.4.8.3. Log fields

This topic describes the log fields in SQL audit and analysis.

| Field | Description | Supported version |
|-----------------|---|-------------------|
| __topic__ | The log topic in the format of <code>drds_audit_log_{instance_id name {db name}}</code> , such as <code>drds_audit_log_drdsxyzabcd_demo_drds_db</code> . | All versions |
| instance_id | The ID of the PolarDB-X instance. | All versions |
| db_name | The name of the PolarDB-X database. | All versions |
| user | The user name used to run the SQL statement. | All versions |
| client_ip | The IP address of the client that accessed the PolarDB-X instance. | All versions |
| client_port | The port of the client that accessed the PolarDB-X instance. | All versions |
| sql | The executed SQL statement. | All versions |
| trace_id | The trace ID of the SQL statement when it was executed. If a transaction was executed, it is tracked by an ID that consists of the trace ID, a hyphen, and a number, for example, <code>drdsabcdxyz-1</code> and <code>drdsabcdxyz-2</code> . | All versions |
| sql_code | The hash value of the template SQL statement. | All versions |
| hint | The hint that was used to execute the SQL statement. | All versions |
| table_name | The name of the table involved in the query. Separate multiple tables by commas (,). | All versions |
| sql_type | The type of the SQL statement. Valid values: SELECT, INSERT, UPDATE, DELETE, SET, ALTER, CREATE, DROP, TRUNCATE, REPLACE, and Other. | All versions |
| sql_type_detail | The name of the SQL parser. | All versions |
| sql_time | The start time for the execution of the SQL statement. The time follows the <code>yyyy-MM-dd HH:mm:ss.SSS</code> format. | All versions |

| Field | Description | Supported version |
|---------------|--|----------------------------------|
| response_time | The response time. Unit: milliseconds. | Version 5.3.4-15378085 and later |
| affect_rows | The number of rows returned when the SQL statement was executed. The number of rows affected when the INSERT, DELETE, or UPDATE statement was executed. | Version 5.3.4-15378085 and later |
| fail | Indicates whether an error occurred in the execution of the SQL statement. Valid values: <ul style="list-style-type: none"> 0: successful 1: failed | Version 5.3.4-15378085 and later |

11.4.8.4. Log analysis

The SQL audit and analysis feature is based on Log Service (SLS) and provides powerful log analytics capabilities. This topic describes SQL statements for log analysis in common scenarios and provides relevant examples.

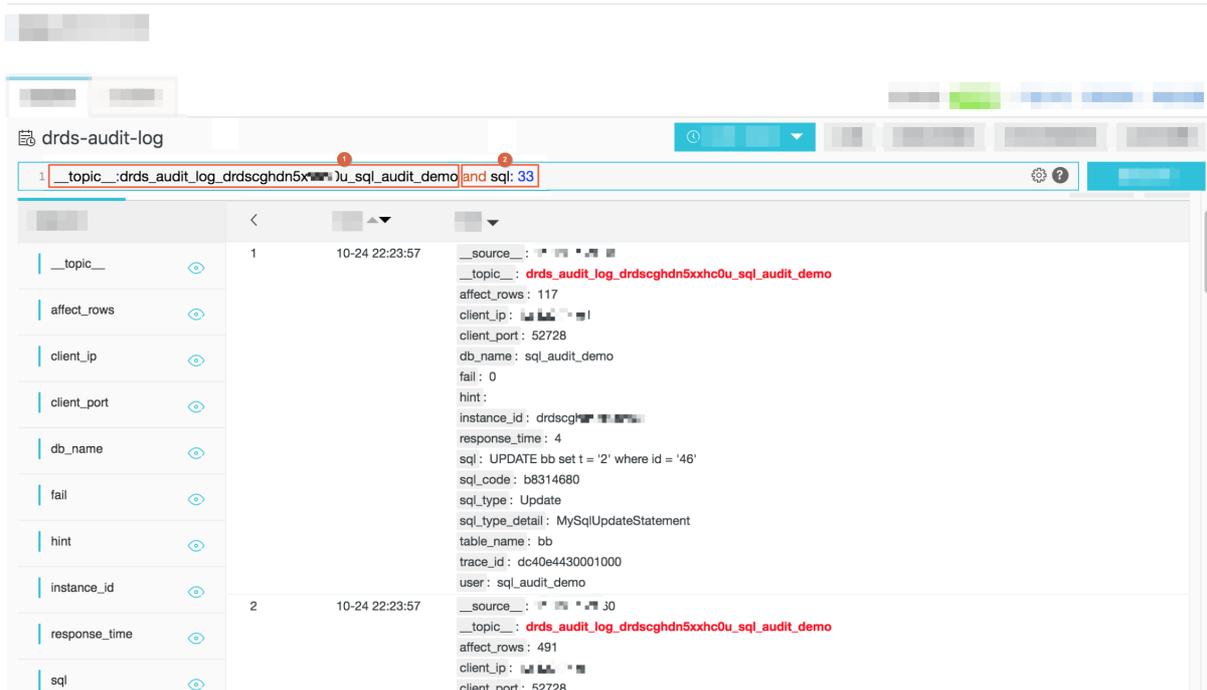
After the SQL audit and analysis feature is enabled, you can perform audit and analysis on SQL log files by using the query and analysis syntax of SLS on the SQL Audit and Analysis page. Based on the query and analysis syntax of SLS, you can find problematic SQL statements on the Log Analysis tab and analyze the SQL statement execution status, performance metrics, and security issues of PolarDB-X. For more information about the query and analysis syntax of SLS, see *Log Service User Guide > Query and Analysis > Query Syntax and Functions > Query Syntax*.

Precautions

All the audit logs of PolarDB-X databases in the same region are written to the same Logstore in SLS. Therefore, by default, the SQL Audit and Analysis page provides the filter conditions based on `__topic__`, to ensure that the searched SQL log files are from PolarDB-X. Therefore, all the statements provided in this topic must be used after the existing filter conditions.

An example is shown in the following figure:

- The ① part is the default filter condition.
- The ② part is the additional filter condition.



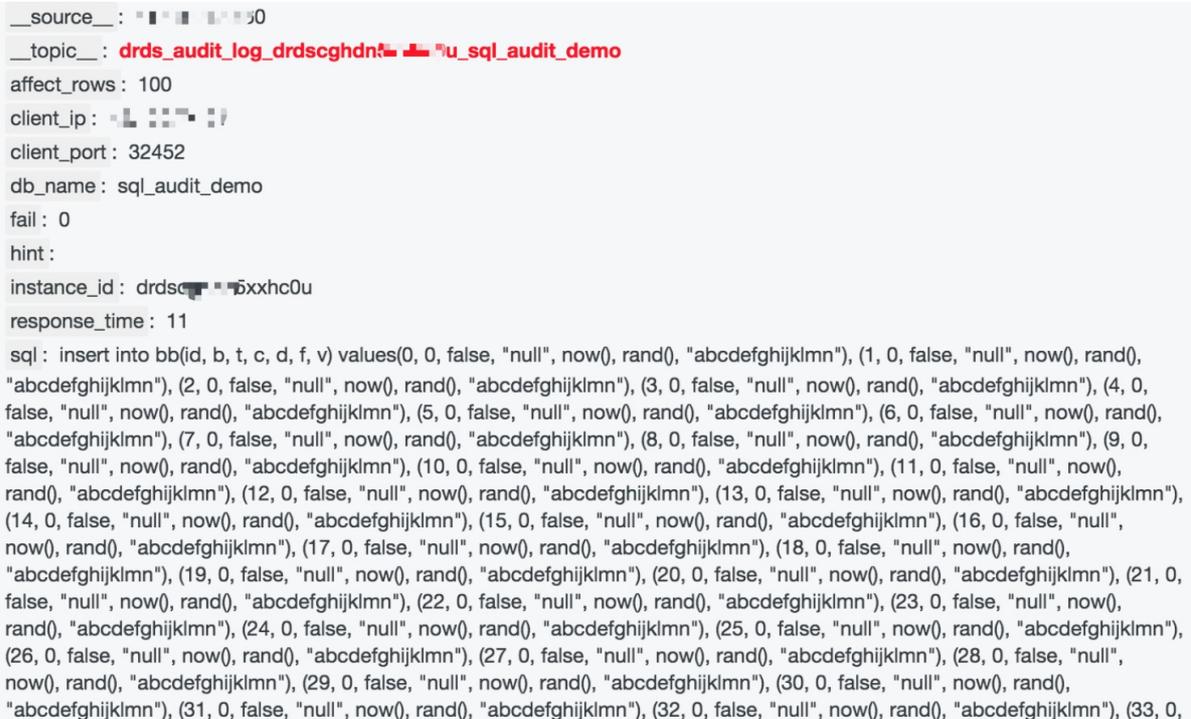
Find problematic SQL statements

- Fuzzy search

For example, to query SQL statements that contain the "34" keyword, enter the following content in the search box:

```
and sql: 34
```

The result is shown in the following figure.



- Field search

Based on built-in index fields, the SQL audit and analysis feature also supports field-based search.

For example, to query SQL statements of the Drop type, execute the following statement:

```
and sql_type:Drop
```

The result is shown in the following figure.

```
__source__: ...
__topic__: drds_audit_log_drdschgdr...0u_sql_audit_demo
affect_rows: 0
client_ip: ...
client_port: 36085
db_name: sql_audit_demo
fail: 0
hint:
instance_id: drdschgdr...0u
response_time: 3172
sql: drop table if exists bb
sql_code: 0cfc96e8
sql_type: Drop
sql_type_detail: SQLDropTableStatement
table_name: bb
trace_id: dc408feedc00000
user: sql_audit_demo
```

- Multi-condition search

You can use the "and" and "or" keywords to perform a multi-condition search.

For example, you can query the delete operation on the rows whose id is 34:

```
and sql:34 and sql_type: Delete
```

- Search based on numeric comparison

affect_rows and response_time in the index field are numeric values and support comparison operators.

For example, you can query the SQL INSERT statements whose response_time is greater than 1s.

```
and response_time > 1507 and sql_type: Insert
```

For example, you can query the SQL statement that deletes more than 100 rows of data:

```
and affect_rows > 100 and sql_type: Delete
```

Analysis of the SQL statement execution status

This section introduces the statements used to query the SQL statement execution status in PolarDB-X.

- Failure rate of SQL statement execution

Execute the following statement to query the failure rate of SQL statement execution:

```
| SELECT sum(case when fail = 1 then 1 else 0 end) * 1.0 / count(1) as fail_ratio
```

The result is shown in the following figure.



0.0010322901477612633

If your business is sensitive to the error rate of SQL statement execution, you can customize the alert information based on the query result. Click **Save as Alert** in the upper-right corner of the page.

In the alert settings shown in the preceding figure, the number of log entries that have an error rate of SQL statement execution greater than 0.01 within 15 minutes is checked within every 15 minutes. You can also customize alerts as needed.

- Total number of rows returned by SELECT statements

Execute the following statement to query the cumulative number of rows queried by SELECT statements:

```
and sql_type: Select | SELECT sum(affect_rows)
```

- SQL statement type distribution

Execute the following statement to query the SQL statement type distribution:

```
| SELECT sql_type, count(sql) as times GROUP BY sql_type
```

- IP address distribution of SQL independent users

Execute the following statement to query the distribution of IP addresses of independent users who execute SQL statements:

```
| SELECT user, client_ip, count(sql) as times GROUP BY user, client_ip
```

SQL performance analysis

This section describes typical SQL statements for SQL performance analysis.

- Average response time of SELECT statements

Execute the following statement to query the average response time of SELECT statements:

```
and sql_type: Select | SELECT avg(response_time)
```

- Distribution of SQL statement response time

Execute the following statement to query the distribution of SQL statement response time:

```
and response_time > 0 | select case when response_time <=10 then '<=10 ms' when response_time > 10 and response_time <= 100 then '10~100 ms' when response_time > 100 and response_time <= 1000 then '100 ms ~ 1s' when response_time > 1000 and response_time <= 10000 then '1s ~ 10s' when response_time > 10000 and response_time <= 60000 then '10s ~ 1 min' > 1 min' end as latency_type, count(1) as cnt group by latency_type order by latency_type DESC
```

The preceding query shows the distribution of SQL statement execution time based on a given time range. You can adjust the time range to obtain finer-grained results.

- Top 50 slow SQL statements

Execute the following statement to query slow SQL statements:

```
| SELECT date_format(from_unixtime(__time__), '%m/%d %H:%i:%s') as time, user, client_ip, client_port, sql_type, affect_rows, response_time, sql ORDER BY response_time desc LIMIT 50
```

The following figure shows the result, which includes the SQL statement execution time, user name, IP address, port number, SQL statement type, number of affected rows, response time, and text of SQL statements.

| time | user | client_ip | client_port | sql_type | affect_rows | response_time | sql |
|----------------|----------------|---------------|-------------|----------|-------------|---------------|-------------------------|
| 09/28 14:04:05 | sql_audit_demo | 192.168.1.101 | 77 | Drop | 0 | 9583 | drop table if exists bb |
| 09/28 14:04:05 | sql_audit_demo | 192.168.1.101 | 77 | Drop | 0 | 9583 | drop table if exists bb |
| 09/28 14:04:05 | sql_audit_demo | 192.168.1.101 | 77 | Drop | 0 | 9583 | drop table if exists bb |
| 09/27 17:38:18 | sql_audit_demo | 192.168.1.101 | 73 | Drop | 0 | 7200 | drop table if exists bb |

- Top 10 costly SQL templates

In most applications, SQL statements are dynamically generated based on several templates, and only the parameters are different. You can find, analyze, and optimize the costly SQL templates based on template IDs. Enter the following query statement:

```
| SELECT sql_code as "SQL template ID", round(total_time * 1.0 / sum(total_time) over() * 100, 2) as "response time share (%)", execute_times as "number of executions", round(avg_time) as "average response time", round(avg_rows) as "average number of affected rows", CASE WHEN length(sql) > 200 THEN concat(substr(sql, 1, 200), '.....') ELSE trim(lpad(sql, 200, 'hour') end as "sample SQL" FROM (SELECT sql_code, count(1) as execute_times, sum(response_time) as total_time, avg(response_time) as avg_time, avg(affect_rows) as avg_rows, arbitrary(sql) as sql FROM log GROUP BY sql_code) ORDER BY "execution time share (%)" desc limit 10
```

The statistics include the SQL template ID, percentage of response time of the statement generated from the template in the total response time of SQL statements, number of executions, average response time, average number of affected rows, and sample SQL statement. For better display effect, each page displays 200 entries. In the preceding query result, statements are ranked by the response time share. However, you can rank the statements by the average response time or the number of executions to troubleshoot relevant issues.

- Average transaction response time

For SQL statements within the same transaction, the preset trace_id field prefixes are the same, and the suffixes are '-' followed by sequence numbers. trace_id of non-transactional SQL statements does not contain '-'. Based on this, you can analyze the performance of transactions.

 **Note** Transaction analysis is less efficient than other query operations because it involves prefix matching.

For example, execute the following statement to query the average response time of transactions:

```
| SELECT sum(response_time) / COUNT(DISTINCT substr(trace_id, 1, strpos(trace_id, '-') - 1)) where strpos(trace_id, '-') > 0
```

- Top 10 slow transactions

You can query the list of slow transactions by response time of transactions. Use the following statement:

```
| SELECT substr(trace_id, 1, strpos(trace_id, '-') - 1) as "transaction ID", sum(response_time) as "response time" where strpos(trace_id, '-') > 0 GROUP BY substr(trace_id, 1, strpos(trace_id, '-') - 1) ORDER BY "response time" DESC LIMIT 10
```

Based on this, you can use the transaction ID to search for all the SQL statements under the transaction and analyze the specific causes of slow execution. Use the following statement:

```
and trace_id: db3226a20402000*
```

- Top 10 transactions with batch operations

Based on the number of rows affected by SQL statements in a transaction, you can obtain the list of transactions that contain batch operations. Use the following statement:

```
| SELECT substr(trace_id, 1, strpos(trace_id, '-') - 1) as "transaction ID", sum(affect_rows) as "number of affected rows" where strpos(trace_id, '-') > 0 GROUP BY substr(trace_id, 1, strpos(trace_id, '-') - 1) ORDER BY "number of affected rows" DESC LIMIT 10
```

SQL security analysis

This section provides typical query statements for SQL security analysis.

- Distribution of types of failed SQL statements

```
and fail > 0 | select sql_type, count(1) as "number of errors" group by sql_type
```

- High-risk SQL statements

High-risk SQL statements are of the Drop or Truncate type. You can also add more conditions as needed.

```
and sql_type: Drop OR sql_type: Truncate
```

- SQL batch delete events

```
and affect_rows > 100 and sql_type: Delete | SELECT date_format(from_unixtime(__time__), '%m/%d %H:%i:%s') as time, user, client_ip, client_port, affect_rows, sql ORDER BY affect_rows desc LIMIT 50
```

11.4.8.5. Log reports

Based on Log Service (SLS), the SQL audit and analysis feature of PolarDB-X provides out-of-the-box report centers, including the Operation Center, Performance Center, and Security Center. This feature allows you to fully understand the performance status, performance metrics, and potential security risks of your PolarDB-X databases.

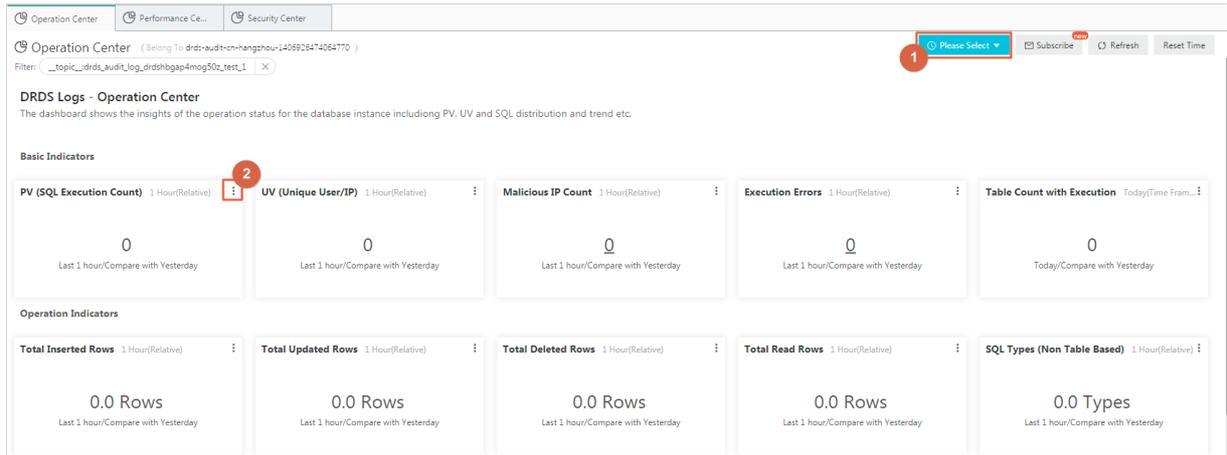
After [Enable SQL audit and analysis](#), click the **Log Reports** tab on the current page. You can view the reports pages provided by SLS, including **Operation Center**, **Performance Center**, and **Security Center**.

Note

- All the audit logs of PolarDB-X databases in the same region are written to the same Logstore in SLS. Therefore, when you view the reports of the current PolarDB-X database, filter conditions based on the `__topic__:drds_audit_log_instance_id_database name` are added by default, which indicates that you are viewing the data of the current database. For example, `drds_audit_log_drdsxyzabcd_demo_drds_db`.
- If the version of the PolarDB-X instance is earlier than Version 5.3.4-15378085, the relevant fields are missing from SQL logs. For more information about log fields, see [Log fields](#). The Log Reports tab provides only a simplified version of Operation Center. To use a full version of reports, upgrade the instance to the latest version.

View reports

Statistics in the charts on the Log Reports tab are generated for different time periods. You can change the time range as needed. You can change the time range for all charts or a single chart.



- Click the time selector (position ① in the figure). In the dialog box that appears, you can change the time range of all the charts on the current page.
- Click the time selector of a chart (position ② in the figure) to modify the time range of the chart.

Operation Center

Operation Center shows the metrics, distribution, and trends of SQL statement execution in PolarDB-X databases.

| Item | Type | Default time range | Description |
|----------------------------|--------------|--------------------|--|
| PV (SQL Execution Count) | Single value | 1 Hour (Relative) | The number of SQL statement executions |
| UV (Unique User/IP) | Single value | 1 Hour (Relative) | The number of unique user-IP groups |
| Malicious IP Count | Single value | 1 Hour (Relative) | The number of malicious IP addresses. For more information about the definition of malicious IP addresses, see security detection functions. |
| Execution Errors | Single value | 1 Hour (Relative) | The number of SQL statements with execution errors |
| Table Count with Execution | Single value | 1 Hour (Relative) | The total number of tables operated by SQL statements |
| Total Inserted Rows | Single value | 1 Hour (Relative) | The total number of rows inserted by INSERT statements |
| Total Updated Rows | Single value | 1 Hour (Relative) | The total number of rows updated by UPDATE statements |
| Total Deleted Rows | Single value | 1 Hour (Relative) | The total number of rows deleted by DELETE statements |

| Item | Type | Default time range | Description |
|--------------------------------------|--------------|--------------------|---|
| Total Read Rows | Single value | 1 Hour (Relative) | The total number of rows returned by SELECT statements |
| SQL Types (Non Table Based) | Single value | 1 Hour (Relative) | The types of SQL statements used for non-table operations, such as SHOW VARIABLES LIKE |
| SQL Execution Trend | Column chart | 1 Hour (Relative) | The distribution trend of SQL statement executions and the distribution trend of failed SQL statements |
| Operated Tables | Flow diagram | 1 Hour (Relative) | The distribution of tables operated by SQL statements |
| SQL Type | Flow diagram | 1 Hour (Relative) | The distribution of SQL statement types by time |
| User Distribution | Pie chart | 1 Hour (Relative) | The distribution of users who execute SQL statements |
| SQL Type Distribution | Area chart | 1 Hour (Relative) | The percentage of SQL statement types in the current time range |
| Tables with Most Operations (Top 50) | Table | 1 Hour (Relative) | The list of top tables by the number of operations, including table names and the number of operations such as read, delete, update, and insert |
| SQL Type (World) | Map | 1 Hour (Relative) | The distribution of IP addresses of clients that execute the SQL statements, on the world map |
| SQL Type (China) | Map | 1 Hour (Relative) | The distribution of IP addresses of clients that execute the SQL statements, on the map of China |

Performance Center

Performance Center shows performance metrics, the distribution of slow and fast SQL statements, and the distribution and sources of costly SQL statements in PolarDB-X databases.

| Item | Data type | Default time range | Description |
|----------------------------|--------------|--------------------|--|
| Peak SQL Execution Traffic | Single value | 1 Hour (Relative) | The maximum number of SQL statements executed per second |

| Item | Data type | Default time range | Description |
|---------------------------------|--------------|--------------------|--|
| Peak Select Traffic | Single value | 1 Hour (Relative) | The maximum number of rows returned by SELECT statements per second |
| Peak Insert Traffic | Single value | 1 Hour (Relative) | The maximum number of rows inserted by INSERT statements per second |
| Peak Update Traffic | Single value | 1 Hour (Relative) | The maximum number of rows updated by UPDATE statements per second |
| Peak Delete Traffic | Single value | 1 Hour (Relative) | The maximum number of rows deleted by DELETE statements per second |
| Average Response Time | Single value | 1 Hour (Relative) | The average response time of SQL statements |
| Select SQL | Single value | 1 Hour (Relative) | The average number of SELECT statements executed per second |
| Insert SQL | Single value | 1 Hour (Relative) | The average number of INSERT statements executed per second |
| Update SQL | Single value | 1 Hour (Relative) | The average number of UPDATE statements executed per second |
| Delete SQL | Single value | 1 Hour (Relative) | The average number of DELETE statements executed per second |
| Select/Update Traffic Trend | Line chart | 1 Hour (Relative) | The distribution of rows affected by the SELECT and UPDATE statements over time |
| SQL Execution Time Distribution | Pie chart | 1 Hour (Relative) | The distribution of execution time of SQL statements |
| Slow SQL Table Distribution | Pie chart | 1 Hour (Relative) | The distribution of tables targeted by slow SQL statements whose response time exceeds 1s |
| Slow SQL User Distribution | Pie chart | 1 Hour (Relative) | The distribution of users who execute slow SQL statements with response time of more than 1s |
| Slow SQL Type Distribution | Pie chart | 1 Hour (Relative) | The distribution of types of slow SQL statements whose response time exceeds 1s |

| Item | Data type | Default time range | Description |
|------------------------------------|-----------|--------------------|--|
| Slow SQL (Top 50) | Table | 1 Hour (Relative) | The table of slow SQL statements whose response time exceeds 1s, including the time, client, response time, PolarDB-X instance, database, table, user, affected rows, SQL type, and SQL text |
| SQL Template Execution Time Top 20 | Table | 1 Hour (Relative) | Statistics of the execution status of the SQL statements in the template based on the specified SQL template, including the SQL template ID, response time share, number of executions, average response time, average number of affected rows, and sample SQL statement |
| Transaction Affected Rows Top 20 | Table | 1 Hour (Relative) | The table of top 20 transaction-by the number of affected rows, including the transaction ID and the number of affected rows |
| Transaction Executed Time Top 20 | Table | 1 Hour (Relative) | The table of top 20 transactions by response time, including the transaction ID and the number of affected rows |

Security Center

Security Center shows failed and malicious SQL statement executions in PolarDB-X databases, and the details, distribution, and trends of malicious SQL batch delete and update events.

| Item | Type | Default time range | Description |
|--------------------------|--------------|--------------------|---|
| Error Count | Single value | 1 Hour (Relative) | The number of failed SQL statement executions |
| Batch Delete Events | Single value | 1 Hour (Relative) | The number of SQL statements for batch delete events (more than 100 rows) |
| Batch Update Events | Single value | 1 Hour (Relative) | The number of SQL statements for batch update events (more than 100 rows) |
| Malicious SQL Executions | Single value | 1 Hour (Relative) | The number of malicious SQL statement executions (Drop and Truncate) |

| Item | Type | Default time range | Description |
|------------------------------------|--------------|--------------------|--|
| Malicious IP Count | Single value | 1 Hour (Relative) | The number of malicious IP addresses. For more information about the definition of malicious IP addresses, see security detection functions. |
| Error Distribution | Area chart | 1 Hour (Relative) | The distribution of types of failed SQL statements |
| Distribution of Client with Errors | Map | 1 Hour (Relative) | The distribution of clients for failed SQL statements on the map of China |
| Client with Most Errors | Table | 1 Hour (Relative) | The table of clients on which the execution of SQL statements failed, including the IP address, number of errors, type of failed SQL statement, and sample failed SQL statement |
| Malicious SQL Executions | Table | 1 Hour (Relative) | The table of malicious SQL statement executions, including the time, IP address, SQL, PolarDB-X instance ID, database, table, and user |
| Batch Delete Events (Top 50) | Table | 1 Hour (Relative) | The table of top SQL batch delete events, including the earliest execution time, most recent execution time, PolarDB-X instance ID, database, table, number of executions, average number of deleted rows, average response time, and sample SQL statement |
| Batch Update Events (Top 50) | Table | 1 Hour (Relative) | The table of top SQL batch update events, including the earliest execution time, most recent execution time, PolarDB-X instance ID, database, table, number of executions, average number of updated rows, average response time, and sample SQL statement |

11.4.9. Monitor PolarDB-X instances

11.4.9.1. View monitoring information

PolarDB-X provides multi-dimensional monitoring. This topic describes how to view monitoring information in the PolarDB-X console.

Procedure

1. [Log on to the PolarDB-X console.](#)
2. On the **DRDS Instance Management** page, find the target instance.
3. Click the instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. On the **Basic Information** page, choose **Monitoring and Alerts > Instance Monitoring** from the left-side navigation pane.
5. On the **Instance Monitoring** page, select a monitoring dimension and the corresponding metrics to view details. For more information about monitoring metrics, see [Monitoring metrics](#).

11.4.9.2. Monitoring metrics

Instance monitoring is divided into resource monitoring and engine monitoring. Engine monitoring metrics are classified into metrics at the PolarDB-X instance level and metrics at the PolarDB-X database level. When some engine monitoring metrics are abnormal, you can directly check the metrics of each database to locate the database with performance problems. The following table describes the metrics of these two types in details.

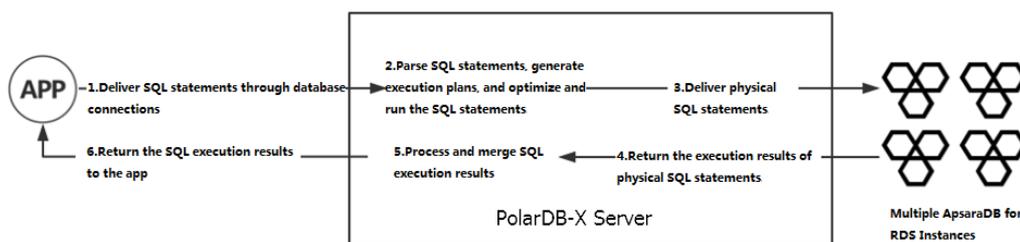
| Monitoring item | Category | Description | Data collection cycle | Data retention period | Description |
|-------------------------|---------------------|---|-----------------------|-----------------------|--|
| CPU Utilization (%) | Resource monitoring | The average CPU utilization of PolarDB-X server nodes. | 1 minute | 3 days | - |
| Memory Usage (%) | Resource monitoring | The memory usage of JVM Old Generation on PolarDB-X server nodes. | 1 minute | 3 days | Memory usage fluctuations are normal. |
| Inbound Traffic (Kbps) | Resource monitoring | The total inbound network traffic of PolarDB-X server nodes. | 1 minute | 3 days | Inbound network traffic is generated when ApsaraDB RDS for MySQL returns data to PolarDB-X. |
| Outbound Traffic (Kbps) | Resource monitoring | The total outbound network traffic of PolarDB-X server nodes. | 1 minute | 3 days | Outbound network traffic is generated when a PolarDB-X instance sends a physical SQL statement to an ApsaraDB RDS for MySQL instance or a PolarDB-X instance returns data to an application. |

| Monitoring item | Category | Description | Data collection cycle | Data retention period | Description |
|------------------|-------------------|---|-----------------------|-----------------------|--|
| Logical QPS | Engine monitoring | The total number of SQL statements processed per second on PolarDB-X server nodes. | 5 seconds | 7 days | - |
| Physical QPS | Engine monitoring | The total number of SQL operations sent from PolarDB-X server nodes to ApsaraDB RDS for MySQL per second. | 5 seconds | 7 days | One logical SQL statement can be partitioned into multiple physical SQL statements. |
| Logical RT (ms) | Engine monitoring | The average response time (RT) for processing each SQL statement by PolarDB-X. | 5 seconds | 7 days | If a logical SQL statement is partitioned into physical SQL statements for delivery, the logical RT of the SQL statement contains the RT of the physical SQL statements. |
| Physical RT (ms) | Engine monitoring | The average RT for transmitting SQL statements from PolarDB-X to ApsaraDB RDS for MySQL. | 5 seconds | 7 days | - |
| Connections | Engine monitoring | The total number of connections established between an application and PolarDB-X. | 5 seconds | 7 days | The connections from PolarDB-X to ApsaraDB RDS for MySQL are not included. |
| Active Threads | Engine monitoring | The number of threads that are used by PolarDB-X to run SQL statements. | 5 seconds | 7 days | - |

11.4.9.3. How metrics work

Before analyzing metrics, you need to understand the execution process of SQL statements on PolarDB-X.

PolarDB-X SQL execution flowchart



In the entire SQL execution process, the execution status of steps 2 through 4 is reflected in various metrics of PolarDB-X.

- In step 2, SQL parsing, optimization, and execution consume CPU resources. A more complex SQL statement (with a complex structure or ultra-long length) consumes more CPU resources. You can run the `TRACE` command to trace the SQL execution process. You can see the time consumed by an SQL statement during optimization. The longer time consumed indicates a higher CPU utilization.
- In step 3, the delivery and execution of physical SQL statements consume I/O resources. You can analyze the execution status of physical SQL statements based on metrics such as logical queries per second (QPS), physical QPS, logical response time (RT), and physical RT. For example, if the physical QPS is low and the physical RT is high, the current ApsaraDB RDS for MySQL instance is processing SQL statements very slowly. You need to check the performance of the ApsaraDB RDS for MySQL instance.
- In step 5, the SQL execution results are processed and integrated. These operations convert the execution results of physical SQL statements. In most cases, only SQL metadata is converted, which consumes few resources. However, the CPU utilization is high for steps such as `heap sort`. For more information about how to determine the consumption of SQL statements at this stage, see [Details about a low SQL statement](#).

11.4.9.4. Prevent performance problems

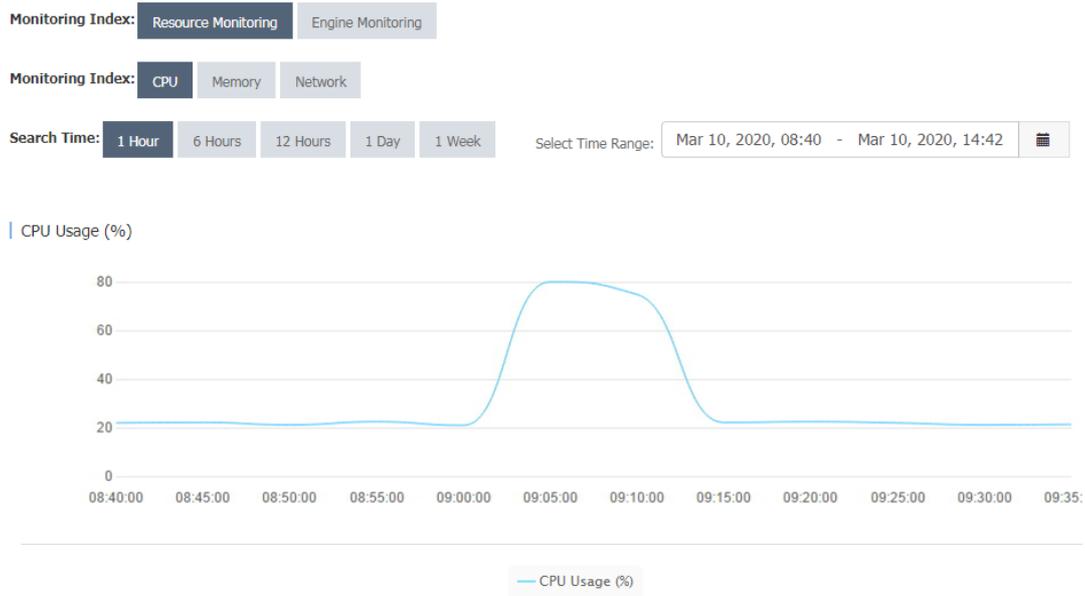
11.4.9.4.1. Example 1: PolarDB-X CPU utilization

Performance metrics change with the system business traffic.

The following describes the CPU utilization in two common cases:

- An application has a shopping spree activity at 09:00 every morning. Therefore, the traffic of the system increases significantly at this time point. According to the monitoring data, the CPU utilization of the PolarDB-X instance increased from 20% to about 80% from about 09:00, with the traffic peak lasting about 10 minutes.

CPU utilization-1



- The system traffic keeps increasing with an application until it reaches a plateau. The monitored CPU utilization of the PolarDB-X instance also reflects this change.

CPU utilization-2



When the load on the PolarDB-X instance changes with the business, you must pay close attention to the changes in metrics. If the CPU utilization exceeds the threshold, you must upgrade the PolarDB-X specifications to alleviate the performance pressure.

You can set alert rules for instances in the PolarDB-X console. When the average CPU utilization exceeds the preset threshold, the system sends short messages to the corresponding contacts. You can set the CPU utilization threshold as needed. We recommend that you set it to 80%.

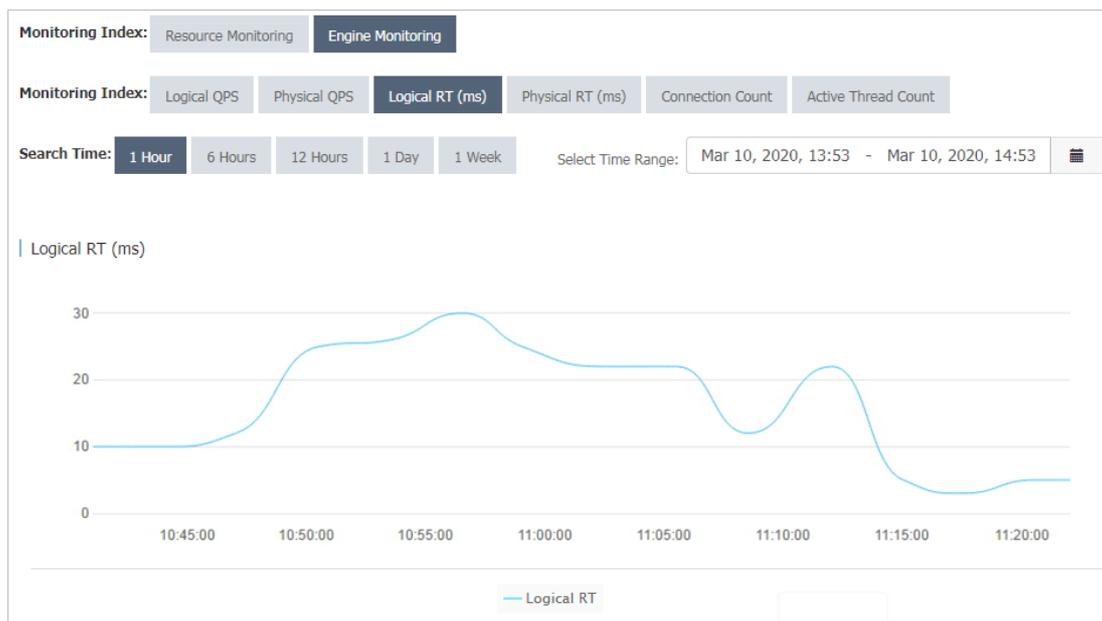
11.4.9.4.2. Example 2: Logical RT and physical RT

You can observe the difference between the logical response time (RT) and physical RT.

Logical RT refers to the time from when a PolarDB-X instance receives a logical SQL statement to when it returns data to an application. Physical RT refers to the time from when a PolarDB-X instance sends a physical SQL statement to an ApsaraDB RDS for MySQL instance to when it receives the data returned by the ApsaraDB RDS for MySQL instance.

If a logical SQL statement is partitioned into one or more physical SQL statements, the logical RT is greater than or equal to the physical RT. Ideally, PolarDB-X performs only a few operations on the data returned by ApsaraDB RDS for MySQL. Therefore, logical RT is slightly longer than physical RT. Under special circumstances, physical SQL queries are run fast, while logical SQL queries take a long time to run. In this case, the logical RT and physical RT are as follows.

Logical RT



Physical RT



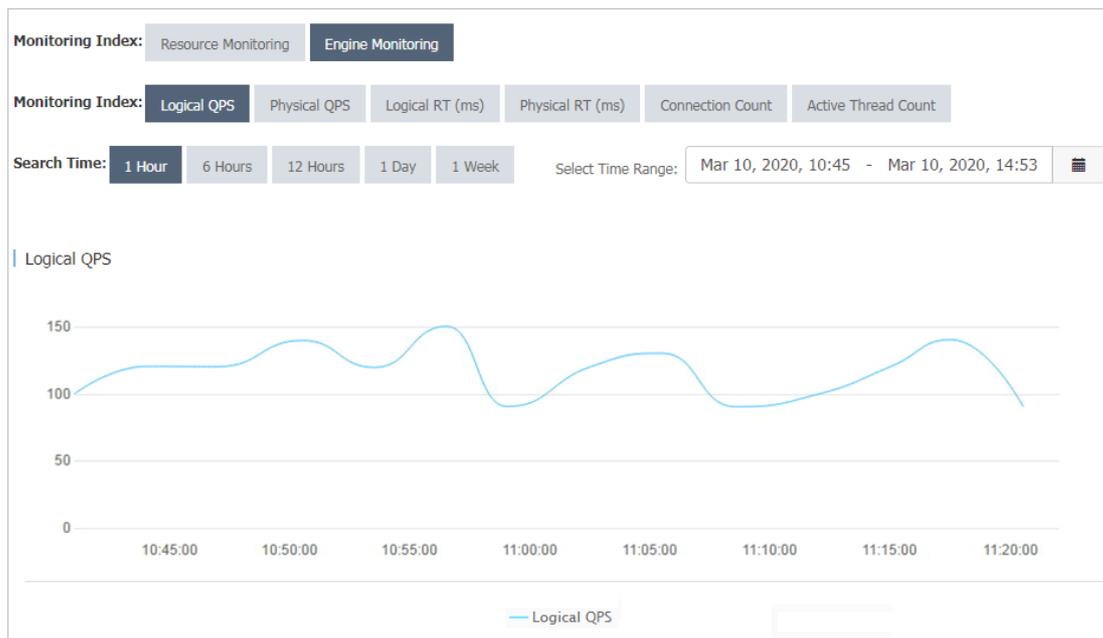
As shown in the preceding figures, the change trends of logical RT and physical RT in the two monitoring charts are basically the same, while logical RT fluctuates between 10 ms and 20 ms and physical RT fluctuates between 2 ms and 5 ms. This means that PolarDB-X has a heavy load, which can be solved by upgrading the PolarDB-X configuration. If both the logical RT and physical RT are high, you can upgrade the ApsaraDB RDS for MySQL configuration or optimize SQL statements on the ApsaraDB RDS for MySQL instance.

11.4.9.4.3. Example 3: Logical QPS and physical QPS

You can observe the difference between the logical queries per second (QPS) and physical QPS.

According to the monitoring data, the logical QPS and physical QPS have the same trends, but the difference between the two is relatively large and in a certain proportion.

Logical QPS



Physical QPS



As shown in the preceding figures, logical QPS fluctuates between 80 and 150, and physical QPS fluctuates between 700 and 1,200.

The reason is that PolarDB-X generates physical SQL statements based on logical SQL statements. The ratio of logical SQL statements to physical SQL statements is not necessarily 1:1. For example, a PolarDB-X logical table is created by using the following statement:

```
CREATE TABLE drds_user
(id int,
name varchar(30))
dbpartition by hash(id);
```

When the query condition contains the database shard key, PolarDB-X pushes the logical SQL statement down to the ApsaraDB RDS for MySQL instance for execution. According to the execution plan, the number of physical SQL statements is 1:

```
mysql> explain select name from drds_user where id = 1;
+-----+-----+-----+
| GROUP_NAME | SQL | PARAMS |
+-----+-----+-----+
| SANGUAN_BSQT_0001_RDS | select `drds_user`.`name` from `drds_user` where (`drds_user`.`id` = 1) | {} |
+-----+-----+-----+
```

When the query does not contain the database shard key, PolarDB-X partitions the logical SQL statement into multiple physical SQL statements. The following execution plan shows that there are eight physical SQL statements:

```
mysql> explain select name from drds_user where name = 'LiLei';
+-----+-----+-----+
| GROUP_NAME | SQL | PARAMS |
+-----+-----+-----+
| SANGUAN_BSQT_0001_RDS | select `drds_user`.`name` from `drds_user` where (`drds_user`.`name` = 'LiLei') | {} |
| SANGUAN_BSQT_0001_RDS | select `drds_user`.`name` from `drds_user` where (`drds_user`.`name` = 'LiLei') | {} |
| SANGUAN_BSQT_0001_RDS | select `drds_user`.`name` from `drds_user` where (`drds_user`.`name` = 'LiLei') | {} |
| SANGUAN_BSQT_0001_RDS | select `drds_user`.`name` from `drds_user` where (`drds_user`.`name` = 'LiLei') | {} |
| SANGUAN_BSQT_0001_RDS | select `drds_user`.`name` from `drds_user` where (`drds_user`.`name` = 'LiLei') | {} |
| SANGUAN_BSQT_0001_RDS | select `drds_user`.`name` from `drds_user` where (`drds_user`.`name` = 'LiLei') | {} |
| SANGUAN_BSQT_0001_RDS | select `drds_user`.`name` from `drds_user` where (`drds_user`.`name` = 'LiLei') | {} |
| SANGUAN_BSQT_0001_RDS | select `drds_user`.`name` from `drds_user` where (`drds_user`.`name` = 'LiLei') | {} |
+-----+-----+-----+
8 rows in set (0.06 sec)
```

Logical or physical QPS indicates the total number of logical or physical SQL statements processed per unit of time. When most SQL statements in the system contain the shard key, the ratio of logical QPS to physical QPS is close to 1:1. If the difference between the logical and physical QPS is too large, many SQL statements of the current application do not contain the shard key. In this case, check the SQL statements of the application to improve performance.

11.4.9.4.4. Example 4: High memory usage

The overly high memory usage of the PolarDB-X instance is mostly caused by the large number of SQL queries in your application and the overlarge result set that is returned. If the memory usage of your PolarDB-X instance remains at about 100%, perform the [Restart a PolarDB-X instance](#) operations to locate and optimize the slow SQL queries of your application.

11.4.10. View the instance version

You can view the PolarDB-X instance version in two ways.

View the instance version in the console

1. [Log on to the PolarDB-X console](#).
2. On the **DRDS Instance Management** page, find the target instance.
3. Click the instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the **Configuration Information** section, view the value of **Current Version**.

View the instance version by using the `version()` function

Connect to the PolarDB-X instance through the MySQL command line and run the `SELECT version()` statement to view the version of the PolarDB-X instance.

```
mysql> select version();
+-----+
| VERSION() |
+-----+
| 5.6.29-TDDL-5.1.28-1320920 |
+-----+
1 row in set (0.00 sec)
```

In the preceding statement, 5.1.28-1320920 indicates the version of the PolarDB-X instance.

11.5. Account management

11.5.1. Terms

This topic describes the terms of the PolarDB-X account and permission system.

The usage of the account and permission system in PolarDB-X is the same as that in MySQL. PolarDB-X supports statements such as `GRANT`, `REVOKE`, `SHOW GRANTS`, `CREATE USER`, `DROP USER`, and `SET PASSWORD`. Currently, PolarDB-X allows you to grant permissions at the database and table levels, but does not allow you to grant permissions at the global or column level.

For more information about the MySQL account and permission system, see [MySQL official documentation](#).

Notice Accounts created by using `CREATE USER` in PolarDB-X only exist in the PolarDB-X instance and will not be synchronized to the backend ApsaraDB RDS for MySQL instances.

Account

An account is specified by the user name and hostname in the `username@'host'` format. Accounts with the same user name but different hostnames are different accounts. For example, `lily@30.9.73.96` and `lily@30.9.73.100` are two different accounts, and their passwords and permissions may be different.

After a database is created in the PolarDB-X console, the system automatically creates two system accounts for the database: administrator account and read-only account. These two accounts are built-in accounts. You cannot delete them or modify their permissions.

- The administrator account name is the same as the database name. For example, if the database name is easydb, the administrator account name is easydb.
- The read-only account name is the database name suffixed with _RO. For example, if the database name is easydb, the read-only account name is easydb_RO.

Assume that the dreamdb and andoradb databases are available. According to the preceding rules, the dreamdb database contains the dreamdb administrator account and the dreamdb_RO read-only account, while the andoradb database contains the andoradb administrator account and the andoradb_RO read-only account.

Account rules

- Each administrator account has all permissions.
- Only the administrator account can create accounts and grant permissions. Other accounts can only be created and granted permissions by the administrator account.
- The administrator account is bound to a database and does not have permissions on other databases. It can only access the bound database, and cannot grant permissions of other databases to an account. For example, the easydb administrator account can only connect to the easydb database, and can only grant permissions of the easydb database or tables in the easydb database to an account.
- A read-only account has only the SELECT permission.

User name rules

- User names are case-insensitive.
- A user name must be 4 to 20 characters in length.
- A user name must start with a letter.
- A user name can contain letters and digits.

Password rules

- A password must be 6 to 20 characters in length.
- A password can contain letters, digits, and special characters (@ # \$ % ^ & + =).

Hostname matching rules

- A hostname must be an IP address. It can contain underscores (_) and wildcards (%). An underscore (_) indicates a character and a wildcard (%) indicates no or more characters. Quote hostnames that contain wildcards with single quotation marks ('), for example, `lily@'30.9.%.%'` and `david@'%'`.
- If there are two user names that can be used to log on to the system, the user name with the longest prefix (the longest IP segment excluding wildcards) prevails. For example, if the `david@'30.9.12_.234'` and `david@'30.9.1%.234'` user names are available in the system, use `david@'30.9.12_.234'` to log on from the 30.9.127.234 host as david.
- When you enable the Virtual Private Cloud (VPC) access feature for a host, the IP address of the host changes. To avoid invalid configurations in the account and permission system, set the hostname to `'%'` to match any IP address.

Permission support

Permission support by level

- Global permission (not supported currently)
- Database-level permission (supported)
- Table-level permission (supported)
- Column-level permission (not supported currently)
- Subprogram-level permission (not supported currently)

Permissions

Currently, eight table-associated basic permissions are supported: `CREATE`, `DROP`, `ALTER`, `INDEX`, `INSERT`, `DELETE`, `UPDATE`, and `SELECT`.

- The `TRUNCATE` statement requires the table-level `DROP` permission.
- The `REPLACE` statement requires the table-level `INSERT` and `DELETE` permissions.
- `CREATE INDEX` and `DROP INDEX` statements are supported.
- The `CREATE SEQUENCE` statement requires the database-level `CREATE` permission.
- The `DROP SEQUENCE` statement requires the database-level `DROP` permission.
- The `ALTER SEQUENCE` statement requires the database-level `ALTER` permission.
- The `INSERT ON DUPLICATE UPDATE` statement requires the table-level `INSERT` and `UPDATE` permissions.

Permission rules

- Permissions are bound to an account (`username@'host'`) rather than a user name (`username`).
- An error occurs if the table does not exist during authorization.
- The database permissions in descending order are as follows: global permissions (not supported currently), database-level permissions, table-level permissions, and column-level permissions (not supported currently). A granted higher-level permission overwrites a lower-level permission. If you remove the higher-level permission, the lower-level permission is also removed.
- `USAGE` authorization is not supported.

11.5.2. Create an account

This topic describes how to create a PolarDB-X account in the PolarDB-X console and by using SQL statements.

Create an account in the console

1. [Log on to the PolarDB-X console](#).
2. On the **DRDS Instance Management** page, find the target instance.
3. Click the instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Configuration and Management > Account Management**.
5. On the **Account Management** page, click **Create Account** in the upper-right corner.
6. Configure the following parameters.

| Parameter | Description |
|----------------------|--|
| Database Account | Enter the account name. The account name must meet the following requirements: <ul style="list-style-type: none"> ◦ The name must be 4 to 20 characters in length. ◦ The name must start with a letter and end with a letter or digit. ◦ The name can contain letters, digits, and underscores (<code>_</code>). |
| New Password | Enter an account password. The password must meet the following requirements: <ul style="list-style-type: none"> ◦ The password must be 8 to 32 characters in length. ◦ The password must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. ◦ Special characters include <code>!@#\$%^&*()_+ -=</code> |
| Confirm New Password | Enter the password again. |

| Parameter | Description |
|-------------------------|--|
| Authorization Databases | <p>You can grant permissions on one or multiple databases to the account.</p> <ol style="list-style-type: none"> Select one or more databases in the left-side section, and click Add to add them to the right-side section. In the right-side section, select Read/Write, Read-only, DDL Only, or DML Only. <div style="background-color: #e1f5fe; padding: 10px; border: 1px solid #cfcfcf;"> <p> Note You can also grant permissions on multiple authorized databases by clicking Set All to Read-only, Set All to DDL Only, Set All to DML Only, or Set All to Read/Write in the upper-right corner of the right-side section.</p> <p>The button in the upper-right corner changes as you click it. For example, after you click Set All to Read-only, the button is changed to Set All to DDL Only.</p> </div> |

Create an account in the command line

Syntax rules:

```
CREATE USER user_specification [, user_specification] ...
user_specification: user [ auth_option ]
auth_option: IDENTIFIED BY 'auth_string'
```

Examples:

Create an account with the user name lily and password 123456, which can be used to log on only from 30.xx.xx.96.

```
CREATE USER lily@30.xx.xx.96 IDENTIFIED BY '123456';
```

Create an account named david with no password, which can be used to log on from any host.

```
CREATE USER david@'%';
```

11.5.3. Reset the password

When using PolarDB-X, you can reset the password of your database account in the PolarDB-X console or by using the command line.

Note

- Accounts with ROOT permissions cannot be deleted or modified.
- For data security, we recommend that you change your password periodically.

Reset the password in the console

- Log on to the PolarDB-X console.
- On the DRDS Instance Management page, find the target instance.
- Click the instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
- In the left-side navigation pane, choose **Configuration and Management > Account Management**.

5. Find the target account, and click **Reset Password**.
6. In the **Reset Account Password** dialog box, set **New Password** and **Confirm New Password**.

- Note** The password must meet the following requirements:
- The password must be 8 to 32 characters in length.
 - The password must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters.
 - Special characters include ! @#\$%^&*()_+ -=

7. After you confirm that the password is correct, click **OK**.

Reset the password in the command line

Syntax rules:

```
SET PASSWORD FOR user = password_option
password_option: {
  PASSWORD('auth_string')
}
```

Examples:

Change the password of the account lily@30.xx.xx.96 to 123456.

```
SET PASSWORD FOR lily@30.xx.xx.96 = PASSWORD('123456')
```

11.5.4. Modify account permissions

You can modify the account permissions of your instances at any time when using PolarDB-X.

Precautions

- Privileged accounts cannot be modified.
- In the console, you can only grant data manipulation language (DML), data definition language (DDL), read-only, and read/write permissions to standard accounts. To grant more permissions, use the command line.

Modify account permissions in the console

1. [Log on to the PolarDB-X console](#).
2. On the **DRDS Instance Management** page, find the target instance.
3. Click the instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Configuration and Management > Account Management**.
5. Find the target account and click **Modify Permission**.
6. In the **Modify Permissions** dialog box, **grant** or **remove** the permissions on one or more databases to or from the account.
 - Authorize a database:

Select one or more databases in the left-side section, and click **Add** to add them to the right-side section.
 - Remove an authorized database:

Select one or more databases in the right-side section, and click **Remove** to remove them. The removed databases are displayed in the left-side section.

- Modify permissions of an authorized database:

In the right-side section, select **Read/Write**, **Read-only**, **DDL Only** or **DML Only**.

 **Note** You can also grant permissions on multiple authorized databases by clicking **Set All to Read-only**, **Set All to DDL Only**, **Set All to DML Only**, or **Set All to Read/Write** in the upper-right corner of the right-side section.

The button in the upper-right corner changes as you click it. For example, after you click **Set All to Read-only**, the button is changed to **Set All to DDL Only**.

7. After the configuration is complete, click **OK**.

GRANT statements

Syntax rules:

```
GRANT
  priv_type[, priv_type] ...
  ON priv_level
  TO user_specification [, user_specification] ...
  [WITH GRANT OPTION]
priv_level: {
  | db_name.*
  | db_name.tbl_name
  | tbl_name
}
user_specification:
  user [ auth_option ]
auth_option: {
  IDENTIFIED BY 'auth_string'
}
```

Notice

- If the account in the GRANT statement does not exist and no IDENTIFIED BY information is provided, an error message indicating that the account does not exist is returned.
- If the account specified in the GRANT statement does not exist but the IDENTIFIED BY information is provided, the account is created and granted with the specified permission.

For example, in the easydb database, create an account named david, which can be used to log on from any host and has all the permissions on easydb.

Method 1: Create an account and then grant permissions to the account.

```
CREATE USER david@%' IDENTIFIED BY 'your#password';
GRANT ALL PRIVILEGES ON easydb.* to david@%';
```

Method 2: Create an account and grant permissions to the account by using one statement.

```
GRANT ALL PRIVILEGES ON easydb.* to david@%' IDENTIFIED BY 'your#password';
```

In the easydb database, create an account named hanson, which can be used to log on from any host and has all the permissions on the easydb.employees table.

```
GRANT ALL PRIVILEGES ON easydb.employees to hanson@%' IDENTIFIED BY 'your#password';
```

In the easydb database, create an account named hanson, which can be used to log on only from 192.xx.xx.10 and has the INSERT and SELECT permissions on the easydb.emp table.

```
GRANT INSERT,SELECT ON easydb.emp to hanson@'192.xx.xx.10' IDENTIFIED BY 'your#password';
```

In the easydb database, create a read-only account named actro, which can be used to log on from any host.

```
GRANT SELECT ON easydb.* to actro@'%' IDENTIFIED BY 'your#password';
```

REVOKE statements

Syntax rules:

- Delete the permissions at a certain level from an account. The permission level is specified by priv_level.

```
REVOKE
priv_type
[, priv_type] ...
ON priv_level
FROM user [, user] ...
```

- Delete all permissions of the account at the database and table levels.

```
REVOKE ALL PRIVILEGES, GRANT OPTION
FROM user [, user] ...
```

Examples:

Delete the CREATE, DROP, and INDEX permissions from hanson@'%' on the easydb.emp table.

```
REVOKE CREATE,DROP,INDEX ON easydb.emp FROM hanson@'%';
```

Delete all permissions from the account lily@30.xx.xx.96.

```
REVOKE ALL PRIVILEGES,GRANT OPTION FROM lily@30.xx.xx.96;
```



Notice GRANT OPTION must be added to the statement for compatibility with MySQL.

SHOW GRANTS statements

Syntax rules:

```
SHOW GRANTS[FOR user@host];
```

Query all permissions:

```
SHOW GRANTS;
```

Query the permissions of an account:

```
SHOW GRANTS FOR user@host;
```

11.5.5. Delete an account

You can delete an account in the Cloud Native Distributed Database PolarDB-X (PolarDB-X) console or by using the command line.

Delete an account in the console

 **Note** You can only delete standard accounts that are created in the console.

1. [Log on to the PolarDB-X console.](#)
2. On the **DRDS Instance Management** page, find the target instance.
3. Click the instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Configuration and Management > Account Management**.
5. Find the target account and click **Delete**.
6. In the **Delete Account** dialog box, click **OK**.

Delete an account by using the command line

Syntax rules:

```
DROP USER user [, user] ...
```

Examples:

Delete the account lily@30.xx.xx.96:

```
DROP USER lily@30.xx.xx.96;
```

11.6. Database management

11.6.1. Create a database

After you create a PolarDB-X instance, create a database that runs on one or more ApsaraDB RDS for MySQL instances.

Prerequisites

- An ApsaraDB RDS for MySQL instance is created in the same department of the PolarDB-X instance.
- The permissions on Resource Access Management (RAM) resources are granted. For more information about how to grant permissions, see *RAM management* in *Apsara Uni-manager user guide*.

1. [Log on to the PolarDB-X console.](#)
2. Find the target instance in the instance list.
3. Click the target instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. On the **Basic Information** page, click **Create Database** in the upper-right corner.
5. In the **Enter Basic Information** step, configure the following parameters:

| Parameter | Description |
|----------------|--------------------------------------|
| Partition Mode | Select Horizontal Partition . |

| Parameter | Description |
|--------------------|--|
| Database Name | Enter a custom name for the PolarDB-X database. The name must meet the following requirements: <ul style="list-style-type: none"> The name must be 2 to 24 characters in length. The name must start with a letter and end with a letter or a digit. The name can contain lowercase letters, digits, and underscores (_). The name must be unique on the instance. |
| Character Set | Select utf8, gbk, latin1, or utf8mb4. |
| DRDS Link Password | Set the password for the database in PolarDB-X. The password must meet the following requirements: <ul style="list-style-type: none"> The password must be 8 to 30 characters in length. The password must contain at least three of the following types: uppercase letters, lowercase letters, digits, and underscores (_). |
| Confirm Password | Enter the password again. |

6. Click **Next**.
7. In the **Select RDS Instance** step, select **Buy New RDS Instance** or **Use Existing RDS Instance**.
 - i. **Buy New RDS Instance**: Click the **Buy New RDS Instance** tab.
 - ii. Set **Storage Type**, **Series**, **Instance Specifications**, **Storage Capacity**, **Availability Zone**, and **Quantity**.
 - iii. Click **Next**.
 - i. **Use Existing RDS Instance**: Click the **Use Existing RDS Instance** tab.
 - ii. In the left-side section, click the ApsaraDB RDS for MySQL instances that you want to add.
 - iii. Click to add the selected instances to the **Selected RDS Instances** section on the right side.
 - iv. Click **Next**.
8. After the precheck is successful in the **Precheck** step, click **Next**.

 **Note** If the precheck fails, fix the issue based on the instructions on the page.

9. In the **Preview** step, click **Next**. Wait until the database is created.

11.6.2. View a database

After the database is created, you can view the basic information of the database in database management in the console.

Procedure

1. [Log on to the PolarDB-X console](#).
2. On the **DRDS Instance Management** page, find the target instance.
3. Click the instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Configuration and Management > Database Management**.
5. Find the target database, and click **Manage**. The **Basic Information** page of the database appears.

 **Note** On the **Basic Information** page, you can delete the database or reset the password.

What's next

PolarDB-X is fully compatible with the MySQL protocol. You can use **Command Line URL** on the MySQL client to connect to the PolarDB-X instance and enter the user name and password to log on to the PolarDB-X database. When using the MySQL client, note the following points:

Note

- Some MySQL clients of earlier versions have limits on the user name length, which cannot be more than 16 characters. The PolarDB-X database name and user name are the same. If the database name exceeds 16 characters, an error is reported.
- When using the MySQL client, you must add the `-c` parameter to the hint command. In PolarDB-X, HINT is implemented by using annotations. If the `-c` parameter is not added, the annotation is lost and the PolarDB-X hint is lost.

11.6.3. Perform smooth scale-out

When the underlying storage of the logical database reaches the physical bottleneck, for example, when the remaining disk space is about 30%, you can smoothly scale it out to improve the performance. The smooth scale-out process is divided into four steps: configuration > migration > switchover > cleanup.

Configuration

Note

In smooth scale-out, ApsaraDB RDS for MySQL instances are added, and some source database shards are migrated to the new ApsaraDB RDS for MySQL instances. In this way, the overall data storage capacity is increased and the number of requests that a single ApsaraDB RDS for MySQL instance needs to process is reduced.

1. [Log on to the PolarDB-X console](#).
2. On the **DRDS Instance Management** page, find the target instance.
3. Click the instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Configuration and Management > Database Management**.
5. Find the target database, and click **Manage**. The **Basic Information** page of the database appears.
6. In the left-side navigation pane, choose **Configuration and Management > Scale-out Management**.
7. In the upper-right corner of the **Scale-out Management** page, click **Scale Out**.
8. Select **Smooth Scale-out**, and then click **Next**.
9. After all prechecks are passed on the **Precheck** page, click **Next**.

Note

If a precheck fails, rectify the configuration as prompted.

10. On the **Select RDS Instance** page, you can click the **Buy New RDS Instance** or **Use Existing RDS Instance** tab.
 - i. **Buy New RDS Instance**: Click the **Buy New RDS Instance** tab.
 - ii. Set **Storage Type**, **Edition**, and **Storage Capacity**.
 - iii. Click **Next**.
 - i. **Use Existing RDS Instance**: Click the **Use Existing RDS Instance** tab.

- ii. Click the ApsaraDB RDS for MySQL instances to be added on the left.
 - iii. Click  to move the selected instances to the **Selected RDS Instances** section on the right.
 - iv. Click **Next**.
11. On the **Preview** page, click **Start Scale-out**.

 **Note** By default, the console evenly distributes the physical database shards to the ApsaraDB RDS for MySQL instances you added. You can also manually add or delete physical database shards to or from the new ApsaraDB RDS for MySQL instances.

12. Click the  icon in the upper-right corner to view the details of the scale-out task.

Migration

Some physical database shards are migrated during smooth scale-out.

The migration does not change the data in the source database, and therefore it does not affect online services. Before the switchover, you can cancel the smooth scale-out operation through a rollback.

Note

- This is because before the switchover, the current scale-out operation does not have a real impact on the data in the source database.
- During scale-out, the binary log files of the source ApsaraDB RDS for MySQL instance are not cleaned, which may result in insufficient disk space. Therefore, you must reserve sufficient disk space on the source ApsaraDB RDS for MySQL instance. Generally, the remaining disk space should be more than 30%. If the disk space cannot be guaranteed, you can submit a ticket to expand the ApsaraDB RDS for MySQL storage space.
- To reduce the pressure of read operations on the source ApsaraDB RDS for MySQL instance, perform scale-out when the load on the source ApsaraDB RDS for MySQL instance is low.
- During the scale-out, do not submit data description language (DDL) tasks in the console or connect to the PolarDB-X instance to directly run DDL statements. Otherwise, the scale-out task may fail.
- Make sure that all tables in the source database have primary keys before scale-out.

After historical data and incremental data are migrated, the migration progress reaches 100%. Then, you can **switch** the read and write traffic to the new ApsaraDB RDS for MySQL instance or **roll back** to cancel this scale-out.

Switchover

The switchover task switches the read and write traffic to the new ApsaraDB RDS for MySQL instance. The whole process takes 3 to 5 minutes. During the switchover process, the service is not affected except for one or two transient disconnections. Perform the switchover during off-peak hours.

1. In the upper-right corner of the **Basic Information** page, click the  icon. The **Task Progress** dialog box appears.
2. In the **Task Progress** dialog box, click **Switch Over** and then **OK**.
During the switchover process, a switchover task is generated and displayed in the **Task Progress** dialog box.
3. After the switchover is complete, the **Clean Up** button appears in the **Task Progress** dialog box, which means that the switchover task is complete.

Cleanup

In this step, the migrated database shards are deleted from the source ApsaraDB RDS for MySQL instance.

1. After switchover is complete, click **Clean Up** next to the target task.
2. Click **OK**. A cleanup task appears in the Task Progress dialog box.

 **Note**

- The cleanup task is an asynchronous task. You can view the execution status in the Task Progress dialog box.
- After the cleanup task is complete, the smooth scale-out process ends. The new ApsaraDB RDS for MySQL instance becomes the storage node of the PolarDB-X logical database.
- Currently, you can implement smooth scale-out by migrating physical database shards. If no further scale-out is allowed after the number of database shards exceeds the capacity of a single ApsaraDB RDS for MySQL instance, you can submit a ticket to apply for increasing the number of database shards and scaling out the database. In this case, Hash calculation is performed again to reallocate data.
- The cleanup task deletes database shards that are no longer used after the current scale-out. You can back up the database shards before running the cleanup task.
- The cleanup operation brings pressure to databases. We recommend that you perform this operation during off-peak hours.

11.6.4. View database monitoring information

PolarDB-X displays the historical monitoring information of a PolarDB-X database in two dimensions: data metrics and query time.

Procedure

1. [Log on to the PolarDB-X console.](#)
2. On the **DRDS Instance Management** page, find the target instance.
3. Click the instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Monitoring and Alerts > Database Monitoring**.
5. Select **Data Indexes** and **Query Time**. Then, the corresponding monitoring information appears.

 **Note** For more information about instance-level monitoring, see [View monitoring information](#).

11.6.5. Set the IP address whitelist

PolarDB-X provides the access control function. Only IP addresses in the whitelist of a database can access the database.

Procedure

1. [Log on to the PolarDB-X console.](#)
2. On the **DRDS Instance Management** page, find the target instance.
3. Click the instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Configuration and Management > Database Management**.
5. Find the target database, and click **Manage**. The **Basic Information** page of the database appears.
6. On the **Basic Information** page of the database, choose **Data Security > Whitelist Settings** from the left-

side navigation pane.

7. On the Whitelist Settings page, click **Manually Modify**.
8. Enter the IP address that is allowed to access the database, and click **OK**.

Note

- The whitelist supports the following formats:
 - IP address, for example, 192.168.1.1.
 - CIDR IP address, for example, 192.168.1.1/24.
 - IP address with an asterisk (*) as the wildcard, for example, 192.168.1.*, indicating that access is allowed from any host with an IP address in the range of 192.168.1.1 to 192.168.1.254.
 - CIDR block, for example, 192.168.1.1-192.168.1.254.
- If you want to add multiple IP addresses or CIDR blocks, separate them with commas (,) without spaces, for example, 192.168.0.1,172.16.213.9.

11.6.6. Delete a database

This topic describes how to delete a database in the Cloud Native Distributed Database PolarDB-X (PolarDB-X) console.

Procedure

1. [Log on to the PolarDB-X console](#).
2. On the **DRDS Instance Management** page, find the target instance.
3. Click the instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.
4. In the left-side navigation pane, choose **Configuration and Management > Database Management**.
5. Find the target database and click **Delete**.

 **Warning** You cannot recover databases that have been deleted. Exercise caution when you perform this operation.

6. In the **Delete Database** dialog box, click **OK**.

11.6.7. Fix database shard connections

Context

When using a PolarDB-X instance, you need to access ApsaraDB RDS for MySQL. If the network configuration of the connected ApsaraDB RDS for MySQL instance changes, for example, if the zone is switched or the network type is changed from classic to Virtual Private Cloud (VPC), the network connection between the PolarDB-X instance and the ApsaraDB RDS for MySQL instance is broken. As a result, the PolarDB-X instance cannot access the ApsaraDB RDS for MySQL instance. In this case, you must manually fix the database shard link in the PolarDB-X console to restore the network connection from the PolarDB-X instance to the ApsaraDB RDS for MySQL instance.

Procedure

1. [Log on to the PolarDB-X console](#).
2. On the **DRDS Instance Management** page, find the target instance.
3. Click the instance ID or choose **More > Manage** from the Actions column of the instance to access the **Basic Information** page.

4. In the left-side navigation pane, choose **Configuration and Management > Database Management**.
5. Find the target database, and click **Manage**. The **Basic Information** page of the database appears.
6. In the **Shortcuts** section, click **Fix Database Shard Connections**.
7. In the pop-up window, click **OK**.

11.7. Custom control commands

PolarDB-X provides a series of auxiliary SQL commands to help you conveniently use PolarDB-X.

11.7.1. Overview

PolarDB-X provides unique auxiliary statements for you to use and maintain PolarDB-X.

Syntax description: The identifier provided by the user is in [] and optional content is in (). In addition, this document applies to the current version. If some statements are unavailable, the version is earlier than required.

11.7.2. Help statements

This topic describes all the auxiliary SQL commands of PolarDB-X and their descriptions.

SHOW HELP statements:

```
mysql> show help;
+-----+-----+-----+
| STATEMENT          | DESCRIPTION                               | EXAMPLE                               |
+-----+-----+-----+
| show rule          | Report all table rule                     |                                       |
| show rule from TABLE | Report table rule                         | show rule from user                   |
| show full rule from TABLE | Report table full rule                   | show full rule from user              |
| show topology from TABLE | Report table physical topology           | show topology from user               |
| show partitions from TABLE | Report table dbPartition or tbPartition columns | show partitions from user             |
|
| show broadcasts   | Report all broadcast tables               |                                       |
| show datasources   | Report all partition db threadPool info   |                                       |
| show node          | Report master/slave read status           |                                       |
| show slow          | Report top 100 slow sql                    |                                       |
| show physical_slow | Report top 100 physical slow sql          |                                       |
| clear slow         | Clear slow data                           |                                       |
| trace SQL          | Start trace sql, use show trace to print profiling data | trace select count(*) from user; show trace |
| show trace         | Report sql execute profiling info         |                                       |
| explain SQL        | Report sql plan info                      | explain select count(*) from user     |
| explain detail SQL | Report sql detail plan info               | explain detail select count(*) from user |
| explain execute SQL | Report sql on physical db plan info       | explain execute select count(*) from user |
| show sequences     | Report all sequences status               |                                       |
| create sequence NAME [start with COUNT] | Create sequence                           | create sequence test start with 0     |
| alter sequence NAME [start with COUNT] | Alter sequence                             | alter sequence test start with 100000 |
| drop sequence NAME | Drop sequence                             | drop sequence test                    |
+-----+-----+-----+
20 rows in set (0.00 sec)
```

11.7.3. Statements for viewing rules and node topologies

SHOW RULE [FROM tablename]

Usage notes:

- `show rule` : shows the partitioning status of each logical table in a database.
- `show rule from tablename` : shows the partitioning status of a specified logical table in a database.

The following describes the meanings of important columns:

- **BROADCAST**: indicates whether the table is a broadcast table. 0 indicates "No" and 1 indicates "Yes".
- **DB_PARTITION_KEY**: indicates the database shard key. If no database shards exist, the parameter value is NULL.
- **DB_PARTITION_POLICY**: indicates the database sharding policy. Options are Hash and date policies such as YYYYMM, YYYYDD, and YYYYWEEK.
- **DB_PARTITION_COUNT**: indicates the number of database shards.
- **TB_PARTITION_KEY**: indicates the table shard key. If no table shards exist, the parameter value is NULL.
- **TB_PARTITION_POLICY**: indicates the table sharding policy. Options are Hash and date policies such as MM, DD, MMDD, and WEEK.
- **TB_PARTITION_COUNT**: indicates the number of table shards.

```
mysql> show rule;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
-+
| ID | TABLE_NAME | BROADCAST | DB_PARTITION_KEY | DB_PARTITION_POLICY | DB_PARTITION_COUNT | TB_PARTITION_KEY | TB_PARTITION_POLICY | TB_PARTITION_COUNT |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
-+
| 0 | dept_manager | 0 | NULL | 1 | NULL | 1 |
| 1 | emp | 0 | emp_no | hash | 8 | id | hash | 2 |
| 2 | example | 0 | shard_key | hash | 8 | NULL | NULL | 1 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
-+
3 rows in set (0.01 sec)
```

SHOW FULL RULE [FROM tablename]

You can run this SQL statement to view the sharding rules of logical tables in a database. It displays more detailed information than the SHOW RULE command.

The following describes the meanings of important columns:

- **BROADCAST**: indicates whether the table is a broadcast table. 0 indicates "No" and 1 indicates "Yes".
- **JOIN_GROUP**: a reserved field.
- **ALLOW_FULL_TABLE_SCAN**: indicates whether to allow data querying when no table shard key is specified for database or table sharding. If this parameter is set to True, each physical table is scanned to find data that meets the condition, which is a full table scan.
- **DB_NAME_PATTERN**: The value 0 between {} in DB_NAME_PATTERN is a placeholder. When the SQL statement is run, this value is replaced by the value of DB_RULES_STR, with the number of digits unchanged. For example, if the value of DB_NAME_PATTERN is SEQ_{0000}_RDS and the value of DB_RULES_STR is [1,2,3,4], four DB_NAME values are generated: SEQ_0001_RDS, SEQ_0002_RDS, SEQ_0003_RDS, and SEQ_0004_RDS.
- **DB_RULES_STR**: indicates the database sharding rule.
- **TB_NAME_PATTERN**: The value 0 between {} in TB_NAME_PATTERN is a placeholder. When the SQL statement is run, this value is replaced by the value of TB_RULES_STR, with the number of digits unchanged. For example, if the value of TB_NAME_PATTERN is table_{00} and the value of TB_RULES_STR is [1,2,3,4,5,6,7,8], eight tables are generated: table_01, table_02, table_03, table_04, table_05, table_06, table_07, and table_08.
- **TB_RULES_STR**: indicates the table sharding rule.
- **PARTITION_KEYS**: indicates a set of database and table shard keys. When database sharding and table sharding coexist, the database shard key is placed before the table shard key.
- **DEFAULT_DB_INDEX**: indicates the database shard in which a single database and a single table are stored.

```
mysql> show full rule;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| ID | TABLE_NAME | BROADCAST | JOIN_GROUP | ALLOW_FULL_TABLE_SCAN | DB_NAME_PATTERN | DB_RULES_STR | TB_NAME_PATTERN | TB_RULES_STR | PARTITION_KEYS | DEFAULT_DB_INDEX |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 0 | dept_manager | 0 | NULL | 0 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0000_RDS | NULL | dept_manager | NULL | NULL | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0000_RDS |
| 1 | emp | 0 | NULL | 1 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_{0000}_RDS | ((#emp_no,1,8#).longValue().abs() % 8) | emp_{0} | ((#id,1,2#).longValue().abs() % 2) | emp_no | id | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0000_RDS |
| 2 | example | 0 | NULL | 1 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_{0000}_RDS | ((#shard_key,1,8#).longValue().abs() % 8).intdiv(1) | example | NULL | shard_key | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0000_RDS |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
3 rows in set (0.01 sec)
```

SHOW TOPOLOGY FROM tablename

You can run this SQL statement to view the topology of a specified logical table, that is, the database shards to which data in the logical table is partitioned and the table shards in each database shard.

```
mysql> show topology from emp;
+-----+-----+-----+
| ID | GROUP_NAME | TABLE_NAME |
+-----+-----+-----+
| 0 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0000_RDS | emp_0 |
| 1 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0000_RDS | emp_1 |
| 2 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0001_RDS | emp_0 |
| 3 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0001_RDS | emp_1 |
| 4 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0002_RDS | emp_0 |
| 5 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0002_RDS | emp_1 |
| 6 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0003_RDS | emp_0 |
| 7 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0003_RDS | emp_1 |
| 8 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0004_RDS | emp_0 |
| 9 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0004_RDS | emp_1 |
| 10 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0005_RDS | emp_0 |
| 11 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0005_RDS | emp_1 |
| 12 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0006_RDS | emp_0 |
| 13 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0006_RDS | emp_1 |
| 14 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0007_RDS | emp_0 |
| 15 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0007_RDS | emp_1 |
+-----+-----+-----+
16 rows in set (0.01 sec)
```

SHOW PARTITIONS FROM tablename

You can run this SQL statement to view the set of database shard keys and table shard keys, which are separated by commas (.). If the final result contains two values, both database sharding and table sharding are performed. The first value is the database shard key and the second value is the table shard key. If only one value is returned, only database sharding is performed. This value is the database shard key.

```
mysql> show partitions from emp;
+-----+
| KEYS |
+-----+
| emp_no,id |
+-----+
1 row in set (0.00 sec)
```

SHOW BROADCASTS

You can run this SQL statement to view the list of broadcast tables.

```
mysql> show broadcasts;
+-----+
| ID | TABLE_NAME |
+-----+
| 0 | brd2 |
| 1 | brd_tbl |
+-----+
2 rows in set (0.01 sec)
```

SHOW DATASOURCES

You can run this SQL statement to view the information about the underlying storage, including the database name, database group name, connection URL, user name, storage type, read/write weight, and connection pool information.

The following describes the meanings of important columns:

- **SCHEMA**: indicates the database name.
- **GROUP**: indicates the database group name. Grouping aims to manage multiple groups of databases that have identical data, such as the primary and secondary databases after data replication through ApsaraDB RDS for MySQL. It is mainly used for read/write splitting and primary/secondary switchovers.
- **URL**: indicates the connection information of the underlying ApsaraDB RDS for MySQL instance.
- **TYPE**: indicates the type of the underlying storage. Currently, only ApsaraDB RDS for MySQL instances are supported.
- **READ_WEIGHT**: indicates the read weight of the database. When the primary ApsaraDB RDS for MySQL instance is under a heavy load of many read requests, you can use the read/write splitting function of PolarDB-X to distribute the read traffic. PolarDB-X automatically identifies the read and write traffic. It directs the write traffic to the primary ApsaraDB RDS for MySQL instance and the read traffic to all ApsaraDB RDS for MySQL instances based on the configured weight.
- **WRITE_WEIGHT**: indicates the write weight. For more information, see **READ_WEIGHT**.

ApsaraDB RDS for MySQL instance, not the configured percentage.

- **SLAVE_READ_PERCENT**: indicates the actual percentage of read-only queries processed by the secondary ApsaraDB RDS for MySQL instances, not the configured percentage.

Note

- Read-only queries in transactions are sent to the primary ApsaraDB RDS for MySQL instance.
- The **MASTER_READ_PERCENT** and **SLAVE_READ_PERCENT** fields indicate the accumulative historical data. After the read/write weight ratio has been changed, these values do not immediately reflect the latest read/write weight ratio, which appears after a long period of time has passed.

```
mysql> show node;
```

| ID | NAME | MASTER_READ_COUNT | SLAVE_READ_COUNT | MASTER_READ_PERCENT | SLAVE_READ_PERCENT |
|----|--|-------------------|------------------|---------------------|--------------------|
| 0 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0000_RDS | 12 | 0 | 100% | 0% |
| 1 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0001_RDS | 0 | 0 | 0% | 0% |
| 2 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0002_RDS | 0 | 0 | 0% | 0% |
| 3 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0003_RDS | 0 | 0 | 0% | 0% |
| 4 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0004_RDS | 0 | 0 | 0% | 0% |
| 5 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0005_RDS | 0 | 0 | 0% | 0% |
| 6 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0006_RDS | 0 | 0 | 0% | 0% |
| 7 | SEQ_TEST_1487767780814RGKKSEQ_TEST_WNJG_0007_RDS | 0 | 0 | 0% | 0% |

8 rows in set (0.01 sec)

11.7.4. SQL tuning statements

SHOW [FULL] SLOW [WHERE expr] [limit expr]

SQL statements that take more than 1 second to execute are slow SQL statements. Slow logical SQL statements are the slow SQL statements sent from an application to a PolarDB-X instance.

- **SHOW SLOW** : You can run this SQL statement to view the 100 slowest logical SQL queries that are recorded since the PolarDB-X instance is started or the last time when **CLEAR SLOW** is executed.

Note The recorded 100 slowest SQL queries are stored in the PolarDB-X system. When the PolarDB-X instance is restarted or executes **CLEAR SLOW**, these queries will be discarded.

- **SHOW FULL SLOW** : You can run this SQL statement to view all the slow logical SQL queries that are recorded and persisted to the built-in database of the PolarDB-X instance since the PolarDB-X instance is started. The upper limit for the number of records is specified in the specifications of the PolarDB-X instance. The PolarDB-X instance scrolls to delete the earliest slow SQL statements. If the specifications of the PolarDB-X instance is 4-core 4 GB, a maximum of 10,000 slow SQL statements can be recorded, including slow logical and physical SQL statements. If the specifications of the PolarDB-X instance is 8-core 8 GB, a maximum of 20,000 slow SQL statements can be recorded, including slow logical and physical SQL statements. The same rule applies to other specifications.

The following describes the meanings of important columns:

- **HOST**: the IP address of the host from which the SQL statement is sent.
- **START_TIME**: the time when the SQL statement starts to be executed.
- **EXECUTE_TIME**: the time when the SQL statement is executed.
- **AFFECT_ROW**: For data manipulation language (DML) statements, this parameter indicates the number of affected rows. For query statements, this parameter indicates the number of returned records.

```
mysql> show slow where execute_time > 1000 limit 1;
+-----+-----+-----+-----+-----+
| HOST | START_TIME | EXECUTE_TIME | AFFECT_ROW | SQL |
+-----+-----+-----+-----+-----+
| 127.0.0.1 | 2016-03-16 13:02:57 | 2785 | 7 | show rule |
+-----+-----+-----+-----+-----+
1 row in set (0.02 sec)
```

SHOW [FULL] PHYSICAL_SLOW [WHERE expr] [limit expr]

SQL statements that take more than 1 second to execute are slow SQL statements. Slow logical SQL statements are the slow SQL statements sent from an application to a PolarDB-X instance.

- **SHOW SLOW** : You can run this SQL statement to view the 100 slowest logical SQL queries that are recorded since the PolarDB-X instance is started or the last time when **CLEAR SLOW** is executed.

Note The recorded 100 slowest SQL queries are stored in the PolarDB-X system. When the PolarDB-X instance is restarted or executes **CLEAR SLOW**, these queries will be discarded.

- **SHOW FULL SLOW** : You can run this SQL statement to view all the slow logical SQL queries that are recorded and persisted to the built-in database of the PolarDB-X instance since the PolarDB-X instance is started. The upper limit for the number of records is specified in the specifications of the PolarDB-X instance. The PolarDB-X instance scrolls to delete the earliest slow SQL statements. If the specifications of the PolarDB-X instance is 4-core 4 GB, a maximum of 10,000 slow SQL statements can be recorded, including slow logical and physical SQL statements. If the specifications of the PolarDB-X instance is 8-core 8 GB, a maximum of 20,000 slow SQL statements can be recorded, including slow logical and physical SQL statements. The same rule applies to other specifications.

The following describes the meanings of important columns:

- **GROUP_NAME**: the name of the group to which the database that executes the SQL statement belongs.
- **START_TIME**: the time when the SQL statement starts to be executed.
- **EXECUTE_TIME**: the time when the SQL statement is executed.
- **AFFECT_ROW**: For data manipulation language (DML) statements, this parameter indicates the number of affected rows. For query statements, this parameter indicates the number of returned records.

```
mysql> show physical_slow;
+-----+-----+-----+-----+-----+-----+-----+
| GROUP_NAME | DBKEY_NAME | START_TIME | EXECUTE_TIME | SQL_EXECUTE_TIME | GETLOCK_CONNECTION_TIME | CREATE_CONNECTION_TIME | AFFECT_ROW | SQL |
+-----+-----+-----+-----+-----+-----+-----+
| TDDL5_00_GROUP | db218249098_sqa_zmf_tddl5_00_3309 | 2016-03-16 13:05:38 | 1057 | 1011 | 0 | 1 | select sleep(1) |
+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.01 sec)
```

CLEAR SLOW

You can run this SQL statement to clear the 100 slowest logical SQL queries and the 100 slowest physical SQL queries that are recorded since the PolarDB-X instance is started or the last time when **CLEAR SLOW** is executed.

Note Both `SHOW SLOW` and `SHOW PHYSICAL_SLOW` can be executed to display the 100 slowest SQL statements. If `CLEAR SLOW` has not been executed for a long time, these SQL statements may have been recorded a long time ago. Therefore, after SQL tuning statements are executed, we recommend that you execute `CLEAR SLOW`. After the system runs for a while, check the tuning results of slow SQL statements.

```
mysql> clear slow;
Query OK, 0 rows affected (0.00 sec)
```

EXPLAIN SQL

You can run this SQL statement to view the execution plan of a specified SQL statement in the PolarDB-X. Note that this SQL statement is not truly executed.

Example:

You can run this SQL statement to view the execution plan of the SQL `select * from doctest` statement. The doctest table is stored in database shards according to values in the id column. According to the execution plan, the SQL statement will be routed to each database shard for execution, and the execution results will be aggregated.

```
mysql> explain select * from doctest;
+-----+-----+-----+
| GROUP_NAME | SQL | PARAMS |
+-----+-----+-----+
| DOCTEST_1488704345426RCUPDOCTEST_CAET_0000_RDS | select `doctest`.`id` from `doctest` | {} |
| DOCTEST_1488704345426RCUPDOCTEST_CAET_0001_RDS | select `doctest`.`id` from `doctest` | {} |
| DOCTEST_1488704345426RCUPDOCTEST_CAET_0002_RDS | select `doctest`.`id` from `doctest` | {} |
| DOCTEST_1488704345426RCUPDOCTEST_CAET_0003_RDS | select `doctest`.`id` from `doctest` | {} |
| DOCTEST_1488704345426RCUPDOCTEST_CAET_0004_RDS | select `doctest`.`id` from `doctest` | {} |
| DOCTEST_1488704345426RCUPDOCTEST_CAET_0005_RDS | select `doctest`.`id` from `doctest` | {} |
| DOCTEST_1488704345426RCUPDOCTEST_CAET_0006_RDS | select `doctest`.`id` from `doctest` | {} |
| DOCTEST_1488704345426RCUPDOCTEST_CAET_0007_RDS | select `doctest`.`id` from `doctest` | {} |
+-----+-----+-----+
8 rows in set (0.00 sec)
```

You can run this SQL statement to view the execution plan of the SQL `select * from doctest where id = 1` statement. The doctest table is stored in database shards according to values in the id column. According to the execution plan, the PolarDB-X instance will calculate a specified database shard based on the shard key, which is id, directly route the SQL statement to the database shard, and aggregate the execution results.

```
mysql> explain select * from doctest where id = 1;
+-----+-----+-----+
| GROUP_NAME | SQL | PARAMS |
+-----+-----+-----+
| DOCTEST_1488704345426RCUPDOCTEST_CAET_0001_RDS | select `doctest`.`id` from `doctest` where (`doctest`.`id` = 1) | {} |
+-----+-----+-----+
1 row in set (0.01 sec)
```

EXPLAIN DETAIL SQL

You can run this SQL statement to view the execution plan of a specified SQL statement in the PolarDB-X. Note that this SQL statement is not truly executed.

```
mysql> explain detail select * from doctest where id = 1;
+-----+
| GROUP_NAME          | SQL
| PARAMS              |
+-----+
| DOCTEST_1488704345426RCUPDOCTEST_CAET_0001_RDS | Query from doctest as doctest
  keyFilter:doctest.id = 1
  queryConcurrency:SEQUENTIAL
  columns:[doctest.id]
  tableName:doctest
  executeOn:DOCTEST_1488704345426RCUPDOCTEST_CAET_0001_RDS
| NULL |
+-----+
1 row in set (0.02 sec)
```

EXPLAIN EXECUTE SQL

You can run this SQL statement to view the execution plan of underlying storage. This statement is equivalent to the MySQL EXPLAIN statement.

```
mysql> explain execute select * from tddl_mgr_log limit 1;
+----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table   | type | possible_keys | key | key_len | ref | rows | Extra |
+----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE     | tddl_mgr_log | ALL | NULL          | NULL | NULL    | NULL | 1 | NULL |
+----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.07 sec)
```

TRACE SQL and SHOW TRACE

You can run these SQL statements to view the execution results of an SQL statement. Note that you must use TRACE SQL and SHOW TRACE together. The difference between TRACE SQL and EXPLAIN SQL is that TRACE SQL is truly executed.

For example, you can run these statements to view the execution results of the select 1 statement.

```
mysql> trace select 1;
+---+
| 1 |
+---+
| 1 |
+---+
1 row in set (0.03 sec)
mysql> show trace;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| ID | TYPE | GROUP_NAME | DBKEY_NAME          | TIME_COST(MS) | CONNECTION_TIME_COST(MS) | ROWS | STATEMENT | PARAMS |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 0 | Optimize | DRDS       | DRDS                | 3             | 0.00                      | 0 | select 1 | NULL |
| 1 | Query   | TDDL5_00_GROUP | db218249098_sqa_zmf_tddl5_00_3309 | 7             | 0.15                      | 1 | select 1 | NULL |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.01 sec)
```

CHECK TABLE tablename

You can run this SQL statement to check a data table. This SQL statement is mainly used when a table failed to be created by using a data definition language (DDL) statement.

- If the data table is a table shard, this SQL statement allows you to check whether any underlying physical table shard is missing and whether the columns and indexes of the underlying physical table are consistent.
- If the data table is a single-database non-partition table, this SQL statement allows you to check whether this table exists.

```
mysql> check table tddl_mgr_log;
+-----+-----+-----+-----+
| TABLE      | OP | MSG_TYPE | MSG_TEXT |
+-----+-----+-----+-----+
| TDDL5_APP.tddl_mgr_log | check | status | OK   |
+-----+-----+-----+-----+
1 row in set (0.56 sec)
mysql> check table tddl_mgr;
+-----+-----+-----+-----+
| TABLE      | OP | MSG_TYPE | MSG_TEXT |
+-----+-----+-----+-----+
| TDDL5_APP.tddl_mgr | check | Error | Table 'tddl5_00.tddl_mgr' doesn't exist |
+-----+-----+-----+-----+
1 row in set (0.02 sec)
```

SHOW TABLE STATUS [LIKE 'pattern' | WHERE expr]

You can run this SQL statement to obtain the information about a table. This command aggregates the data of all underlying physical table shards.

The following describes the meanings of important columns:

- NAME: indicates the name of the table.
- ENGINE: indicates the storage engine of the table.
- VERSION: indicates the version of the storage engine of the table.
- ROW_FORMAT: indicates the format of the rows in the table. Valid values include Dynamic, Fixed, and Compressed. The value Dynamic indicates that the row length is variable, for example, is a VARCHAR or BLOB field. The value Fixed indicates that the row length is constant, for example, is a CHAR or INTEGER field.
- ROWS: indicates the number of rows in the table.
- AVG_ROW_LENGTH: indicates the average number of bytes in each row.
- DATA_LENGTH: indicates the data volume of the entire table, in bytes.
- MAX_DATA_LENGTH: indicates the maximum volume of data that can be stored in the table.
- INDEX_LENGTH: indicates the size of the disk space occupied by indexes.
- CREATE_TIME: indicates the time when the table was created.
- UPDATE_TIME: indicates the time when the table was last updated.
- COLLATION: indicates the default character set and character sorting rule of the table.
- CREATE_OPTIONS: indicates all the other options specified when the table was created.

```
mysql> show table status like 'multi_db_multi_tbl';
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| NAME      | ENGINE | VERSION | ROW_FORMAT | ROWS | AVG_ROW_LENGTH | DATA_LENGTH | MAX_DATA_LENGTH | INDEX_LENGTH | DATA_FREE | AUTO_INCREMENT | CREATE_TIME      | UPDATE_TIME | CHECK_TIME | COLLATION  | CHECKSUM |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| multi_db_multi_tbl | InnoDB | 10 | Compact | 2 | 16384 | 16384 | 0 | 16384 | 0 | 100000 | 2017-03-27 17:43:57.0 | NULL | NULL | utf8_general_ci | NULL |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.03 sec)
```

The combination of the SHOW TABLE STATUS statement and the PolarDB-X SCAN hint allows you to view the data volume of each physical table shard.

```
mysql> /*! TDDL:SCAN='multi_db_multi_tbl'*/show table status like 'multi_db_multi_tbl';
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Name      | Engine | Version | Row_format | Rows | Avg_row_length | Data_length | Max_data_length | Index_length | D
ata_free | Auto_increment | Create_time   | Update_time | Check_time | Collation   | Checksum | Create_options | Com
ment | Block_format |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| multi_db_multi_tbl_1 | InnoDB | 10 | Compact | 0 | 0 | 16384 | 0 | 16384 | 0 | 1 | 2017-03-27 1
7:43:57 | NULL | NULL | utf8_general_ci | NULL | | | Original |
| multi_db_multi_tbl_0 | InnoDB | 10 | Compact | 0 | 0 | 16384 | 0 | 16384 | 0 | 1 | 2017-03-27 1
7:43:57 | NULL | NULL | utf8_general_ci | NULL | | | Original |
| multi_db_multi_tbl_1 | InnoDB | 10 | Compact | 0 | 0 | 16384 | 0 | 16384 | 0 | 1 | 2017-03-27 1
7:43:57 | NULL | NULL | utf8_general_ci | NULL | | | Original |
| multi_db_multi_tbl_0 | InnoDB | 10 | Compact | 1 | 16384 | 16384 | 0 | 16384 | 0 | 2 | 2017-03-2
7 17:43:57 | NULL | NULL | utf8_general_ci | NULL | | | Original |
| multi_db_multi_tbl_1 | InnoDB | 10 | Compact | 0 | 0 | 16384 | 0 | 16384 | 0 | 1 | 2017-03-27 1
7:43:57 | NULL | NULL | utf8_general_ci | NULL | | | Original |
| multi_db_multi_tbl_0 | InnoDB | 10 | Compact | 0 | 0 | 16384 | 0 | 16384 | 0 | 1 | 2017-03-27 1
7:43:57 | NULL | NULL | utf8_general_ci | NULL | | | Original |
| multi_db_multi_tbl_1 | InnoDB | 10 | Compact | 0 | 0 | 16384 | 0 | 16384 | 0 | 1 | 2017-03-27 1
7:43:57 | NULL | NULL | utf8_general_ci | NULL | | | Original |
| multi_db_multi_tbl_0 | InnoDB | 10 | Compact | 0 | 0 | 16384 | 0 | 16384 | 0 | 1 | 2017-03-27 1
7:43:57 | NULL | NULL | utf8_general_ci | NULL | | | Original |
| multi_db_multi_tbl_1 | InnoDB | 10 | Compact | 0 | 0 | 16384 | 0 | 16384 | 0 | 1 | 2017-03-27 1
7:43:57 | NULL | NULL | utf8_general_ci | NULL | | | Original |
| multi_db_multi_tbl_0 | InnoDB | 10 | Compact | 0 | 0 | 16384 | 0 | 16384 | 0 | 1 | 2017-03-27 1
7:43:57 | NULL | NULL | utf8_general_ci | NULL | | | Original |
| multi_db_multi_tbl_1 | InnoDB | 10 | Compact | 0 | 0 | 16384 | 0 | 16384 | 0 | 1 | 2017-03-27 1
7:43:57 | NULL | NULL | utf8_general_ci | NULL | | | Original |
| multi_db_multi_tbl_0 | InnoDB | 10 | Compact | 0 | 0 | 16384 | 0 | 16384 | 0 | 1 | 2017-03-27 1
7:43:57 | NULL | NULL | utf8_general_ci | NULL | | | Original |
| multi_db_multi_tbl_1 | InnoDB | 10 | Compact | 0 | 0 | 16384 | 0 | 16384 | 0 | 1 | 2017-03-27 1
7:43:57 | NULL | NULL | utf8_general_ci | NULL | | | Original |
| multi_db_multi_tbl_0 | InnoDB | 10 | Compact | 1 | 16384 | 16384 | 0 | 16384 | 0 | 3 | 2017-03-2
7 17:43:57 | NULL | NULL | utf8_general_ci | NULL | | | Original |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
16 rows in set (0.04 sec)
```

11.7.5. Statistics query statements

SHOW [FULL] STATS

You can run this SQL statement to view the overall statistics of a Distributed Relational Database Service (DRDS) instance. The statistics are instantaneous values. Note that the results of `DRDS SHOW FULL STATS` vary with the version of the DRDS instance.

The following describes the meanings of important columns:

- QPS: the number of queries per second (QPS) sent from an application to the DRDS instance. These queries are usually called logical QPS.
- RDS_QPS: the number of QPS sent from the DRDS instance to an ApsaraDB RDS for MySQL instance. These queries are usually called physical QPS.

- **ERROR_PER_SECOND**: the total number of errors that occur on the DRDS instance per second. These errors include various errors such as SQL syntax errors, primary key conflicts, system errors, and connectivity errors.
- **VIOLATION_PER_SECOND**: the number of conflicts that occur on primary keys or unique keys per second.
- **MERGE_QUERY_PER_SECOND**: the number of queries processed on multiple tables through database sharding and table sharding per second.
- **ACTIVE_CONNECTIONS**: the number of active connections to the DRDS instance.
- **CONNECTION_CREATE_PER_SECOND**: the number of connections that are created for the DRDS instance per second.
- **RT(MS)**: the time to respond to an SQL query sent from an application to the DRDS instance. This response time (RT) is usually called logical RT.
- **RDS_RT(MS)**: the time to respond to an SQL query sent from the DRDS instance to an ApsaraDB RDS for MySQL instance. This RT is usually called physical RT.
- **NET_IN(KB/S)**: the amount of inbound traffic of the DRDS instance per second.
- **NET_OUT(KB/S)**: the amount of outbound traffic of the DRDS instance per second.
- **THREAD_RUNNING**: the number of threads that are running in the DRDS instance.
- **HINT_USED_PER_SECOND**: the number of SQL queries that contain hints and are processed by the DRDS instance per second.
- **HINT_USED_COUNT**: the total number of SQL queries that contain hints and have been processed by the DRDS instance since startup.
- **AGGREGATE_QUERY_PER_SECOND**: the number of aggregate SQL queries processed by the DRDS instance per second.
- **AGGREGATE_QUERY_COUNT**: the total number of aggregate SQL queries that have been processed by the DRDS instance.
- **TEMP_TABLE_CREATE_PER_SECOND**: the number of temporary tables created in the DRDS instance per second.
- **TEMP_TABLE_CREATE_COUNT**: the total number of temporary tables that have been created in the DRDS instance since startup.
- **MULTI_DB_JOIN_PER_SECOND**: the number of multi-database JOIN queries processed by the DRDS instance per second.
- **MULTI_DB_JOIN_COUNT**: the number of multi-database JOIN queries that have been processed by the DRDS instance since startup.

```
mysql> show stats;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+
| QPS | RDS_QPS | SLOW_QPS | PHYSICAL_SLOW_QPS | ERROR_PER_SECOND | MERGE_QUERY_PER_SECOND | ACTIVE_CONNECTIONS | RT(MS) | RDS_RT(MS) | NET_IN(KB/S) | NET_OUT(KB/S) | THREAD_RUNNING |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1.77 | 1.68 | 0.03 | 0.03 | 0.02 | 0.00 | 7 | 157.13 | 51.14 | 134.49 | 1.48 | 1 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.01 sec)
mysql> show full stats;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+
| QPS | RDS_QPS | SLOW_QPS | PHYSICAL_SLOW_QPS | ERROR_PER_SECOND | VIOLATION_PER_SECOND | MERGE_QUERY_PER_SECOND | ACTIVE_CONNECTIONS | CONNECTION_CREATE_PER_SECOND | RT(MS) | RDS_RT(MS) | NET_IN(KB/S) | NET_OUT(KB/S) | THREAD_RUNNING | HINT_USED_PER_SECOND | HINT_USED_COUNT | AGGREGATE_QUERY_PER_SECOND | AGGREGATE_QUERY_COUNT | TEMP_TABLE_CREATE_PER_SECOND | TEMP_TABLE_CREATE_COUNT | MULTI_DB_JOIN_PER_SECOND | MULTI_DB_JOIN_COUNT | CPU | FREEMEM | FULLGCCOUNT | FULLGCTIME |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1.63 | 1.68 | 0.03 | 0.03 | 0.02 | 0.00 | 0.00 | 6 | 0.01 | 157.13 | 51.14 | 134.33 | 1.21 | 1 | 0.00 | 54 | 0.00 | 663 | 0.00 | 512 | 0.00 | 516 | 0.09% | 6.96% | 76446 | 21326906 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.01 sec)
```

SHOW DB STATUS

You can run this SQL statement to view the capacity and performance information of a physical database, which is also called a database shard. All the returned values indicate the real-time information. The capacity information is obtained from the ApsaraDB RDS for MySQL system table, and therefore may be different from the actual capacity information.

The following describes the meanings of important columns:

- NAME: the internal tag that represents a logical database corresponding to the database shard. The value is different from the name of the logical database.
- CONNECTION_STRING: the information about a connection from the DRDS instance to the database shard.
- PHYSICAL_DB: the name of the database shard. The TOTAL row indicates the total amount of capacity of all the database shards corresponding to the logical database.
- SIZE_IN_MB: the size of the space occupied by the data in the database shard. Unit: MB
- RATIO: the ratio of the data volume of the database shard to the total data volume of the current logical database.
- THREAD_RUNNING: the number of threads that are running in the ApsaraDB RDS for MySQL instance to which the physical database belongs. The meaning of this parameter is the same as that of the THREAD_RUNNING parameter returned by the MySQL SHOW GLOBAL STATUS command. For more information, see [MySQL Documentation](#).

```
mysql> show db status;
+-----+-----+-----+-----+-----+-----+-----+
| ID | NAME                | CONNECTION_STRING | PHYSICAL_DB | SIZE_IN_MB | RATIO | THREAD_RUNNING |
+-----+-----+-----+-----+-----+-----+-----+
| 1 | drds_db_1516187088365dau | 100.100.64.1:59077 | TOTAL      | 13.109375 | 100% | 3              |
| 2 | drds_db_1516187088365dau | 100.100.64.1:59077 | drds_db_xzip_0000 | 1.578125 | 12.04% |                |
| 3 | drds_db_1516187088365dau | 100.100.64.1:59077 | drds_db_xzip_0001 | 1.4375 | 10.97% |                |
| 4 | drds_db_1516187088365dau | 100.100.64.1:59077 | drds_db_xzip_0002 | 1.4375 | 10.97% |                |
| 5 | drds_db_1516187088365dau | 100.100.64.1:59077 | drds_db_xzip_0003 | 1.4375 | 10.97% |                |
| 6 | drds_db_1516187088365dau | 100.100.64.1:59077 | drds_db_xzip_0004 | 1.734375 | 13.23% |                |
| 7 | drds_db_1516187088365dau | 100.100.64.1:59077 | drds_db_xzip_0005 | 1.734375 | 13.23% |                |
| 8 | drds_db_1516187088365dau | 100.100.64.1:59077 | drds_db_xzip_0006 | 2.015625 | 15.38% |                |
| 9 | drds_db_1516187088365dau | 100.100.64.1:59077 | drds_db_xzip_0007 | 1.734375 | 13.23% |                |
+-----+-----+-----+-----+-----+-----+-----+
```

SHOW FULL DB STATUS [LIKE {tablename}]

You can run this SQL statement to view the capacity and performance information of a table shard in a physical database, which is also called a database shard. All the returned values indicate the real-time information. The capacity information is obtained from the ApsaraDB RDS for MySQL system table, and therefore may be different from the actual capacity information.

The following describes the meanings of important columns:

- **NAME:** the internal tag that represents a logical database corresponding to the database shard. The value is different from the name of the logical database.
- **CONNECTION_STRING:** the information about a connection from the DRDS instance to the database shard.
- **PHYSICAL_DB:** the name of the database shard. If the LIKE keyword is used for filtering in the statement, the TOTAL row indicates the total amount of capacity of the database shard. If the LIKE keyword is not used for filtering in the statement, the TOTAL row indicates the total amount of capacity of all database shards.
- **PHYSICAL_TABLE:** the name of the table shard in the database shard. If the LIKE keyword is used for filtering in the statement, the TOTAL row indicates the total amount of capacity of the table shard. If the LIKE keyword is not used for filtering in the statement, the TOTAL row indicates the total amount of capacity of all table shards in the database shard.
- **SIZE_IN_MB:** the size of the space occupied by the data in the database shard. Unit: MB
- **RATIO:** the ratio of the data volume of the table shard to the total data volume of all the table shards obtained through filtering.
- **THREAD_RUNNING:** the number of threads that are running in the ApsaraDB RDS for MySQL instance to which the physical database belongs. The meaning of this parameter is the same as that of the THREAD_RUNNING parameter returned by the MySQL `SHOW GLOBAL STATUS` command. For more information, see [MySQL Documentation](#).

```
mysql> show full db status like hash_tb;
+-----+-----+-----+-----+-----+-----+-----+
| ID | NAME          | CONNECTION_STRING | PHYSICAL_DB | PHYSICAL_TABLE | SIZE_IN_MB | RATIO | THREAD_RUNNING |
+-----+-----+-----+-----+-----+-----+-----+
| 1 | drds_db_1516187088365dai | 100.100.64.1:59077 | TOTAL      |                | 19.875 | 100% | 3                |
| 2 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0000 | TOTAL      | 3.03125 | 15.25% |                  |
| 3 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0000 | hash_tb_00 | 1.515625 | 7.63% |                  |
| 4 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0000 | hash_tb_01 | 1.515625 | 7.63% |                  |
| 5 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0001 | TOTAL      | 2.0 | 10.06% |                  |
| 6 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0001 | hash_tb_02 | 1.515625 | 7.63% |                  |
| 7 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0001 | hash_tb_03 | 0.484375 | 2.44% |                  |
| 8 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0002 | TOTAL      | 3.03125 | 15.25% |                  |
| 9 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0002 | hash_tb_04 | 1.515625 | 7.63% |                  |
| 10 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0002 | hash_tb_05 | 1.515625 | 7.63% |                  |
| 11 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0003 | TOTAL      | 1.953125 | 9.83% |                  |
| 12 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0003 | hash_tb_06 | 1.515625 | 7.63% |                  |
| 13 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0003 | hash_tb_07 | 0.4375 | 2.2% |                  |
| 14 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0004 | TOTAL      | 3.03125 | 15.25% |                  |
| 15 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0004 | hash_tb_08 | 1.515625 | 7.63% |                  |
| 16 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0004 | hash_tb_09 | 1.515625 | 7.63% |                  |
| 17 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0005 | TOTAL      | 1.921875 | 9.67% |                  |
| 18 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0005 | hash_tb_11 | 1.515625 | 7.63% |                  |
| 19 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0005 | hash_tb_10 | 0.40625 | 2.04% |                  |
| 20 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0006 | TOTAL      | 3.03125 | 15.25% |                  |
| 21 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0006 | hash_tb_12 | 1.515625 | 7.63% |                  |
| 22 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0006 | hash_tb_13 | 1.515625 | 7.63% |                  |
| 23 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0007 | TOTAL      | 1.875 | 9.43% |                  |
| 24 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0007 | hash_tb_14 | 1.515625 | 7.63% |                  |
| 25 | drds_db_1516187088365dai | 100.100.64.1:59077 | drds_db_xzip_0007 | hash_tb_15 | 0.359375 | 1.81% |                  |
+-----+-----+-----+-----+-----+-----+-----+
```

11.7.6. SHOW PROCESSLIST and KILL commands

Note

- If the version of PolarDB-X is 5.1.28-1408022 or later, PolarDB-X support the SHOW PROCESSLIST and KILL commands for both logical and physical connections. For more information, see this topic.
- If the version of PolarDB-X is earlier than 5.1.28-1408022, PolarDB-X support the SHOW PROCESSLIST and KILL commands only for physical connections. For more information, see [SHOW PROCESSLIST and KILL commands in earlier versions](#).

SHOW PROCESSLIST command

In a PolarDB-X instance, you can run the `SHOW PROCESSLIST` command to view information such as connections to the PolarDB-X instance and SQL statements that are being executed in the PolarDB-X instance.

Syntax

```
SHOW [FULL] PROCESSLIST
```

Examples

```
mysql> SHOW PROCESSLIST\G
  ID: 1971050
  USER: admin
  HOST: 111.111.111.111:4303
  DB: drds_test
  COMMAND: Query
  TIME: 0
  STATE:
  INFO: show processlist
1 row in set (0.01 sec)
```

The following describes the meanings of the fields in the result set:

- **ID**: the ID of the connection. The value is a long-type number.
- **USER**: the name of the user who sets up the connection.
- **HOST**: the IP address and port number of the host that sets up the connection.
- **DB**: the name of the database accessed by the connection.
- **COMMAND**: the usage state of the connection. Currently, this field can be set to the following values:
 - **Query**: the current connection is executing an SQL statement.
 - **Sleep**: the current connection is idle.
- **TIME**: the duration when the connection is in the current state:
 - When the value of **COMMAND** is **Query**, this field indicates how long the SQL statement has been being executed on the connection.
 - When the value of **COMMAND** is **Sleep**, this field indicates how long the connection has been in the idle state.
- **STATE**: currently, no meaning has been assigned for this field. The value is constantly empty.
- **INFO**:
 - When the value of **COMMAND** is **Query**, this field indicates the content of the SQL statement that is being executed on the connection. If the **FULL** parameter is not specified, a maximum of the first 30 characters of the SQL statement are returned. If the **FULL** parameter is specified, a maximum of the first 1000 characters of the SQL statement are returned.
 - When the value of **COMMAND** is other values, this field is meaningless and left empty.

SHOW PHYSICAL_PROCESSLIST command

In a PolarDB-X instance, you can run the `SHOW PHYSICAL_PROCESSLIST` command to view information about all the SQL statements that are being executed on underlying ApsaraDB RDS for MySQL instances.

Syntax

```
SHOW [FULL] PHYSICAL_PROCESSLIST
```

When an SQL statement is excessively long, the responses of the `SHOW PHYSICAL_PROCESSLIST` command may be truncated. In this case, you can run the `SHOW FULL PHYSICAL_PROCESSLIST` command to obtain the complete SQL statement.

The meaning of each column in the responses is equivalent to that in the responses of the `SHOW PROCESSLIST` command. For more information, see [SHOW PROCESSLIST Syntax](#).

 **Note** Different from ApsaraDB RDS for MySQL, the PolarDB-X instance returns a string instead of a number in the ID column of a physical connection.

```
mysql> SHOW PHYSICAL_PROCESSLIST\G
***** 1. row *****
  ID: 0-0-521414
  USER: tddl5
  DB: tddl5_00
  COMMAND: Query
  TIME: 0
  STATE: init
  INFO: show processlist
***** 2. row *****
  ID: 0-0-521570
  USER: tddl5
  DB: tddl5_00
  COMMAND: Query
  TIME: 0
  STATE: User sleep
  INFO: /*DRDS /88.88.88.88/b67a0e4d8800000/ */ select sleep(1000)
2 rows in set (0.01 sec)
```

KILL command

The KILL command is used to terminate an SQL statement that is being executed.

The PolarDB-X instance connects to an ApsaraDB RDS for MySQL instance by using the username created by the PolarDB-X instance on the ApsaraDB RDS for MySQL instance. Therefore, if you directly connect to the ApsaraDB RDS for MySQL instance, you do not have the permission to run the KILL command on a request initiated by the PolarDB-X instance.

To terminate an SQL statement that is being executed on the PolarDB-X instance, you must use tools such as the MySQL command line and to connect to the PolarDB-X instance, and then run the KILL command on the PolarDB-X instance.

Syntax

```
KILL PROCESS_ID | 'PHYSICAL_PROCESS_ID' | 'ALL'
```

The KILL command can be used in the following ways:

- Run `KILL PROCESS_ID` to terminate a specified logical SQL statement.

The `PROCESS_ID` parameter is obtained from the ID column in the responses of the `SHOW [FULL] PROCESSLIST` command.

Running the `KILL PROCESS_ID` command in the PolarDB-X instance will terminate both logical and physical SQL statements that are being executed on this connection, and disconnect this connection.

The PolarDB-X instance does not support the `KILL QUERY` command.

- Run `KILL 'PHYSICAL_PROCESS_ID'` to terminate a specified physical SQL statement.

The `PHYSICAL_PROCESS_ID` parameter is obtained from the ID column in the responses of the `SHOW PHYSICAL_PROCESS_ID` command.

Note The `PHYSICAL_PROCESS_ID` column is a string instead of a number. Therefore, the `PHYSICAL_PROCESS_ID` parameter must be enclosed in single quotation marks (') in the KILL command.

Examples:

```
mysql> KILL '0-0-521570';
Query OK, 0 rows affected (0.01 sec)
```

- Run `KILL 'ALL'` to terminate all the physical SQL statements that are executed by the PolarDB-X instance in the

current logical database.

When the underlying ApsaraDB RDS for MySQL instance is overloaded due to some SQL statements, you can use the `KILL 'ALL'` command to terminate all the physical SQL statements that are being executed in the current logical PolarDB-X database.

All physical SQL statements indicated by `PROCESS` that meet the following conditions can be terminated by running `KILL 'ALL'` :

- The value of the `User` parameter for the physical SQL statement indicated by `PROCESS` is a username created by the PolarDB-X instance in the ApsaraDB RDS for MySQL instance.
- The physical SQL statement indicated by `PROCESS` is executing a query. In other words, the value of `COMMAND` is `Query`.

11.7.7. SHOW PROCESSLIST and KILL commands in earlier versions

Note

- If the version of PolarDB-X is 5.1.28-1408022 or later, PolarDB-X supports the `SHOW PROCESSLIST` and `KILL` commands for both logical and physical connections. For more information, see [SHOW PROCESSLIST and KILL commands](#).
- If the version of PolarDB-X is earlier than 5.1.28-1408022, PolarDB-X only supports the `SHOW PROCESSLIST` and `KILL` commands for physical connections. For more information, see this topic.

SHOW PROCESSLIST command

In a PolarDB-X instance, you can run the `SHOW PROCESSLIST` command to view information about all the SQL statements that are being executed on the ApsaraDB RDS for MySQL instances.

Syntax

```
SHOW [FULL] PROCESSLIST
```

When an SQL statement is excessively long, the responses of the `SHOW PROCESSLIST` command may be truncated. In this case, you can run the `SHOW FULL PROCESSLIST` command to obtain the complete SQL statement.

The meaning of each column in the responses is equivalent to that in the responses of the `SHOW PROCESSLIST` statement. For more information, see [SHOW PROCESSLIST Syntax](#).

```
mysql> SHOW PROCESSLIST\G
***** 1. row *****
  ID: 0-0-521414
  USER: tddl5
  DB: tddl5_00
  COMMAND: Query
  TIME: 0
  STATE: init
  INFO: show processlist
  ROWS_SENT: NULL
  ROWS_EXAMINED: NULL
  ROWS_READ: NULL
***** 2. row *****
  ID: 0-0-521570
  USER: tddl5
  DB: tddl5_00
  COMMAND: Query
  TIME: 0
  STATE: User sleep
  INFO: /*DRDS /88.88.88.88/b67a0e4d8800000/ */ select sleep(1000)
  ROWS_SENT: NULL
  ROWS_EXAMINED: NULL
  ROWS_READ: NULL
2 rows in set (0.01 sec)
```

KILL

You can execute the KILL statement to terminate an SQL statement that is being executed.

The PolarDB-X instance connects to an ApsaraDB RDS for MySQL instance by using the username created by the PolarDB-X instance on the ApsaraDB RDS for MySQL instance. Therefore, if you directly connect to the ApsaraDB RDS for MySQL instance, you are not authorized to execute the KILL statement to terminate a request initiated by the PolarDB-X instance.

To terminate an SQL statement that is being executed on the PolarDB-X instance, you must use tools to connect to the PolarDB-X instance. You can use tools such as the MySQL command line and . Then, execute the KILL statement on the PolarDB-X instance.

Syntax

```
KILL 'PROCESS_ID' | 'ALL'
```

The KILL command can be used in the following ways:

- Run `KILL 'PROCESS_ID'` to terminate a specified SQL statement.

The `PROCESS_ID` parameter is obtained from the ID column in the responses of the `SHOW PROCESSLIST` command.

Note Different from ApsaraDB RDS for MySQL, the PolarDB-X instance returns a string instead of a number in the ID column. Therefore, the `PROCESS_ID` parameter must be enclosed in single quotation marks (') in the KILL command.

Examples

```
mysql> KILL '0-0-521570';
Query OK, 0 rows affected (0.01 sec)
```

- Run `KILL 'ALL'` to terminate all the SQL statements executed by the PolarDB-X instance in the current logical database.

When the underlying ApsaraDB RDS for MySQL instance is overloaded due to several SQL statements, you can use the `KILL 'ALL'` command to terminate all the SQL statements that are being executed in the current logical PolarDB-X database.

All SQL statements indicated by `PROCESS` that meet the following conditions can be terminated by running `KILL 'ALL'` :

- The value of the `User` parameter for the physical SQL statement indicated by `PROCESS` is a username created by the PolarDB-X instance in the ApsaraDB RDS for MySQL instance.
- The physical SQL statement indicated by `PROCESS` is executing a query, which means that the value of `COMMAND` is `Query`.

PolarDB-X instances in earlier versions do not support the `KILL 'ALL'` command. An error will be reported if this command is being executed in these instances. To resolve this problem, you can upgrade the version of the PolarDB-X instance.

11.8. Custom hints

 **Note** This topic is applicable to PolarDB-X 5.3 and later. For earlier versions, see [PolarDB-X 5.2 hints](#).

11.8.1. Introduction to hints

As a supplement to the SQL syntax, hints play a critical role in relational databases. They allow you to influence execution plans of SQL statements by using relevant syntax, to specially optimize the SQL statements. PolarDB-X also provides special hint syntax.

For example, if you know the target data is stored in table shards in certain database shards and you need to route the SQL statement directly to the database shards for execution, you can use custom hints provided by PolarDB-X.

```
SELECT /*+TDDL:node('node_name')*/ * FROM table_name;
```

In the preceding SQL statement, the part between `/*` and `*/`, namely, `+TDDL:node('node_name')`, is a PolarDB-X hint. The hint specifies the ApsaraDB RDS for MySQL database shard where the SQL statement is to be executed.

Note

- PolarDB-X hints can be in the formats of `/*+TDDL:hint_command*/` and `/*!+TDDL:hint_command*/`.
- In the MySQL command-line client, if you need to run an SQL statement that contains a PolarDB-X hint in the format of `/*+TDDL:hint_command*/`, add the `-c` parameter to the logon command, because PolarDB-X hints are based on the [MySQL Comment Syntax](#). Otherwise, the client deletes the PolarDB-X hint and then sends the SQL statement to the server for execution, which causes the hint to fail to take effect. For more information, see [MySQL Client Options](#).

PolarDB-X hint syntax

Basic syntax

```
/*+TDDL: hint_command [hint_command ...] */  
/*!+TDDL: hint_command [hint_command ...] */
```

PolarDB-X hints are based on the [MySQL Comment Syntax](#). The hint statements are located between `/*` and `*/` or between `/*!` and `*/`, and must begin with `+TDDL:`. The `hint_command` parameter indicates a PolarDB-X hint command related to the specific operation. Multiple `hint_command` parameters are separated by spaces.

Examples

```
# Query the names of physical tables in each database shard.
/*+TDDL:scan()*/SHOW TABLES;
# Route the query to database shard 0000 of a read-only ApsaraDB RDS for MySQL instance.
/*+TDDL:node(0) slave()*/SELECT * FROM t1;
```

In the example, `/*+TDDL:scan()*/` and `/*+TDDL:node(0) slave()*/` are PolarDB-X hints that begin with `+TDDL: .` The `scan()`, `node(0)`, and `slave()` functions are PolarDB-X hint commands. Hint commands are separated by spaces.

- Use one hint in an SQL statement:

PolarDB-X allows you to use hints in data manipulation language (DML), data definition language (DDL), and data access language (DAL) statements. The following describes the syntax in detail.

- For all statements that support hints, you can specify a hint at the beginning of the statements, for example:

```
/*+TDDL: ... */ SELECT ...
/*+TDDL: ... */ INSERT ...
/*+TDDL: ... */ REPLACE ...
/*+TDDL: ... */ UPDATE ...
/*+TDDL: ... */ DELETE ...
/*+TDDL: ... */ CREATE TABLE ...
/*+TDDL: ... */ ALTER TABLE ...
/*+TDDL: ... */ DROP TABLE ...
/*+TDDL: ... */ SHOW ...
...
```

- For DML statements, you can specify a hint behind the first keyword of the statements, for example:

```
SELECT /*+TDDL: ... */ ...
INSERT /*+TDDL: ... */ ...
REPLACE /*+TDDL: ... */ ...
UPDATE /*+TDDL: ... */ ...
DELETE /*+TDDL: ... */ ...
...
```

Note Different hints may be applicable to different syntaxes. For more information about the applicable syntaxes, see the documentation of hint commands.

- Use multiple hint commands in an SQL statement:

PolarDB-X allows you to use multiple hint commands in SQL statements that contain hints.

```
SELECT /*+TDDL:node(0) slave()*/ ...;
```

PolarDB-X has the following limitations on the use of multiple hint commands:

```
# A single SQL statement cannot contain multiple hint statements.
SELECT /*+TDDL:node(0)*/ /*+TDDL:slave()*/ ...;
# An SQL statement that contains a hint cannot contain duplicate hint commands.
SELECT /*+TDDL:node(0) node(1)*/ ...;
```

PolarDB-X hint classification

PolarDB-X hints are classified into the following major categories according to operation types:

- Read/write splitting
- Specify a timeout period for an SQL statement
- Specify a database shard to run an SQL statement
- Scan all or some of database shards and table shards

11.8.2. Read/write splitting

PolarDB-X provides transparent read/write splitting at the application layer. Data synchronization between primary and read-only ApsaraDB RDS for MySQL instances has a delay of several milliseconds. If you need to read changed data immediately after the primary ApsaraDB RDS for MySQL instance is changed, you must ensure that the SQL statement for reading data is routed to the primary ApsaraDB RDS for MySQL instance. To meet this demand, PolarDB-X provides custom hints for read/write splitting, to route SQL statements to a specified primary or read-only ApsaraDB RDS for MySQL instance.

 **Note** This topic is applicable to PolarDB-X 5.3 and later. For earlier versions, see [Read/write splitting](#).

Syntax

```
/*+TDDL:  
  master()  
  | slave()  
*/
```

With this custom hint, you can specify whether to run an SQL statement on a primary or read-only ApsaraDB RDS for MySQL instance. With the custom hint `/*+TDDL:slave()*/`, if a primary ApsaraDB RDS for MySQL instance is configured with multiple read-only ApsaraDB RDS for MySQL instances, the PolarDB-X instance randomly selects a read-only ApsaraDB RDS for MySQL instance based on its weight, to run the SQL statement.

Note

- PolarDB-X hints can be in the formats of `/*+TDDL:hint_command*/` and `/*!+TDDL:hint_command*/`.
- In the MySQL command-line client, you may need to execute an SQL statement that contains a PolarDB-X hint in the format of `/*+TDDL:hint_command*/`. In this case, add the `-c` parameter to the logon command. Otherwise, the client deletes the PolarDB-X hint before it sends the SQL statement to the server for execution because the PolarDB-X hint is in the format of a [MySQL comment](#). In this case, the hint fails to take effect. For more information, see [MySQL Client Options](#).

Examples

- Specify a primary ApsaraDB RDS for MySQL instance to run an SQL statement:

```
SELECT /*+TDDL:master()*/ * FROM table_name;
```

After the custom hint `/*+TDDL:master()*/` is added behind the first keyword in the SQL statement, this SQL statement is routed to the primary ApsaraDB RDS for MySQL instance for execution.

- Specify a read-only ApsaraDB RDS for MySQL instance to run an SQL statement:

```
SELECT /*+TDDL:slave()*/ * FROM table_name;
```

After the custom hint `/*+TDDL:slave()*/` is added behind the first keyword in the SQL statement, this SQL statement is randomly routed to a read-only ApsaraDB RDS for MySQL instance based on the allocated weight.

Note

- The custom hints for read-write splitting are only applicable to read SQL statements for non-transactional data. SQL statements for transactional data and write SQL statements are still routed to the primary ApsaraDB RDS for MySQL instance for execution.
- The PolarDB-X hint `/*+TDDL:slave()*/` allows you to route the SQL statement randomly to a read-only ApsaraDB RDS for MySQL instance based on the configured weight for execution. If no read-only ApsaraDB RDS for MySQL instance is available, no error is reported. Instead, the primary ApsaraDB RDS for MySQL instance is selected to run the SQL statement.

11.8.3. Specify a timeout period for an SQL statement

In PolarDB-X, the SQL statements for PolarDB-X instances and ApsaraDB RDS for MySQL instances are timed out after 900 seconds (which can be adjusted) by default. However, for some slow SQL statements, the execution duration may exceed 900 seconds. For these slow SQL statements, PolarDB-X provides a custom hint to adjust their timeout periods. You can use this custom hint to adjust the SQL execution duration as needed.

 **Note** This topic is applicable to PolarDB-X 5.3 and later. For earlier versions, see [Specify a timeout period for an SQL statement](#).

Syntax

The syntax of the PolarDB-X hint for specifying a timeout period for an SQL statement is as follows:

```
/*+TDDL:SOCKET_TIMEOUT(time)*/
```

The `SOCKET_TIMEOUT` parameter is measured in milliseconds. With this custom hint, you can adjust the timeout period for the SQL statement based on business requirements.

 **Note**

- PolarDB-X hints can be in the formats of `/*+TDDL:hint_command*/` and `/*!+TDDL:hint_command*/`.
- In the MySQL command-line client, you may need to execute an SQL statement that contains a PolarDB-X hint in the format of `/*+TDDL:hint_command*/`. In this case, add the `-c` parameter to the `logon` command. Otherwise, the client deletes the PolarDB-X hint before it sends the SQL statement to the server for execution because the PolarDB-X hint is in the format of a [MySQL comment](#). In this case, the hint fails to take effect. For more information, see [MySQL Client Options](#).

Examples

Set the timeout period of an SQL statement to 40 seconds:

```
/*+TDDL:SOCKET_TIMEOUT(40000)*/SELECT * FROM t_item;
```

 **Note** A longer timeout period causes database resources to be occupied for a longer period of time. If excessive SQL statements are executed over a long time within the same period, a large number of database resources may be consumed. This will make users unable to use PolarDB-X properly. In this case, we need to use this custom hint to optimize the SQL statements that take a long time to execute.

11.8.4. Specify a database shard to run an SQL statement

When running SQL commands in a PolarDB-X instance, you may find that some SQL statements are not supported by the PolarDB-X instance. In this case, you can use the `NODE HINT` provided by PolarDB-X, to route the SQL statements to one or more database shards for execution. In addition, if you need to query the data in a specified database shard or the data in a specified table shard in a known database shard, you can use the `NODE HINT` to directly route the SQL statement to the database shard for execution.

 **Note** This topic is applicable to PolarDB-X 5.3 and later. For earlier versions, see [Specify a database shard to run an SQL statement](#).

Syntax

The `NODE HINT` allows you to specify a database shard by using a shard name, to run the SQL statement in the database shard. A shard name uniquely identifies a database shard in a PolarDB-X instance. You can run the `SHOW NODE` statement to obtain the shard name.

Specify a database shard by using a shard name, to run an SQL statement

This custom hint allows you to specify one or more database shards to run an SQL statement.

Note If the hint for specifying a database shard is used in an `INSERT` statement that contains a sequence for the target table, the sequence will not take effect. For more information, see [Limits and precautions for sequences](#).

- Specify one database shard to run an SQL statement:

```
/*+TDDL:node('node_name')*/
```

Specifically, `node_name` indicates the shard name. This PolarDB-X hint enables you to route the SQL statement to the database shard specified by `node_name`.

- Specify multiple database shards to run an SQL statement:

```
/*+TDDL:node('node_name',['node_name1','node_name2'])*/
```

You can specify multiple shard names in the parameters and route the SQL statement to multiple database shards for execution. Separate multiple shard names with commas (,).

Note

- When this custom hint is used, the PolarDB-X instance directly routes the SQL statement to the specified database shards for execution. Therefore, the specified shard names in the SQL statement must correspond to existing database shards.
- The `NODE HINT` can be used in data manipulation language (DML), data definition language (DDL), and data access language (DAL) statements.
- PolarDB-X hints can be in the formats of `/*+TDDL:hint_command*/` and `/*!+TDDL:hint_command*/`.
- In the MySQL command-line client, you may need to execute an SQL statement that contains a PolarDB-X hint in the format of `/*+TDDL:hint_command*/`. In this case, add the `-c` parameter to the logon command. Otherwise, the client deletes the PolarDB-X hint before it sends the SQL statement to the server for execution because the PolarDB-X hint is in the format of a [MySQL comment](#). In this case, the hint fails to take effect. For more information, see [MySQL Client Options](#).

Examples

The following shows the responses of the `SHOW NODE` statement for a logical database named `drds_test` in a PolarDB-X instance.

```
mysql> SHOW NODE\G
***** 1. row *****
      ID: 0
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0000_RDS
      MASTER_READ_COUNT: 212
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
***** 2. row *****
      ID: 1
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0001_RDS
      MASTER_READ_COUNT: 29
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
***** 3. row *****
      ID: 2
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0002_RDS
      MASTER_READ_COUNT: 29
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
***** 4. row *****
      ID: 3
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0003_RDS
      MASTER_READ_COUNT: 29
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
***** 5. row *****
      ID: 4
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0004_RDS
      MASTER_READ_COUNT: 29
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
***** 6. row *****
      ID: 5
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0005_RDS
      MASTER_READ_COUNT: 29
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
***** 7. row *****
      ID: 6
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0006_RDS
      MASTER_READ_COUNT: 29
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
***** 8. row *****
      ID: 7
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0007_RDS
      MASTER_READ_COUNT: 29
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
8 rows in set (0.02 sec)
```

As you can see, each database shard has the `NAME` attribute, which indicates the shard name corresponding to the database shard. Each shard name uniquely corresponds to one database shard name. For example, the shard name `DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0003_RDS` corresponds to the database shard name `drds_test_vtla_0003`. Therefore, after obtaining the shard name, you can use the PolarDB-X hint to specify the corresponding database shard to run the SQL statement.

- Specify database shard 0 to run an SQL statement:

```
SELECT /*TDDL:node('DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0000_RDS')*/ * FROM table_name;
```

- Specify multiple database shards to run an SQL statement:

```
SELECT /*TDDL:node('DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0000_RDS','DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0006_RDS')*/ * FROM table_name;
```

This SQL statement will be executed in the database shards whose shard names are `DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0000_RDS` and `DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0006_RDS`.

- View the execution plan of an SQL statement in database shard 0:

```
/*TDDL:node('DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0000_RDS')*/ EXPLAIN SELECT * FROM table_name;
```

11.8.5. Scan all or some of database shards and table shards

In addition to routing an SQL statement to one or more database shards for execution, PolarDB-X provides the `SCAN HINT` to scan all or some of database shards and table shards. With the `SCAN HINT`, you can route an SQL statement to each database shard at a time. For example, you can view all the table shards in a specified database shard or view the data volume of each physical table of a specified logical table.

 **Note** This topic is applicable to PolarDB-X 5.3 and later. For earlier versions, see [Scan all database shards and table shards](#).

With the `SCAN HINT`, you can specify the following SQL execution manners:

- Run an SQL statement in all table shards in all database shards.
- Run an SQL statement in all table shards in a specified database shard.
- Run an SQL statement in the specified table shard in the specified database shard by calculating the name of the physical table based on conditions.
- Run an SQL statement in the specified table shard in the specified database shard by explicitly specifying the name of the physical table.

The `SCAN HINT` can be used in data manipulation language (DML) statements, data definition language (DDL) statements, and some data access language (DAL) statements.

 **Note**

- PolarDB-X hints can be in the formats of `/*+TDDL:hint_command*/` and `/*!+TDDL:hint_command*/`.
- In the MySQL command-line client, you may need to execute an SQL statement that contains a PolarDB-X hint in the format of `/*+TDDL:hint_command*/`. In this case, add the `-c` parameter to the logon command. Otherwise, the client deletes the PolarDB-X hint before it sends the SQL statement to the server for execution because the PolarDB-X hint is in the format of a [MySQL comment](#). In this case, the hint fails to take effect. For more information, see [MySQL Client Options](#).

Syntax

```

# SCAN HINT
# Route an SQL statement to all table shards in all database shards.
SCAN()
# Route an SQL statement to all table shards in a specified database shard.
SCAN(NODE="node_list") # Specify the database shard.
# Route an SQL statement to the specified table shard in the specified database shard by calculating the name of the physical table based on conditions.
SCAN(
  [TABLE="table_name_list" # Specify the name of the logical table.
  , CONDITION="condition_string" # Calculate the names of physical databases based on the content of TABLE and CONDITION.
  [, NODE="node_list"]) # Filter the results obtained based on the content of CONDITION, to retain only the results of the specified physical database.
# Route an SQL statement to the specified table shard in the specified database shard by explicitly specifying the name of the physical table.
SCAN(
  [TABLE="table_name_list" # Specify the name of the logical table.
  , REAL_TABLE=("table_name_list") # Specify the name of the physical table. The same physical table names are applied to all physical databases.
  [, NODE="node_list"]) # Filter the results obtained based on the content of CONDITION, to retain only the results of the specified physical database.
# Specify physical table names or logical table names.
table_name_list:
  table_name [, table_name]...
# Specify physical databases by using GROUP_KEY and GROUP_INDEX, which can be obtained by running the SHOW NODE statement.
node_list:
  {group_key | group_index} [, {group_key | group_index}]...
# Run an SQL WHERE statement. When using this syntax, you must specify conditions for each table, for example, t1.id = 2 and t2.id = 2.
condition_string:
  where_condition

```

Examples

- Run the following SQL statement in all table shards in all database shards:

```
SELECT /*+TDDL:scan()*/ COUNT(1) FROM t1
```

After this statement is executed, the SQL statement is routed to all the physical tables corresponding to the logical table `t1`, and the result sets are merged and returned.

- Run the following SQL statement in all table shards in specified database shards:

```
SELECT /*+TDDL:scan(node='0,1,2')*/ COUNT(1) FROM t1
```

After this statement is executed, all physical tables corresponding to the logical table `t1` in database shards 0000, 0001, and 0002 are calculated, the SQL statement is routed to the physical tables, and the result sets are merged and returned.

- Run the following SQL statement in specified table shards based on conditions:

```
SELECT /*+TDDL:scan('t1', condition='t1.id = 2')*/ COUNT(1) FROM t1
```

After this statement is executed, all physical tables that correspond to the logical table `t1` and meet the conditions are calculated, the SQL statement is routed to the physical tables, and the result sets are merged and returned.

- Run the following SQL JOIN statement in the specified table shards based on conditions:

```
SELECT /*+TDDL:scan('t1, t2', condition='t1.id = 2 and t2.id = 2')*/ FROM t1 a JOIN t2 b ON a.id = b.id WHERE b.name = "test"
```

After this statement is executed, all physical tables that correspond to the logical tables `t1` and `t2` and meet the conditions are calculated, the SQL statement is routed to the physical tables, and the result sets are merged and returned.

 **Notice** Before using this custom hint, you must ensure that the logical tables `t1` and `t2` are partitioned into the same number of database shards and the same number of table shards. Otherwise, the database shards calculated by the PolarDB-X instance based on the conditions are different, and an error will be returned.

- Run the following SQL statement in the specified table shards in database shards by explicitly specifying the names of the physical tables:

```
SELECT /*+TDDL:scan('t1', real_table=('t1_00', 't1_01'))*/ COUNT(1) FROM t1
```

After this statement is executed, the SQL statement is routed to the table shards `t1_00` `t1_01` in all database shards, and the result sets are merged and returned.

- Run the following SQL JOIN statement in the specified table shards in database shards by explicitly specifying the names of the physical tables:

```
SELECT /*+TDDL:scan('t1, t2', real_table=('t1_00,t2_00', 't1_01,t2_01'))*/ FROM t1 a JOIN t2 b ON a.id = b.id WHERE b.name = "test";
```

After this statement is executed, the SQL statement is routed to the table shards `t1_00` , `t2_00` , `t1_01` , and `t2_01` in all database shards, and the result sets are merged and returned.

11.8.6. INDEX HINT

- PolarDB-X supports global secondary indexes. The INDEX hint allows you to obtain query results from a specified GSI.
- The INDEX hint takes effect only for SQL SELECT statements.

 **Note** This custom hint is applicable to only MySQL 5.7 and later and PolarDB-X 5.4.1 and later.

Syntax

```
# FORCE INDEX
tbl_name [[AS] alias] [index_hint]
index_hint:
    FORCE INDEX({index_name})
# INDEX()
/*+TDDL:
    INDEX({table_name | table_alias}, {index_name})
*/
```

PolarDB-X INDEX hint can be used in two ways:

- `FORCE INDEX()` : This syntax is the same as that of [MySQL FORCE INDEX](#).
- `INDEX()` : In this syntax, a global secondary index is specified using a table name (or alias) and an index name. This hint does not take effect in the following cases:
 - The query does not contain the specified table name or alias.
 - The specified global secondary index is not in the specified table.

Note

- PolarDB-X hints can be in the formats of `/*+TDDL:hint_command*/` and `/*!+TDDL:hint_command*/`.
- In the MySQL command-line client, you may need to execute an SQL statement that contains a PolarDB-X hint in the format of `/*+TDDL:hint_command*/`. In this case, add the `-c` parameter to the logon command. Otherwise, the client deletes the PolarDB-X hint before it sends the SQL statement to the server for execution because the PolarDB-X hint is in the format of a [MySQL comment](#). In this case, the hint fails to take effect. For more information, see [MySQL Client Options](#).

Examples

```
CREATE TABLE t_order (
  `id` bigint(11) NOT NULL AUTO_INCREMENT,
  `order_id` varchar(20) DEFAULT NULL,
  `buyer_id` varchar(20) DEFAULT NULL,
  `seller_id` varchar(20) DEFAULT NULL,
  `order_snapshot` longtext DEFAULT NULL,
  `order_detail` longtext DEFAULT NULL,
  PRIMARY KEY (`id`),
  GLOBAL INDEX `g_i_seller` (`seller_id`) dbpartition by hash(`seller_id`),
  UNIQUE GLOBAL INDEX `g_i_buyer` (`buyer_id`) COVERING(`seller_id`, `order_snapshot`)
  dbpartition by hash(`buyer_id`) tpartition by hash(`buyer_id`) tpartitions 3
) ENGINE=InnoDB DEFAULT CHARSET=utf8 dbpartition by hash(`order_id`);
```

Specify the global secondary index `g_i_seller` by using `FORCE INDEX` in the `FROM` clause:

```
SELECT a.*, b.order_id
FROM t_seller a
JOIN t_order b FORCE INDEX(g_i_seller) ON a.seller_id = b.seller_id
WHERE a.seller_nick="abc";
```

Specify the global secondary index `g_i_buyer` by using `INDEX+table alias`:

```
/*+TDDL:index(a, g_i_buyer)*/SELECT * FROM t_order a WHERE a.buyer_id = 123
```

11.9. PolarDB-X 5.2 hints

11.9.1. Introduction to hints

As a supplement to the SQL syntax, hints play a critical role in relational databases. They allow you to affect execution plans of SQL statements by using relevant syntax, to specially optimize the SQL statements.

Overview of PolarDB-X hints

PolarDB-X provides special hint syntax.

For example, if you know the target data is stored in table shards in certain database shards and you need to route the SQL statement directly to the database shards for execution, you can use custom hints provided by PolarDB-X.

```
/*! TDDL:NODE IN('node_name', ...) */SELECT * FROM table_name;
```

In the preceding SQL statement, the part between `/*!` and `*/`, namely, `TDDL:node in('node_name', ...)`, is a PolarDB-X hint. The hint specifies the ApsaraDB RDS for MySQL database shard where the SQL statement is to be executed.

 Note

- PolarDB-X hints can be in the formats of `/*+TDDL:hint_command*/` and `/*!+TDDL:hint_command*/`.
- In the MySQL command-line client, you may need to execute an SQL statement that contains a PolarDB-X hint in the format of `/*+TDDL:hint_command*/`. In this case, add the `-c` parameter to the logon command. Otherwise, the client deletes the PolarDB-X hint before it sends the SQL statement to the server for execution because the PolarDB-X hint is in the format of a [MySQL comment](#). In this case, the hint fails to take effect. For more information, see [MySQL Client Options](#).

PolarDB-X hint syntax

Basic syntax:

```
/*! TDDL:hint command*/
```

PolarDB-X hints are based on [MySQL Comment Syntax](#). Therefore, an SQL statement that contains a PolarDB-X hint is located between `/*!` and `*/`, and must begin with `TDDL:`. The `hint command` indicates a PolarDB-X hint command related to the specific operation. For example, a PolarDB-X hint is added to the following SQL statement to display the name of each database shard.

```
/*! TDDL:SCAN*/SHOW TABLES;
```

In this SQL statement, `/*! TDDL:SCAN*/` is the PolarDB-X hint that begins with `TDDL:`, and `SCAN` is a PolarDB-X hint command.

11.9.2. Read/write splitting

PolarDB-X provides transparent read/write splitting at the application layer. Data synchronization between primary and read-only ApsaraDB RDS for MySQL instances has a delay of several milliseconds. If you need to read changed data immediately after the primary ApsaraDB RDS for MySQL instance is changed, you must ensure that the SQL statement for reading data is routed to the primary ApsaraDB RDS for MySQL instance. To meet this demand, PolarDB-X provides custom hints for read/write splitting, to route SQL statements to a specified primary or read-only ApsaraDB RDS for MySQL instance.

Syntax

```
/*! TDDL:MASTER|SLAVE*/
```

With this custom hint, you can specify whether to run an SQL statement on a primary or read-only ApsaraDB RDS for MySQL instance. With the custom hint `/*!TDDL:SLAVE*/`, if a primary ApsaraDB RDS for MySQL instance is configured with multiple read-only ApsaraDB RDS for MySQL instances, the PolarDB-X instance randomly selects a read-only ApsaraDB RDS for MySQL instance based on its weight, to run the SQL statement.

 Note

- PolarDB-X hints can be in the formats of `/*+TDDL:hint_command*/` and `/*!+TDDL:hint_command*/`.
- In the MySQL command-line client, you may need to execute an SQL statement that contains a PolarDB-X hint in the format of `/*+TDDL:hint_command*/`. In this case, add the `-c` parameter to the logon command. Otherwise, the client deletes the PolarDB-X hint before it sends the SQL statement to the server for execution because the PolarDB-X hint is in the format of a [MySQL comment](#). In this case, the hint fails to take effect. For more information, see [MySQL Client Options](#).

Examples

- Specify a primary ApsaraDB RDS for MySQL instance to run an SQL statement:

```
/*! TDDL:MASTER*/SELECT * FROM table_name;
```

After the custom hint `/*! TDDL:MASTER*/` is added at the beginning of the SQL statement, this SQL statement is routed to the primary ApsaraDB RDS for MySQL instance for execution.

- Specify a read-only ApsaraDB RDS for MySQL instance to run an SQL statement:

```
/*! TDDL:SLAVE*/SELECT * FROM table_name;
```

After the custom hint `/*! TDDL:SLAVE*/` is added at the beginning of the SQL statement, this SQL statement is randomly routed to a read-only ApsaraDB RDS for MySQL instance based on the allocated weight.

Considerations

- The custom hints for read/write splitting are only applicable to read SQL statements for non-transactional data. SQL statements for transactional data and write SQL statements are still routed to the primary ApsaraDB RDS for MySQL instance.
- When you use the `/*+TDDL:slave()*/` hint, the PolarDB-X instance routes the SQL statement randomly to a read-only ApsaraDB RDS for MySQL instance based on the allocated weight. If no read-only ApsaraDB RDS for MySQL instance is available, no error is reported. Instead, the primary ApsaraDB RDS for MySQL instance is selected to execute the SQL statement.

11.9.3. Prevent the delay from a read-only ApsaraDB RDS for MySQL instance

Normally, if you have configured a read-only ApsaraDB for RDS instance for the primary ApsaraDB RDS for MySQL instance of a logical database in a PolarDB-X instance and set read traffic for both the primary and read-only ApsaraDB RDS for MySQL instances, PolarDB-X routes SQL statements to the primary and read-only ApsaraDB RDS for MySQL instances based on the read/write ratio. However, if asynchronous data replication between the primary and read-only ApsaraDB RDS for MySQL instances has a high delay, an error is reported or error results are returned when PolarDB-X routes the SQL statements to the read-only ApsaraDB RDS for MySQL instance.

To address this issue, the PolarDB-X instance provides a custom hint to cut off the delay of the read-only instance. Specifically, based on the maximum delay of primary/secondary replication, PolarDB-X determines whether to route the SQL statement to the primary or the read-only ApsaraDB RDS for MySQL instance.

Syntax

```
/*! TDDL:SQL_DELAY_CUTOFF=time*/
```

With this custom hint, you can specify the value of `SQL_DELAY_CUTOFF`. When the value of `SQL_DELAY` (primary/secondary replication delay of ApsaraDB RDS for MySQL) for the read-only ApsaraDB RDS for MySQL instance reaches or exceeds the value of `time` (which is measured in seconds), the SQL statement is routed to the primary ApsaraDB RDS for MySQL instance.

Note

- PolarDB-X hints can be in the formats of `/*+TDDL:hint_command*/` and `/*!+TDDL:hint_command*/`.
- In the MySQL command-line client, you may need to execute an SQL statement that contains a PolarDB-X hint in the format of `/*+TDDL:hint_command*/`. In this case, add the `-c` parameter to the logon command. Otherwise, the client deletes the PolarDB-X hint before it sends the SQL statement to the server for execution because the PolarDB-X hint is in the format of a [MySQL comment](#). In this case, the hint fails to take effect. For more information, see [MySQL Client Options](#).

Examples

- Set the primary/secondary replication delay to 5 seconds:

```
/*! TDDL:SQL_DELAY_CUTOFF=5*/SELECT * FROM table_name;
```

In this SQL statement, the value of `SQL_DELAY_CUTOFF` is set to 5. Therefore, when the value of `SQL_DELAY` for the read-only ApsaraDB RDS for MySQL instance reaches or exceeds 5 seconds, the SQL statement is routed to the primary ApsaraDB RDS for MySQL instance.

- Use the custom hint for delay cut off with other custom hints:

```
/*! TDDL:SLAVE AND SQL_DELAY_CUTOFF=5*/SELECT * FROM table_name;
```

The custom hint for cutting off the delay of the read-only ApsaraDB RDS for MySQL instance can be used with other hints. By default, the SQL query request is routed to a read-only ApsaraDB RDS for MySQL instance. However, when the primary/secondary replication delay reaches or exceeds 5 seconds, the SQL query request is routed to the primary ApsaraDB RDS for MySQL instance.

11.9.4. Specify a timeout period for an SQL statement

In PolarDB-X, the SQL statements for PolarDB-X instances and ApsaraDB RDS for MySQL instances are timed out after 900 seconds (which can be adjusted) by default. However, for some slow SQL statements, the execution duration may exceed 900 seconds. For these slow SQL statements, PolarDB-X provides a custom hint to adjust their timeout periods. You can use this custom hint to adjust the SQL execution duration as needed.

Syntax

The syntax of the PolarDB-X hint for specifying a timeout period for an SQL statement is as follows:

```
/*! TDDL:SOCKET_TIMEOUT=time*/
```

The `SOCKET_TIMEOUT` parameter is measured in milliseconds. With this custom hint, you can adjust the timeout period for the SQL statement based on business requirements.

Note

- PolarDB-X hints can be in the formats of `/*+TDDL:hint_command*/` and `/*!+TDDL:hint_command*/`.
- In the MySQL command-line client, you may need to execute an SQL statement that contains a PolarDB-X hint in the format of `/*+TDDL:hint_command*/`. In this case, add the `-c` parameter to the logon command. Otherwise, the client deletes the PolarDB-X hint before it sends the SQL statement to the server for execution because the PolarDB-X hint is in the format of a [MySQL comment](#). In this case, the hint fails to take effect. For more information, see [MySQL Client Options](#).

Examples

Set the timeout period of an SQL statement to 40 seconds:

```
/*! TDDL:SOCKET_TIMEOUT=40000*/SELECT * FROM t_item;
```

 **Note** A longer timeout period causes database resources to be occupied for a longer period of time. If excessive SQL statements are executed over a long time within the same period, a large number of database resources may be consumed. This will make users unable to use PolarDB-X properly. In this case, we need to use this custom hint to optimize the SQL statements that take a long time to execute.

11.9.5. Specify a database shard to run an SQL statement

When running SQL commands in a PolarDB-X instance, you may find that some SQL statements are not supported by the PolarDB-X instance. In this case, you can use the custom hint provided by PolarDB-X to route the SQL statements to one or more database shards for execution. In addition, if you need to query the data in a specified database shard or the data in a specified table shard, you can use the custom hint to directly route the SQL statement to the database shard for execution.

Syntax

This custom hint allows you to specify a database shard by using a shard name or the value of the database shard key, to run an SQL statement in the database shard. A shard name uniquely identifies a database shard in a PolarDB-X instance. You can run the `SHOW NODE` command to obtain the shard name.

Note If the hint for specifying a database shard is used in an `INSERT` statement that contains a sequence for the target table, the sequence will not take effect. For more information, see [Limits and precautions for sequences](#).

- Specify a database shard by using a shard name, to run an SQL statement

This custom hint allows you to specify one or more database shards to run an SQL statement.

- Specify one database shard to run an SQL statement:

```
/*! TDDL:NODE='node_name'*/
```

Specifically, `node_name` indicates the shard name. This PolarDB-X hint enables you to route the SQL statement to the database shard specified by `node_name`.

- Specify multiple database shards to run an SQL statement:

```
/*! TDDL:NODE IN ('node_name','node_name1','node_name2')*/
```

The `IN` keyword is used to specify multiple shard names. This custom hint allows you to route the SQL statement to multiple database shards. Separate multiple shard names with commas (,).

Note When this custom hint is used, the PolarDB-X instance directly routes the SQL statement to the specified database shards for execution. Therefore, the specified shard names in the SQL statement must correspond to existing database shards.

- Specify a database shard by using the value of the database shard key, to run an SQL statement

```
/*! TDDL:table_name.partition_key=value [and table_name1.partition_key=value1]*/
```

In this PolarDB-X hint, `table_name` indicates the name of a logical table, and this table is a partitioned table. In addition, `partition_key` indicates a shard key, and `value` indicates the value specified for the shard key. In this custom hint, you can use the `and` keyword to specify the shard keys of multiple partitioned tables. When this PolarDB-X hint is used, the PolarDB-X instance calculates the database shards and table shards where the SQL statement is to be executed, and routes the SQL statement to the corresponding database shards.

 Note

- PolarDB-X hints can be in the formats of `/*+TDDL:hint_command*/` and `/*!+TDDL:hint_command*/`.
- In the MySQL command-line client, you may need to execute an SQL statement that contains a PolarDB-X hint in the format of `/*+TDDL:hint_command*/`. In this case, add the `-c` parameter to the logon command. Otherwise, the client deletes the PolarDB-X hint before it sends the SQL statement to the server for execution because the PolarDB-X hint is in the format of a [MySQL comment](#). In this case, the hint fails to take effect. For more information, see [MySQL Client Options](#).

Examples

The following shows the responses of the `SHOW NODE` statement for a logical database named `drds_test` in a PolarDB-X instance.

```

mysql> SHOW NODE\G
***** 1. row *****
      ID: 0
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0000_RDS
      MASTER_READ_COUNT: 212
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
***** 2. row *****
      ID: 1
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0001_RDS
      MASTER_READ_COUNT: 29
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
***** 3. row *****
      ID: 2
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0002_RDS
      MASTER_READ_COUNT: 29
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
***** 4. row *****
      ID: 3
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0003_RDS
      MASTER_READ_COUNT: 29
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
***** 5. row *****
      ID: 4
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0004_RDS
      MASTER_READ_COUNT: 29
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
***** 6. row *****
      ID: 5
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0005_RDS
      MASTER_READ_COUNT: 29
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
***** 7. row *****
      ID: 6
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0006_RDS
      MASTER_READ_COUNT: 29
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
***** 8. row *****
      ID: 7
      NAME: DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0007_RDS
      MASTER_READ_COUNT: 29
      SLAVE_READ_COUNT: 0
      MASTER_READ_PERCENT: 100%
      SLAVE_READ_PERCENT: 0%
8 rows in set (0.02 sec)

```

As you can see, each database shard has the `NAME` attribute, which indicates the shard name corresponding to the database shard. Each shard name uniquely corresponds to one database shard name. For example, the shard name `DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0003_RDS` corresponds to the database shard name `drds_test_vtla_0003`. Therefore, after obtaining the shard name, you can use the PolarDB-X hint to specify the corresponding database shard to run the SQL statement.

- Specify database shard 0 to run an SQL statement:

```
#!/TDDL:NODE='DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0000_RDS'*/SELECT * FROM table_name;
```

- Specify multiple database shards to run an SQL statement:

```
#!/TDDL:NODE IN('DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0000_RDS','DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0006_RDS')*/SELECT * FROM table_name;
```

This SQL statement will be executed in the database shards whose shard names are `DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0000_RDS` and `DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0006_RDS`.

- View the execution plan of an SQL statement in a specified database shard:

```
#!/TDDL:NODE='DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0000_RDS'*/EXPLAIN SELECT * FROM table_name;
```

After this SQL statement is executed, the execution plan of the `SELECT` statement in the database shard corresponding to the shard name `DRDS_TEST_1473471355140LRPRDRDS_TEST_VTLA_0000_RDS` will be returned.

- Specify a database shard by using the value of the database shard key, to run an SQL statement:

PolarDB-X does not support subqueries in the `SET` clause of an `UPDATE` statement, because a shard key must be specified for `UPDATE` statements in PolarDB-X. To address this issue, PolarDB-X provides a custom hint to route the statement to a database shard for execution.

For example, the following shows the `CREATE TABLE` statement for creating two logical tables `t1` and `t2`, which are partitioned into table shards in database shards:

```
CREATE TABLE `t1` (
  `id` bigint(20) NOT NULL,
  `name` varchar(20) NOT NULL,
  `val` varchar(20) DEFAULT NULL,
  PRIMARY KEY (`id`)
) ENGINE=InnoDB DEFAULT CHARSET=utf8 dbpartition by hash(`id`) tpartition by hash(`name`) tpartitions 3
CREATE TABLE `t2` (
  `id` bigint(20) NOT NULL,
  `name` varchar(20) NOT NULL,
  `val` varchar(20) DEFAULT NULL,
  PRIMARY KEY (`id`)
) ENGINE=InnoDB DEFAULT CHARSET=utf8 dbpartition by hash(`id`) tpartition by hash(`name`) tpartitions 3
```

The following SQL statement is to be executed for the two tables:

```
UPDATE t1 SET val=(SELECT val FROM t2 WHERE id=1) WHERE id=1;
```

If this statement is directly executed in a PolarDB-X instance, an error will be returned indicating that this statement is not supported. In this case, you can add the PolarDB-X hint to this SQL statement before submitting it to the PolarDB-X instance for execution. The SQL statements are as follows:

```
#!/TDDL:t1.id=1 and t2.id=1*/UPDATE t1 SET val=(SELECT val FROM t2 WHERE id=1) WHERE id=1;
```

This statement will be routed to database shards of `t1`, with the `id` of the database shards being 1. You can run the following `EXPLAIN` command to view the execution plan of this SQL statement:

```
mysql> explain /*! TDDL:t1.id=1 and t2.id=1*/UPDATE t1 SET val=(SELECT val FROM t2 WHERE id=1) WHERE id=1\G
***** 1. row *****
GROUP_NAME: TEST_DRDS_1485327111630IXLWTEST_DRDS_IGHF_0001_RDS
SQL: UPDATE `t1_2` AS `t1` SET `val` = (SELECT val FROM `t2_2` AS `t2` WHERE `id` = 1) WHERE `id` = 1
PARAMS: {}
***** 2. row *****
GROUP_NAME: TEST_DRDS_1485327111630IXLWTEST_DRDS_IGHF_0001_RDS
SQL: UPDATE `t1_1` AS `t1` SET `val` = (SELECT val FROM `t2_1` AS `t2` WHERE `id` = 1) WHERE `id` = 1
PARAMS: {}
***** 3. row *****
GROUP_NAME: TEST_DRDS_1485327111630IXLWTEST_DRDS_IGHF_0001_RDS
SQL: UPDATE `t1_0` AS `t1` SET `val` = (SELECT val FROM `t2_0` AS `t2` WHERE `id` = 1) WHERE `id` = 1
PARAMS: {}
3 rows in set (0.00 sec)
```

According to the result set of the `EXPLAIN` command, the SQL statement is rewritten into three statements, which are then routed to the database shards for execution. You can further specify a table shard by using the value of the table shard key, to narrow the execution scope of the SQL statement to a specified table shard.

```
mysql> explain /*! TDDL:t1.id=1 and t2.id=1 and t1.name='1'*/UPDATE t1 SET val=(SELECT val FROM t2 WHERE id=1) WHERE id=1\G
***** 1. row *****
GROUP_NAME: TEST_DRDS_1485327111630IXLWTEST_DRDS_IGHF_0001_RDS
SQL: UPDATE `t1_1` AS `t1` SET `val` = (SELECT val FROM `t2_1` AS `t2` WHERE `id` = 1) WHERE `id` = 1
PARAMS: {}
1 row in set (0.00 sec)
```

Note Before using this custom hint, you must ensure that the logical tables `t1` and `t2` are partitioned into the same number of database shards and the same number of table shards. Otherwise, the database shards calculated by the PolarDB-X instance based on the conditions are different, and an error will be returned.

11.9.6. Scan all database shards and table shards

In addition to routing an SQL statement to one or more database shards for execution, PolarDB-X provides a custom hint to allow you to scan all database shards and table shards. With this custom hint, you can route an SQL statement to each database shard at a time. For example, you can use this custom hint to view all the table shards in a specified database shard. In addition, you can use this custom hint to view the data volume of table shards in each database shard corresponding to a specified logical table.

Syntax

With this PolarDB-X hint, you can route an SQL statement to all database shards for execution and route an SQL statement to all database shards to perform an operation on a specified logical table.

- Route an SQL statement to all database shards for execution:

```
/*! TDDL:SCAN*/
```

- Perform an operation on a specified logical table:

```
/*! TDDL:SCAN='table_name'*/
```

The `table_name` parameter indicates the name of a logical table in the logical database of a PolarDB-X instance. This custom hint is provided for table shards in database shards. Ensure that the value of `table_name` is the name of a table shard in database shards.

 Note

- PolarDB-X hints can be in the formats of `/*+TDDL:hint_command*/` and `/*!+TDDL:hint_command*/`.
- In the MySQL command-line client, you may need to execute an SQL statement that contains a PolarDB-X hint in the format of `/*+TDDL:hint_command*/`. In this case, add the `-c` parameter to the logon command. Otherwise, the client deletes the PolarDB-X hint before it sends the SQL statement to the server for execution because the PolarDB-X hint is in the format of a [MySQL comment](#). In this case, the hint fails to take effect. For more information, see [MySQL Client Options](#).

Examples

- View the data volume of a specified broadcast table in each database shard:

```
/*! TDDL:SCAN*/SELECT COUNT(1) FROM table_name
```

In this SQL statement, `table_name` indicates a broadcast table. This hint causes the PolarDB-X instance to route the SQL statement to each database shard for execution. Therefore, the result sets include the total data volume of the broadcast table `table_name` in all database shards. This statement allows you to conveniently check whether the data of a broadcast table is normal.

- Scan a single-database non-partition logical table:

```
/*! TDDL:SCAN*/SELECT COUNT(1) FROM table_name
```

This hint causes the PolarDB-X instance to route the SQL `select count(1) from table_name` statement to each database shard for execution. The `table_name` parameter indicates a logical table in a logical database of a PolarDB-X instance. Before using this hint, ensure that each database shard contains the table shard `table_name`. In other words, the table shard `table_name` is a logical table that is only partitioned into database shards, but not partitioned into table shards. Otherwise, an error that indicates that the table is not found will be returned.

- Scan a partitioned logical table in database shards:

```
/*! TDDL:SCAN='table_name'*/SELECT COUNT(1) FROM table_name
```

When executing this statement, the PolarDB-X instance first calculates all the database shards and table shards corresponding to the logical table `table_name`, and then generates a COUNT clause for each table shard in each database shard.

- View the execution plans of all database shards:

```
/*! TDDL:SCAN='table_name'*/EXPLAIN SELECT * FROM table_name;
```

11.10. Distributed transactions

11.10.1. Distributed transactions based on MySQL 5.7

 Note

- When you use MySQL 5.7 or later and PolarDB-X 5.3.4 or later, XA distributed transactions are automatically enabled. The user experience of the XA distributed transactions is the same as that of single-database transactions in MySQL. No special commands are required to enable XA distributed transactions.
- When you use MySQL and a PolarDB-X instance in other versions, see [Distributed transactions based on MySQL 5.6](#).

How it works

When you use MySQL 5.7 or later, the PolarDB-X instance processes distributed transactions based on the XA protocol by default.

Use method

The user experience of distributed transactions in a PolarDB-X instance is the same as that of single-database transactions in MySQL, for example, in terms of the following commands:

- `SET AUTOCOMMIT=0` : Start a transaction.
- `COMMIT` : Commit the current transaction.
- `ROLLBACK` : Roll back the current transaction.

If the SQL statement in a transaction involves only a single shard, the PolarDB-X instance routes the transaction directly to the ApsaraDB RDS for MySQL instance as a single-database transaction. If the SQL statement in the transaction is to modify the data of multiple shards, the PolarDB-X instance automatically upgrades the current transaction to a distributed transaction.

11.10.2. Distributed transactions based on MySQL 5.6

How it works

The XA protocol for MySQL 5.6 is not mature. Therefore, the PolarDB-X instance independently implements two-phase commit (2PC) transaction policies for distributed transactions. When you use MySQL 5.7 or later, we recommend that you use XA transaction policies.

 **Note** The distributed transactions described in this topic are intended for users who use MySQL 5.6 or PolarDB-X earlier than 5.3.4. When you use MySQL 5.7 or later and a PolarDB-X instance in 5.3.4 or later, see [Distributed transactions based on MySQL 5.7](#).

Use method

If a transaction involves multiple database shards, you must declare the current transaction as a distributed transaction. If a transaction involves only a single database shard, you do not need to enable distributed transactions, but can process the transaction as a single-database transaction in MySQL. No additional operations are required.

To enable distributed transactions, do as follows:

After transactions are enabled, run `SET drds_transaction_policy = '...'`.

To enable 2PC transactions in the MySQL command-line client, run the following statements:

```
SET AUTOCOMMIT=0;
SET drds_transaction_policy = '2PC'; -- We recommend that you use MySQL 5.6 to run this command.
.... -- Here, you can run your business SQL statement.
COMMIT; -- You can alternatively write ROLLBACK.
```

To enable 2PC transactions by using the Java database connectivity (JDBC) API, write the code as follows:

```
conn.setAutoCommit(false);
try (Statement stmt = conn.createStatement()) {
    stmt.execute("SET drds_transaction_policy = '2PC'");
}
// ... Here, you can run your business SQL statement.
conn.commit(); // You can alternatively write rollback().
```

FAQ

Q: How can I use PolarDB-X distributed transactions in the Spring framework?

A: If you enable transactions by using the Spring `@Transactional` annotation, you can enable PolarDB-X distributed transactions by extending the transaction manager.

Sample code:

```
import org.springframework.jdbc.datasource.DataSourceTransactionManager;
import org.springframework.transaction.TransactionDefinition;
import javax.sql.DataSource;
import java.sql.Connection;
import java.sql.SQLException;
import java.sql.Statement;
public class DrdsTransactionManager extends DataSourceTransactionManager {
    public DrdsTransactionManager(DataSource dataSource) {
        super(dataSource);
    }
    @Override
    protected void prepareTransactionalConnection(Connection con, TransactionDefinition definition) throws SQLException {
        try (Statement stmt = con.createStatement()) {
            stmt.executeUpdate("SET drds_transaction_policy = '2PC'"); // A 2PC transaction is used as an example.
        }
    }
}
```

After that, instantiate the preceding class in the Spring configuration. For example, you can write the code as follows:

```
<bean id="drdsTransactionManager" class="my.app.DrdsTransactionManager">
    <property name="dataSource" ref="yourDataSource" />
</bean>
```

To enable PolarDB-X distributed transactions for a class, you can add the `@Transactional("drdsTransactionManager")` annotation.

11.11. DDL operations

11.11.1. DDL statements

The data definition language (DDL) statement `CREATE TABLE` in a Distributed Relational Database Service (DRDS) instance is similar to that in a MySQL database, and is extended based on the syntax in a MySQL database. To create a table shard in a DRDS instance, you must specify the table sharding manner and the database sharding manner in the `drds_partition_options` parameter. The valid values include `DBPARTITION BY`, `TBPARTITION BY`, `TBPARTITIONS`, and `BROADCAST`.

Currently, you can run a DDL statement in the following ways:

- Run the DDL statement through the MySQL command-line client, for example, by using MySQL command lines, Navicat, or MySQL Workbench.
- Connect to the specified DRDS instance by using program code and then call the DDL statement for execution.

For the syntax of the `CREATE TABLE` statement in a MySQL database, see [MySQL CREATE TABLE Statement](#).

11.11.2. CREATE TABLE statement

11.11.2.1. Overview

This topic describes the syntax, clauses, parameters, and basic methods for creating a table by using a data definition language (DDL) statement.

Note PolarDB-X instances do not allow you to directly create a database by using a DDL statement. To create a database, you can [Log on to the PolarDB-X console](#). For the information about how to create a database, see [Create a database](#).

Syntax

```
CREATE [TEMPORARY] TABLE [IF NOT EXISTS] tbl_name
  (create_definition,...)
  [table_options]
  [drds_partition_options]
  [partition_options]
drds_partition_options:
  DBPARTITION BY
  HASH([column])
  [TBPARTITION BY
  { HASH(column)
  | {MM|DD|WEEK|MMDD}(column)}
  [TBPARTITIONS num]
  ]
```

Clauses and parameters for database and table sharding

- DBPARTITION BY hash(partition_key)** : This parameter specifies the shard key and the sharding algorithm for database sharding. Database sharding by time is not supported.
- TBPARTITION BY { HASH(column) | {MM|DD|WEEK|MMDD}(column)** : (Optional) This parameter specifies the method of mapping data to a physical table. The value is the same as that of DBPARTITION BY by default.
- TBPARTITIONS num** : (Optional) This parameter specifies the number of physical tables to be created in each database shard. The default value is 1. If no table sharding is required, you do not need to specify this parameter.

11.11.2.2. Create a single-database non-partition table

This topic describes how to create a single-database non-partition table.

Create a single-database non-partition table

```
CREATE TABLE single_tbl(
  id int,
  name varchar(30),
  primary key(id)
);
```

According to the node topology of the logical table, you can see that a single-database non-partition logical table is created in database 0.

```
mysql> show topology from single_tbl;
+-----+-----+-----+
| ID | GROUP_NAME | TABLE_NAME |
+-----+-----+-----+
| 0 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | single_tbl |
+-----+-----+-----+
1 row in set (0.01 sec)
```

Specify parameters

You can also specify the `select_statement` parameter when creating a single-database non-partition table. If you need to create table shards, you cannot specify this parameter.

```
CREATE [TEMPORARY] TABLE [IF NOT EXISTS] tbl_name  
  [(create_definition,...)]  
  [table_options]  
  [partition_options]  
  select_statement
```

For example, you can run the following statement to create a single-database non-partition table `single_tbl2` to store the data from the `single_tbl` table. In this case, no sharding is required.

```
CREATE TABLE single_tbl2(  
  id int,  
  name varchar(30),  
  primary key(id)  
) select * from single_tbl;
```

11.11.2.3. Create a non-partition table in database shards

This topic describes how to create a non-partition table in database shards.

Assume that eight database shards have been created. You can run the following command to create a non-partition table in the database shards by calculating the hash function based on the `userId` shard key.

```
CREATE TABLE multi_db_single_tbl(  
  id int,  
  name varchar(30),  
  primary key(id)  
) dbpartition by hash(id);
```

According to the node topology of the logical table, you can see that a table shard is created in each database shard. In other words, the table is only distributed to database shards.

```
mysql> show topology from multi_db_single_tbl;  
+-----+-----+-----+-----+  
| ID | GROUP_NAME | TABLE_NAME |  
+-----+-----+-----+-----+  
| 0 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | multi_db_single_tbl |  
| 1 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | multi_db_single_tbl |  
| 2 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0002_RDS | multi_db_single_tbl |  
| 3 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0003_RDS | multi_db_single_tbl |  
| 4 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0004_RDS | multi_db_single_tbl |  
| 5 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0005_RDS | multi_db_single_tbl |  
| 6 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0006_RDS | multi_db_single_tbl |  
| 7 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | multi_db_single_tbl |  
+-----+-----+-----+-----+  
8 rows in set (0.01 sec)
```

11.11.2.4. Create table shards in database shards

This topic describes how to create table shards in database shards in different sharding manners.

- Use HASH for sharding

- Use RANGE_HASH for sharding
- Use date functions for sharding

In the following examples, it is assumed that eight database shards have been created.

Use HASH for sharding

Create a table that is split into table shards in database shards, with each database shard containing three physical tables. The database sharding process is calculating the hash function by using id as the shard key, and the table sharding process is calculating the hash function by using bid as the shard key. Specifically, a hash operation is performed on the data of the table based on the id column, to distribute the data to multiple database shards. Then, a hash operation is performed on the data in each database shard based on the bid column, to distribute the data to the three physical tables.

```
CREATE TABLE multi_db_multi_tbl(
  id int auto_increment,
  bid int,
  name varchar(30),
  primary key(id)
) dbpartition by hash(id) tpartition by hash(bid) tpartitions 3;
```

According to the node topology of the logical table, you can see that three table shards are created in each database shard.

```
mysql> show topology from multi_db_multi_tbl;
+-----+-----+-----+-----+
| ID | GROUP_NAME | TABLE_NAME |
+-----+-----+-----+
| 0 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | multi_db_multi_tbl_00 |
| 1 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | multi_db_multi_tbl_01 |
| 2 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | multi_db_multi_tbl_02 |
| 3 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | multi_db_multi_tbl_03 |
| 4 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | multi_db_multi_tbl_04 |
| 5 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | multi_db_multi_tbl_05 |
| 6 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0002_RDS | multi_db_multi_tbl_06 |
| 7 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0002_RDS | multi_db_multi_tbl_07 |
| 8 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0002_RDS | multi_db_multi_tbl_08 |
| 9 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0003_RDS | multi_db_multi_tbl_09 |
| 10 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0003_RDS | multi_db_multi_tbl_10 |
| 11 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0003_RDS | multi_db_multi_tbl_11 |
| 12 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0004_RDS | multi_db_multi_tbl_12 |
| 13 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0004_RDS | multi_db_multi_tbl_13 |
| 14 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0004_RDS | multi_db_multi_tbl_14 |
| 15 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0005_RDS | multi_db_multi_tbl_15 |
| 16 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0005_RDS | multi_db_multi_tbl_16 |
| 17 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0005_RDS | multi_db_multi_tbl_17 |
| 18 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0006_RDS | multi_db_multi_tbl_18 |
| 19 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0006_RDS | multi_db_multi_tbl_19 |
| 20 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0006_RDS | multi_db_multi_tbl_20 |
| 21 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | multi_db_multi_tbl_21 |
| 22 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | multi_db_multi_tbl_22 |
| 23 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | multi_db_multi_tbl_23 |
+-----+-----+-----+
24 rows in set (0.01 sec)
```

According to the sharding rule of the logical table, you can see that both database sharding and table sharding are running the hash function, except that the database shard key is id and the table shard key is bid.

```
mysql> show rule from multi_db_multi_tbl;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| ID | TABLE_NAME | BROADCAST | DB_PARTITION_KEY | DB_PARTITION_POLICY | DB_PARTITION_COUNT | TB_PARTITION_KEY | TB_PARTITION_POLICY | TB_PARTITION_COUNT |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 0 | multi_db_multi_tbl | 0 | id | hash | 8 | bid | hash | 3 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.01 sec)
```

Use RANGE_HASH for sharding

- Requirements

The shard key must be a character or a number.

- Routing method

Calculate a hash value based on the last N digits of any shard key, and then calculate the route by using RANGE_HASH. The number N is the third parameter in the function. For example, during calculation of the RANGE_HASH(COL1, COL2, N) function, COL1 is preferentially selected and then truncated to obtain the last N digits for calculation. If COL1 does not exist, COL2 is selected and truncated for calculation.

- Scenarios

RANGE_HASH is applicable to scenarios where two shard keys are used for sharding but only the values of one shard is used for SQL query. Assume that a DRDS database is partitioned into eight physical databases. Our customer has the following requirements:

- The order table of each service needs to be split into database shards by buyer ID and order ID.
- The query is executed based on either the buyer ID or order ID as the condition.

In this case, you can run the following DDL statement to create the order table:

```
create table test_order_tb (
  id int,
  seller_id varchar(30) DEFAULT NULL,
  order_id varchar(30) DEFAULT NULL,
  buyer_id varchar(30) DEFAULT NULL,
  create_time datetime DEFAULT NULL,
  primary key(id)
) ENGINE=InnoDB DEFAULT CHARSET=utf8 dbpartition by RANGE_HASH(buyer_id, order_id, 10) tpartition by RANGE_HASH(buyer_id, order_id, 10) tpartitions 3;
```

 Note

- Neither of the two shard keys can be modified.
- Data insertion fails if the two shard keys point to different database shards or table shards.

Use date functions for sharding

In addition to using the hash function as the sharding algorithm, you can also use the date functions MM, DD, WEEK, and MMDD as the table sharding algorithms. For more information, see the following examples.

- Create a table and then split the table into table shards in database shards. The database sharding process is calculating the hash function by using userid as the shard key, and the table sharding process is calculating DAY_OF_WEEK through the WEEK(actionDate) function and then splitting the table into table shards based on the actionDate column, with one week counted as seven days.

For example, if the value in the actionDate column is 2017-02-27, which is on Monday, the value obtained by calculating the WEEK(actionDate) function is 2. In this case, the record is stored in table shard 2, because 2 % 7 = 2. This table shard is located in a database shard and is named user_log_2. For another example, if the value in the actionDate column is 2017-02-26, which is on Sunday, the value obtained by calculating the WEEK(actionDate) function is 1. In this case, the record is stored in table shard 1, because 1 % 7 = 1. This table shard is located in a database shard and is named user_log_1.

```
CREATE TABLE user_log(
  userID int,
  name varchar(30),
  operation varchar(30),
  actionDate DATE
) dbpartition by hash(userID) tpartition by WEEK(actionDate) tpartitions 7;
```

According to the node topology of the logical table, you can see that seven table shards are created in each database shard, because one week is counted as seven days in the function. The responses are very long, and therefore are omitted by using an ellipsis (...).

```
mysql> show topology from user_log;
+-----+-----+-----+-----+
| ID | GROUP_NAME | TABLE_NAME |
+-----+-----+-----+-----+
| 0 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log_0 |
| 1 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log_1 |
| 2 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log_2 |
| 3 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log_3 |
| 4 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log_4 |
| 5 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log_5 |
| 6 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log_6 |
| 7 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log_0 |
| 8 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log_1 |
| 9 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log_2 |
| 10 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log_3 |
| 11 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log_4 |
| 12 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log_5 |
| 13 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log_6 |
...
| 49 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log_0 |
| 50 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log_1 |
| 51 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log_2 |
| 52 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log_3 |
| 53 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log_4 |
| 54 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log_5 |
| 55 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log_6 |
+-----+-----+-----+-----+
56 rows in set (0.01 sec)
```

According to the sharding rule of the logical table, you can see that the database sharding process is calculating the hash function by using userID as the shard key, and the table sharding process is calculating the WEEK function by using actionDate as the shard key.

```
mysql> show rule from user_log;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
--+
| ID | TABLE_NAME | BROADCAST | DB_PARTITION_KEY | DB_PARTITION_POLICY | DB_PARTITION_COUNT | TB_PARTITION_KEY | TB_PARTITION_POLICY | TB_PARTITION_COUNT |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
--+
| 0 | user_log | 0 | userId | hash | 8 | actionDate | week | 7 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
--+
1 row in set (0.00 sec)
```

According to the specified database and table shard key parameters, you can see the specific physical table in the specific physical database to which the SQL statement is routed.

View the route of the SQL statement

```
mysql> explain select name from user_log where userId = 1 and actionDate = '2017-02-27'\G
***** 1 row *****
GROUP_NAME: SANGUAN_1490167540907XNDVSANGUAN_BSQT_0001_RDS
SQL: select `user_log`.`name` from `user_log_2` `user_log` where ((`user_log`.`userId` = 1) AND (`user_log`.`actionDate` = '2017-02-27'))
PARAMS: {}
1 row in set (0.01 sec)
```

- Create a table that is split into table shards in database shards. The database sharding process is calculating the hash function by using userId as the shard key, and the table sharding process is calculating MONTH_OF_YEAR through the MM(actionDate) function and then splitting the table into table shards based on the actionDate column, with one year counted as 12 months.

For example, if the value in the actionDate column is 2017-02-27, the value obtained by calculating the MM(actionDate) function is 02. In this case, the record is stored in table shard 02, because 02 % 12 = 02. This table shard is located in a database shard and is named user_log_02. For another example, if the value in the actionDate column is 2016-12-27, the value obtained by calculating the MM(actionDate) function is 12. In this case, the record is stored in table shard 00, because 12 % 12 = 00. This table shard is located in a database shard and is named user_log_00.

```
CREATE TABLE user_log2(
  userId int,
  name varchar(30),
  operation varchar(30),
  actionDate DATE
) dbpartition by hash(userId) tpartition by MM(actionDate) tpartitions 12;
```

According to the node topology of the logical table, you can see that 12 table shards are created in each database shard, because one year is counted as 12 months in the function. The responses are very long, and therefore are omitted by using an ellipsis (...).

```
mysql> show topology from user_log2;
+-----+-----+-----+-----+
| ID | GROUP_NAME          | TABLE_NAME |
+-----+-----+-----+-----+
| 0 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log2_00 |
| 1 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log2_01 |
| 2 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log2_02 |
| 3 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log2_03 |
| 4 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log2_04 |
| 5 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log2_05 |
| 6 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log2_06 |
| 7 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log2_07 |
| 8 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log2_08 |
| 9 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log2_09 |
| 10 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log2_10 |
| 11 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log2_11 |
| 12 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log2_00 |
| 13 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log2_01 |
| 14 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log2_02 |
| 15 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log2_03 |
| 16 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log2_04 |
| 17 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log2_05 |
| 18 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log2_06 |
| 19 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log2_07 |
| 20 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log2_08 |
| 21 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log2_09 |
| 22 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log2_10 |
| 23 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | user_log2_11 |
...
| 84 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log2_00 |
| 85 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log2_01 |
| 86 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log2_02 |
| 87 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log2_03 |
| 88 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log2_04 |
| 89 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log2_05 |
| 90 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log2_06 |
| 91 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log2_07 |
| 92 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log2_08 |
| 93 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log2_09 |
| 94 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log2_10 |
| 95 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log2_11 |
+-----+-----+-----+-----+
96 rows in set (0.02 sec)
```

According to the sharding rule of the logical table, you can see that the database sharding process is calculating the hash function by using userid as the shard key, and the table sharding process is calculating the MM function by using actionDate as the shard key.

```
mysql> show rule from user_log2;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
--+
| ID | TABLE_NAME | BROADCAST | DB_PARTITION_KEY | DB_PARTITION_POLICY | DB_PARTITION_COUNT | TB_PARTITI
ON_KEY | TB_PARTITION_POLICY | TB_PARTITION_COUNT |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
--+
| 0 | user_log2 | 0 | userId | hash | 8 | actionDate | mm | 12 | |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
--+
1 row in set (0.00 sec)
```

- Create a table that is split into table shards in database shards. The database sharding process is calculating the hash function by using userId as the shard key, and the table sharding process is calculating DAY_OF_MONTH through the DD(actionDate) function and then splitting the table into table shards, with one month counted as 31 days.

For example, if the value in the actionDate column is 2017-02-27, the value obtained by calculating the DD(actionDate) function is 27. In this case, the record is stored in table shard 27, because $27 \% 31 = 27$. This table shard is located in a database shard and is named user_log_27.

```
CREATE TABLE user_log3(
  userId int,
  name varchar(30),
  operation varchar(30),
  actionDate DATE
) dbpartition by hash(userId) tpartition by DD(actionDate) tpartitions 31;
```

According to the node topology of the logical table, you can see that 31 table shards are created in each database shard, because one month is counted as 31 days in the function. The responses are very long, and therefore are omitted by using an ellipsis (...).

```
mysql> show topology from user_log3;
+-----+-----+-----+-----+
| ID | GROUP_NAME | TABLE_NAME |
+-----+-----+-----+
| 0 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_00 |
| 1 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_01 |
| 2 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_02 |
| 3 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_03 |
| 4 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_04 |
| 5 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_05 |
| 6 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_06 |
| 7 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_07 |
| 8 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_08 |
| 9 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_09 |
| 10 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_10 |
| 11 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_11 |
| 12 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_12 |
| 13 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_13 |
| 14 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_14 |
| 15 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_15 |
| 16 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_16 |
| 17 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_17 |
| 18 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_18 |
| 19 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_19 |
| 20 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_20 |
| 21 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_21 |
| 22 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_22 |
| 23 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_23 |
| 24 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_24 |
| 25 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_25 |
| 26 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_26 |
| 27 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_27 |
| 28 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_28 |
| 29 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_29 |
| 30 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log3_30 |
...
| 237 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log3_20 |
| 238 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log3_21 |
| 239 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log3_22 |
| 240 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log3_23 |
| 241 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log3_24 |
| 242 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log3_25 |
| 243 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log3_26 |
| 244 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log3_27 |
| 245 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log3_28 |
| 246 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log3_29 |
| 247 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log3_30 |
+-----+-----+-----+-----+
248 rows in set (0.01 sec)
```

According to the sharding rule of the logical table, you can see that the database sharding process is calculating the hash function by using `userId` as the shard key, and the table sharding process is calculating the DD function by using `actionDate` as the shard key.

```
mysql> show rule from user_log3;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
--+
| ID | TABLE_NAME | BROADCAST | DB_PARTITION_KEY | DB_PARTITION_POLICY | DB_PARTITION_COUNT | TB_PARTITI
ON_KEY | TB_PARTITION_POLICY | TB_PARTITION_COUNT |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
--+
| 0 | user_log3 | 0 | userId | hash | 8 | actionDate | dd | 31 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
--+
1 row in set (0.01 sec)
```

- Create a table that is split into table shards in database shards. The database sharding process is calculating the hash function by using userId as the shard key, and the table sharding process is calculating DAY_OF_YEAR % 365 through the MMDD(actionDate) tpartitions 365 function and then splitting the table into 365 physical tables, with one year counted as 365 days.

For example, if the value in the actionDate column is 2017-02-27, the value obtained by calculating the MMDD(actionDate) function is 58. In this case, the record is stored in table shard 58. This table shard is located in a database shard and is named user_log_58.

```
CREATE TABLE user_log4(
  userId int,
  name varchar(30),
  operation varchar(30),
  actionDate DATE
) dbpartition by hash(userId) tpartition by MMDD(actionDate) tpartitions 365;
```

According to the node topology of the logical table, you can see that 365 table shards are created in each database shard, because one year is counted as 365 days in the function. The responses are very long, and therefore are omitted by using an ellipsis (...).

```
mysql> show topology from user_log4;
```

| ID | GROUP_NAME | TABLE_NAME |
|------|--|---------------|
| ... | | |
| 2896 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_341 |
| 2897 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_342 |
| 2898 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_343 |
| 2899 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_344 |
| 2900 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_345 |
| 2901 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_346 |
| 2902 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_347 |
| 2903 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_348 |
| 2904 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_349 |
| 2905 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_350 |
| 2906 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_351 |
| 2907 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_352 |
| 2908 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_353 |
| 2909 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_354 |
| 2910 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_355 |
| 2911 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_356 |
| 2912 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_357 |
| 2913 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_358 |
| 2914 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_359 |
| 2915 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_360 |
| 2916 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_361 |
| 2917 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_362 |
| 2918 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_363 |
| 2919 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log4_364 |

2920 rows in set (0.07 sec)

According to the sharding rule of the logical table, you can see that the database sharding process is calculating the hash function by using userId as the shard key, and the table sharding process is calculating the MMDD function by using actionDate as the shard key.

```
mysql> show rule from user_log4;
```

| ID | TABLE_NAME | BROADCAST | DB_PARTITION_KEY | DB_PARTITION_POLICY | DB_PARTITION_COUNT | TB_PARTITION_KEY | TB_PARTITION_POLICY | TB_PARTITION_COUNT |
|----|------------|-----------|------------------|---------------------|--------------------|------------------|---------------------|--------------------|
| 0 | user_log4 | 0 | userId | hash | 8 | actionDate | mmdd | 365 |

1 row in set (0.02 sec)

- Create a table that is split into table shards in database shards. The database sharding process is calculating the hash function by using userId as the shard key, and the table sharding process is calculating DAY_OF_YEAR % 10 through the MMDD(actionDate) tpartitions 10 function and then splitting the table into 10 physical tables, with one year counted as 365 days.

```
CREATE TABLE user_log5(
  userId int,
  name varchar(30),
  operation varchar(30),
  actionDate DATE
) dbpartition by hash(userId) tpartition by MMDD(actionDate) tpartitions 10;
```

According to the node topology of the logical table, you can see that 10 table shards are created in each database shard, because one year is counted as 365 days in the function and the table data is routed to 10 physical tables. The responses are very long, and therefore are omitted by using an ellipsis (...).

```
mysql> show topology from user_log5;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| ID | GROUP_NAME | TABLE_NAME |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 0 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log5_00 |
| 1 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log5_01 |
| 2 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log5_02 |
| 3 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log5_03 |
| 4 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log5_04 |
| 5 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log5_05 |
| 6 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log5_06 |
| 7 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log5_07 |
| 8 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log5_08 |
| 9 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | user_log5_09 |
...
| 70 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log5_00 |
| 71 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log5_01 |
| 72 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log5_02 |
| 73 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log5_03 |
| 74 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log5_04 |
| 75 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log5_05 |
| 76 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log5_06 |
| 77 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log5_07 |
| 78 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log5_08 |
| 79 | SANGUAN_TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | user_log5_09 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
80 rows in set (0.02 sec)
```

According to the sharding rule of the logical table, you can see that the database sharding process is calculating the hash function by using userId as the shard key, and the table sharding process is calculating the MMDD function by using actionDate as the shard key, and then routing the table data to 10 physical tables.

```
mysql> show rule from user_log5;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
--+
| ID | TABLE_NAME | BROADCAST | DB_PARTITION_KEY | DB_PARTITION_POLICY | DB_PARTITION_COUNT | TB_PARTITION_KEY | TB_PARTITION_POLICY | TB_PARTITION_COUNT |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
--+
| 0 | user_log5 | 0 | userId | hash | 8 | actionDate | mmdd | 10 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
--+
1 row in set (0.01 sec)
```

11.11.2.5. Use the primary key as the shard key

When no shard key is specified for the sharding algorithm, the system uses the primary key as the shard key by default. The following illustrates how to use the primary key as the database shard key and the table shard key.

Use the primary key as the database shard key

```
CREATE TABLE prmkey_tbl(  
  id int,  
  name varchar(30),  
  primary key(id)  
) dbpartition by hash();
```

Use the primary key as the database shard key and the table shard key

```
CREATE TABLE prmkey_multi_tbl(  
  id int,  
  name varchar(30),  
  primary key(id)  
) dbpartition by hash() tpartition by hash() tpartitions 3;
```

11.11.2.6. Create a broadcast table

The BROADCAST clause is used to specify a broadcast table to be created. A broadcast table is replicated to each database shard and data consistency is ensured between the database shards by using a synchronization mechanism with a delay of several seconds. This feature allows you to route a JOIN operation from a Cloud Native Distributed Database PolarDB-X (PolarDB-X) instance to an underlying ApsaraDB RDS for MySQL instance to prevent the JOIN operation from being performed in multiple databases. [Overview](#) describes how to optimize SQL statements by using broadcast tables.

The following is an example statement for creating a broadcast table:

```
CREATE TABLE brd_tbl(  
  id int,  
  name varchar(30),  
  primary key(id)  
) ENGINE=InnoDB DEFAULT CHARSET=utf8 BROADCAST;
```

11.11.2.7. Other attributes of the MySQL CREATE TABLE statement

When creating table shards in database shards, you can also specify other attributes of the table shards in the MySQL CREATE TABLE statement. For example, you can specify other attributes as follows:

```
CREATE TABLE multi_db_multi_tbl(  
  id int,  
  name varchar(30),  
  primary key(id)  
) ENGINE=InnoDB DEFAULT CHARSET=utf8 dbpartition by hash(id) tpartition by hash(id) tpartitions 3;
```

11.11.3. ALTER TABLE statement

The syntax of the ALTER TABLE statement used to modify a table is as follows:

```
ALTER [ONLINE|OFFLINE] [IGNORE] TABLE tbl_name  
  [alter_specification [, alter_specification] ...]  
  [partition_options]
```

In a Distributed Relational Database Service (DRDS) instance, you can use this data definition language (DDL) statement to perform routine DDL operations, such as adding a column, adding an index, and modifying a data definition. For more information about the syntax, see [MySQL CREATE TABLE Statement](#).

 **Note** If you need to modify a table shard, you are not allowed to modify the shard key.

- Add a column:

```
ALTER TABLE user_log
ADD COLUMN idcard varchar(30);
```

- Add an index:

```
ALTER TABLE user_log
ADD INDEX idcard_idx (idcard);
```

- Delete an index:

```
ALTER TABLE user_log
DROP INDEX idcard_idx;
```

- Modify a field:

```
ALTER TABLE user_log
MODIFY COLUMN idcard varchar(40);
```

11.11.4. DROP TABLE statement

The syntax of the DROP TABLE statement used to delete a table is as follows:

```
DROP [TEMPORARY] TABLE [IF EXISTS]
tbl_name [, tbl_name] ...
[RESTRICT | CASCADE]
```

The DROP TABLE statement in a Distributed Relational Database Service (DRDS) instance is the same as the DROP TABLE statement in a MySQL database. After the statement is executed, the system automatically deletes the corresponding physical table. For more information about the syntax, see [MySQL DROP TABLE Statement](#).

For example, you can run the following statement to delete the user_log table:

```
DROP TABLE user_log;
```

11.11.5. FAQ about DDL statements

What can I do if an error occurs during table creation?

Data definition language (DDL) statements in a PolarDB-X instance are processed in a distributed manner. If an error occurs, the structures of all table shards are inconsistent from each other. Therefore, you need to perform manual cleanup.

Perform the following steps:

1. Check the basic error descriptions provided by the PolarDB-X instance, such as syntax errors. If the error message is too long, the system will prompt you to call the SHOW WARNINGS command to view the failure cause of each database shard.
2. Run the SHOW TOPOLOGY command to view the topology of physical tables.

```
SHOW TOPOLOGY FROM multi_db_multi_tbl;
+-----+-----+-----+
| ID | GROUP_NAME | TABLE_NAME |
+-----+-----+-----+
| 0 | corona_qatest_0 | multi_db_multi_tbl_00 |
| 1 | corona_qatest_0 | multi_db_multi_tbl_01 |
| 2 | corona_qatest_0 | multi_db_multi_tbl_02 |
| 3 | corona_qatest_1 | multi_db_multi_tbl_03 |
| 4 | corona_qatest_1 | multi_db_multi_tbl_04 |
| 5 | corona_qatest_1 | multi_db_multi_tbl_05 |
| 6 | corona_qatest_2 | multi_db_multi_tbl_06 |
| 7 | corona_qatest_2 | multi_db_multi_tbl_07 |
| 8 | corona_qatest_2 | multi_db_multi_tbl_08 |
| 9 | corona_qatest_3 | multi_db_multi_tbl_09 |
| 10 | corona_qatest_3 | multi_db_multi_tbl_10 |
| 11 | corona_qatest_3 | multi_db_multi_tbl_11 |
+-----+-----+-----+
12 rows in set (0.21 sec)
```

- Run the `CHECK TABLE tablename` command to check whether the logical table has been created. For example, the following response indicates that a physical table corresponding to the logical table `multi_db_multi_tbl` failed to be created.

```
mysql> check table multi_db_multi_tbl;
+-----+-----+-----+-----+
| TABLE | OP | MSG_TYPE | MSG_TEXT |
+-----+-----+-----+-----+
| andor_mysql_qatest.multi_db_multi_tbl | check | Error | Table 'corona_qatest_0.multi_db_multi_tbl_02' doesn't exist |
+-----+-----+-----+-----+
1 row in set (0.16 sec)
```

- Continue to create or delete the table in idempotent mode to create or delete the remaining physical tables.

```
CREATE TABLE IF NOT EXISTS table1
(id int, name varchar(30), primary key(id))
dbpartition by hash(id);
DROP TABLE IF EXISTS table1;
```

What can I do if I failed to create an index or add a column?

The method for handling the failure in creating an index or adding a column is similar to that for the failure in creating a table. For more information, see [Handle DDL exceptions](#).

11.11.6. DDL functions for sharding

11.11.6.1. Overview

PolarDB-X is a database service that supports both database sharding and table sharding.

Support for PolarDB-X database sharding and table sharding

The following table lists the support for database sharding and table sharding in PolarDB-X data definition language (DDL) sharding functions.

| Sharding function | Description | Support for database sharding | Support for table sharding |
|-------------------|-------------|-------------------------------|----------------------------|
|-------------------|-------------|-------------------------------|----------------------------|

| Sharding function | Description | Support for database sharding | Support for table sharding |
|-------------------|---|-------------------------------|----------------------------|
| HASH | Performs a simple modulus operation. | Yes | Yes |
| UNI_HASH | Performs a simple modulus operation. | Yes | Yes |
| RIGHT_SHIFT | Shifts the value to the right. | Yes | Yes |
| RANGE_HASH | Performs double hashing. | Yes | Yes |
| MM | Performs hashing by month. | No | Yes |
| DD | Performs hashing by date. | No | Yes |
| WEEK | Performs hashing by week. | No | Yes |
| MMDD | Performs hashing by month and date. | No | Yes |
| YYYYMM | Performs hashing by year and month. | Yes | Yes |
| YYYYWEEK | Performs hashing by year and week. | Yes | Yes |
| YYYYDD | Performs hashing by year and date. | Yes | Yes |
| YYYYMM_OPT | Performs optimized hashing by year and month. | Yes | Yes |
| YYYYWEEK_OPT | Performs optimized hashing by year and week. | Yes | Yes |
| YYYYDD_OPT | Performs optimized hashing by year and date. | Yes | Yes |

-  **Note** When using database sharding and table sharding in PolarDB-X, note the following:
- In a PolarDB-X instance, the sharding method of a logical table is defined jointly by a sharding function and a shard key. The sharding function contains the number of shards to be created and the routing algorithm. The shard key also specifies the MySQL data type of the shard key.
 - When the database sharding function is the same as the table sharding function and the database shard key is the same as the table shard key in a PolarDB-X instance, the same sharding method is used for database sharding and table sharding. This allows the PolarDB-X instance to uniquely locate one physical table in a physical database based on the value of the shard key.
 - If the database sharding method and the table sharding method of a logical table are different and an SQL query does not contain both database shard key and table shard key, the PolarDB-X instance scans all database shards or all table shards when processing the SQL query.

Support for data types of PolarDB-X DDL sharding functions

Different PolarDB-X DDL sharding functions support different data types. The following table lists the support for various data types in PolarDB-X sharding functions (✓ indicates supported and × indicates not supported).

Support for data types in PolarDB-X DDL sharding functions

| Sharding function | BIGINT | INT | MEDIUMINT | SMALLINT | TINYINT | VARCHAR | CHAR | DATE | DATETIME | TIMESTAMP | Other types |
|-------------------|--------|-----|-----------|----------|---------|---------|------|------|----------|-----------|-------------|
| HASH | √ | √ | √ | √ | √ | √ | √ | x | x | x | x |
| UNI_HASH | √ | √ | √ | √ | √ | √ | √ | x | x | x | x |
| RANGE_HASH | √ | √ | √ | √ | √ | √ | √ | x | x | x | x |
| RIGHT_SHIFT | √ | √ | √ | √ | √ | x | x | x | x | x | x |
| MM | x | x | x | x | x | x | x | √ | √ | √ | x |
| DD | x | x | x | x | x | x | x | √ | √ | √ | x |
| WEEK | x | x | x | x | x | x | x | √ | √ | √ | x |
| MMDD | x | x | x | x | x | x | x | √ | √ | √ | x |
| YYYYMM | x | x | x | x | x | x | x | √ | √ | √ | x |
| YYYYWEEK | x | x | x | x | x | x | x | √ | √ | √ | x |
| YYYYDD | x | x | x | x | x | x | x | √ | √ | √ | x |
| YYYYMM_OPT | x | x | x | x | x | x | x | √ | √ | √ | x |
| YYYYWEEK_OPT | x | x | x | x | x | x | x | √ | √ | √ | x |
| YYYYDD_OPT | x | x | x | x | x | x | x | √ | √ | √ | x |

Syntax description for PolarDB-X DDL sharding functions

PolarDB-X is compatible with the CREATE TABLE statement in MySQL, and additionally provides the `drds_partition_options` keyword to support database sharding and table sharding:

```
CREATE [TEMPORARY] TABLE [IF NOT EXISTS] tbl_name
    (create_definition,...)
    [table_options]
    [drds_partition_options]
    [partition_options]
CREATE [TEMPORARY] TABLE [IF NOT EXISTS] tbl_name
    [(create_definition,...)]
    [table_options]
    [drds_partition_options]
    [partition_options]
    select_statement
drds_partition_options:
    DBPARTITION BY
        {{HASH|YYYYMM|YYYYWEEK|YYYYDD|YYYYMM_OPT|YYYYWEEK_OPT|YYYYDD_OPT}}{(column)}
    [TBPARTITION BY
        {{HASH|MM|DD|WEEK|MMDD|YYYYMM|YYYYWEEK|YYYYDD|YYYYMM_OPT|YYYYWEEK_OPT|YYYYDD_OPT}}{(column)}
    [TBPARTITIONS num]
    ]
```

11.11.6.2. HASH

Requirements

- The shard key must be an integer or a string.
- This sharding function has no requirements on the version of a Distributed Relational Database Service (DRDS) instance. It supports all DRDS instances by default.

Routing method

When the HASH function is run by using different shard keys for database sharding and table sharding, perform the remainder operation on the value of the database shard key based on the number of database shards. If the value of the shard key is a string, the string is converted to a hash value before route calculation. For example, HASH('8') is equivalent to $8 \% D$, where D indicates the number of database shards.

When the UNI_HASH function is run by using the same shard key for both database sharding and table sharding, perform the remainder operation on the value of the shard key based on the total number of table shards. For example, assume that two database shards are created, each database shard contains four table shards, table shards 0 to 3 are stored in database shard 0, and table shards 4 to 7 are stored in database shard 1. If a key value is 15, the key value 15 is distributed to table shard 7 in database shard 1, because $15 \% (2 \times 4) = 7$.

Scenarios

- HASH is applicable when database sharding is implemented by user ID or order ID.
- HASH is also applicable when the shard key is a string.

Examples

If you need to create a non-partition table in database shards by using the HASH function based on the ID column, you can use the following CREATE TABLE statement:

```
create table test_hash_tb (  
  id int,  
  name varchar(30) DEFAULT NULL,  
  create_time datetime DEFAULT NULL,  
  primary key(id)  
) ENGINE=InnoDB DEFAULT CHARSET=utf8 dbpartition by HASH(ID);
```

Notes

The HASH is a simple modulus operation. The output distribution of the HASH function can be even only when the values in the shard column are evenly distributed.

11.11.6.3. UNI_HASH

Requirements

- The shard key must be an integer or a string.
- The version of the PolarDB-X instance must be 5.1.28-1508068 or later. For more information about the PolarDB-X release notes, see [View the instance version](#).

Routing method

When the UNI_HASH function is used for database sharding, perform a remainder operation on the value of the database shard key based on the number of database shards. If the value of the shard key is a string, the string is converted to a hash value before route calculation. For example, HASH('8') is equivalent to $8 \% D$, where D indicates the number of database shards.

When the UNI_HASH function is run by using the same shard key for both database sharding and table sharding, perform the remainder operation on the value of the database shard key based on the number of database shards first (this step is different from that in the HASH function). Then, the data is evenly distributed to the table shards in the database shard.

Scenarios

- UNI_HASH is applicable when database sharding is implemented by user ID or order ID.
- UNI_HASH is also applicable when the shard key is an integer or a string.
- UNI_HASH can be used when the following conditions are met: Two logical tables need to be partitioned into different numbers of table shards in database shards based on the same shard key. In addition, the two tables

are frequently joined by using a JOIN statement based on the shard key.

Comparison with HASH

When you use the UNI_HASH function to create a non-partition table in database shards, the routing method is the same as that used in the HASH function. Specifically, the route is calculated by performing the remainder operation on the key value of the database shard key based on the number of database shards.

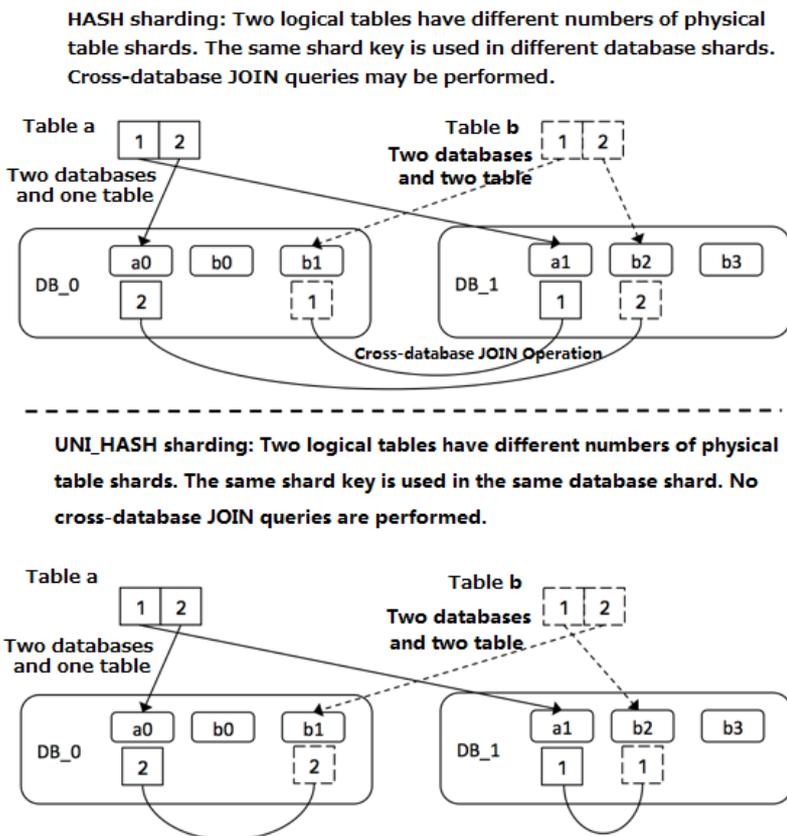
When the UNI_HASH function is run by using the same shard key for both database sharding and table sharding, as the number of table shards changes, the database shard route calculated based on the same key value may also change.

When the UNI_HASH function is run by using the same shard key for both database sharding and table sharding, the database shard route calculated based on the same key value is always the same regardless of the number of table shards.

If two logical tables need to be partitioned into different table shards in database shards based on the same shard key, when the two tables are joined by using the HASH function based on the shard key, multi-database join may occur. However, when the two tables are joined by using the UNI_HASH function based on the shard key, multi-database join does not occur.

Assume that you have two database shards and two logical tables, and each database shard in logical table a stores one table shard and each database shard in logical table b stores two table shards. The following figures separately show the results of a JOIN query for logical tables a and b after the HASH function is used for sharding and the results of a JOIN query for logical tables a and b after the HASH function is used for sharding.

Comparison between HASH and UNI_HASH



Examples

If you need to create four table shards in each database shard by using the UNI_HASH function based on the ID column, you can run the following CREATE TABLE statement:

```
create table test_hash_tb (  
  id int,  
  name varchar(30) DEFAULT NULL,  
  create_time datetime DEFAULT NULL,  
  primary key(id)  
) ENGINE=InnoDB DEFAULT CHARSET=utf8  
dbpartition by UNI_HASH(ID)  
tbpartment by UNI_HASH(ID) tbpartitions 4;
```

Precautions

The UNI_HASH is a simple modulus operation. The output distribution of the UNI_HASH function can be even only when the values in the shard column are evenly distributed.

11.11.6.4. RIGHT_SHIFT

Requirements

- The shard key must be an integer.
- The version of the PolarDB-X instance must be 5.1.28-1320920 or later. For more information, see [View the instance version](#).

Routing method

Shift the value of the database shard key to the right by a specified number of binary digits, and then perform the remainder operation on the obtained integer based on the number of database shards or table shards. In particular, you can specify the number of shifted digits by running a data definition language (DDL) statement.

Scenarios

RIGHT_SHIFT is applicable to improve the evenness of the hash results when the lower-digit parts of most shard key values are very similar to each other but the higher-digit parts vary greatly.

Assume that four shard key values are available: 12340000, 12350000, 12460000, and 12330000. The four lower digits of the four values are all 0000. Directly hashing the values of the shard keys outputs poor results. However, if you run the RIGHT_SHIFT(shardKey, 4) statement to shift the values of the shard keys to the right by four digits, to obtain 1234, 1235, 1246, and 1233, the hashing results are improved.

Examples

If you need to use the ID column as a shard key and shift the values of the ID column to the right by four binary digits to obtain hash values, you can run the following CREATE TABLE statement:

```
create table test_hash_tb (  
  id int,  
  name varchar(30) DEFAULT NULL,  
  create_time datetime DEFAULT NULL,  
  primary key(id)  
) ENGINE=InnoDB DEFAULT CHARSET=utf8  
dbpartition by RIGHT_SHIFT(id, 4)  
tbpartment by RIGHT_SHIFT(id, 4) tbpartitions 2;
```

Precautions

The number of shifted digits cannot exceed the number of digits occupied by the integer.

11.11.6.5. RANGE_HASH

Requirements

- The shard key must be a character or a number.
- The version of the PolarDB-X instance must be 5.1.28-1320920 or later. For more information, see [View the instance version](#).

Routing method

Calculate the hash value based on the last N digits of any shard key and then perform the remainder operation on the hash value based on the number of database shards. This completes the route computing. The number N is the third parameter in the function.

For example, during calculation of the RANGE_HASH(COL1, COL2, N) function, COL1 is preferentially selected and then truncated to obtain the last N digits for calculation. If COL1 does not exist, COL2 is selected and truncated for calculation.

Scenarios

RANGE_HASH is applicable to scenarios where a table needs to be partitioned by two shard keys but query is performed only based on the value of one shard key.

Examples

Assume that a PolarDB-X database is partitioned into eight physical databases. Our customer has the following requirements:

The order table of a business needs to be partitioned into database shards by buyer ID and order ID. The query is executed based on either the buyer ID or order ID as the condition.

In this case, you can run the following DDL statement to create the order table:

```
create table test_order_tb (  
  id int,  
  buyer_id varchar(30) DEFAULT NULL,  
  order_id varchar(30) DEFAULT NULL,  
  create_time datetime DEFAULT NULL,  
  primary key(id)  
) ENGINE=InnoDB DEFAULT CHARSET=utf8  
  dbpartition by RANGE_HASH(buyer_id,order_id, 10)  
  tbpartition by RANGE_HASH (buyer_id,order_id, 10) tbpartitions 3;
```

Precautions

- Neither of the two shard keys can be modified.
- Data insertion fails if the two shard keys point to different database shards or table shards.

11.11.6.6. MM

Requirements

- The shard key must be of the DATE, DATETIME, or TIMESTAMP type.
- MM is only applicable to table sharding.
- The version of the PolarDB-X instance must be 5.1.28-1320920 or later. For more information, see [View the instance version](#).

Routing method

Perform the remainder operation based on the month that corresponds to the time value of the database shard key to obtain the table shard subscript.

Scenarios

MM can be used to partition tables by month. The table shard name indicates a specific month.

Examples

Assume that we need to perform database sharding by ID, perform table sharding for the create_time column by month, and map every month to a physical table. The data definition language (DDL) statement is as follows:

```
create table test_mm_tb (  
  id int,  
  name varchar(30) DEFAULT NULL,  
  create_time datetime DEFAULT NULL,  
  primary key(id)  
) ENGINE=InnoDB DEFAULT CHARSET=utf8  
dbpartition by HASH(id)  
tbpartment by MM(create_time) tpartitions 12;
```

Precautions

When you partition tables with MM, ensure that each database shard has no more than 12 table shards because a year has 12 months.

11.11.6.7. DD

Requirements

- The shard key must be of the DATE, DATETIME, or TIMESTAMP type.
- DD is only applicable to table sharding.
- The version of the PolarDB-X instance must be 5.1.28-1320920 or later. For more information, see [View the instance version](#).

Routing method

Perform the remainder operation based on the day of the month that corresponds to the time value of the database shard key to obtain the table shard subscript.

Scenarios

DD can be used to partition tables based on a specified number of days in a month, that is, a date. The subscript of the table shard name indicates the day in a month. A month has 31 days at most.

Examples

Assume that we need to perform database sharding by ID, perform table sharding for the create_time column by day, and map every day to a physical table. The data definition language (DDL) statement is as follows:

```
create table test_dd_tb (  
  id int,  
  name varchar(30) DEFAULT NULL,  
  create_time datetime DEFAULT NULL,  
  primary key(id)  
) ENGINE=InnoDB DEFAULT CHARSET=utf8  
dbpartition by HASH(id)  
tbpartment by DD(create_time) tpartitions 31;
```

Precautions

When you partition tables with DD, ensure that each database shard has no more than 31 table shards because a month has 31 days at most.

11.11.6.8. WEEK

Requirements

- The shard key must be of the DATE, DATETIME, or TIMESTAMP type.
- WEEK is only applicable to table sharding.
- The version of the PolarDB-X instance must be 5.1.28-1320920 or later. For more information, see [View the instance version](#).

Routing method

Perform the remainder operation based on the day of a week that corresponds to the time value of the database shard key to obtain the table shard subscript.

Scenarios

WEEK can be used to partition tables based on days in a week. The subscript of the table shard name corresponds to each day of a week, from Monday to Sunday.

Examples

Assume that we need to perform database sharding by ID, perform table sharding for the create_time column by week, and map every day of a week (from Monday to Sunday) to a physical table. The data definition language (DDL) statement is as follows:

```
create table test_week_tb (  
  id int,  
  name varchar(30) DEFAULT NULL,  
  create_time datetime DEFAULT NULL,  
  primary key(id)  
) ENGINE=InnoDB DEFAULT CHARSET=utf8  
dbpartition by HASH(name)  
tbpartment by WEEK(create_time) tpartitions 7;
```

Precautions

When you partition tables with WEEK, ensure that each database shard has no more than seven table shards because a week has seven days.

11.11.6.9. MMDD

Requirements

- The shard key must be of the DATE, DATETIME, or TIMESTAMP type.
- MMDD is only applicable to table sharding.
- The version of the PolarDB-X instance must be 5.1.28-1320920 or later. For more information, see [View the instance version](#).

Routing method

Perform the remainder operation based on the number of days in a year that corresponds to the time value of the database shard key to obtain the table sharding subscript.

Scenarios

MMDD can be used to partition tables based on the number of days in a year that corresponds to a date in that year. The subscript of the table shard name indicates the day in that year, with a maximum of 366 days in a year.

Examples

Assume that we need to perform database sharding by ID, create tables for the create_time column by date (month-day), and map every day of a year to a physical table. The data definition language (DDL) statement is as follows.

```
create table test_mmdd_tb (
  id int,
  name varchar(30) DEFAULT NULL,
  create_time datetime DEFAULT NULL,
  primary key(id)
) ENGINE=InnoDB DEFAULT CHARSET=utf8
dbpartition by HASH(name)
tbpartition by MMDD(create_time) tbpartitions 365;
```

Precautions

When you partition tables with MMDD, ensure that each database shard has no more than 366 table shards because a year has 366 days at most.

11.11.6.10. YYYYMM

Requirements

- The shard key must be of the DATE, DATETIME, or TIMESTAMP type.
- The version of the PolarDB-X instance must be 5.1.28-1320920 or later. For more information, see [View the instance version](#).

Routing method

Calculate the hash value based on the year and months of the year in the time value of the database shard key and then perform the remainder operation on the hash value based on the number of database shards. This completes route computing.

For example, YYYYMM('2012-12-31 12:12:12') is equivalent to $(2012 \times 12 + 12) \% D$, where D indicates the number of database shards.

Scenarios

YYYYMM can be used to partition databases by year and month. We recommend that you use YYYYMM with tbpartition YYYYMM(ShardKey).

Assume that a PolarDB-X database is partitioned into eight physical databases. Our customer has the following requirements:

- Perform database sharding for a service by year and month.
- Distribute data from every month within two years to a separate table shard.
- Distribute a query with the database and table shard keys to a physical table shard of a physical database shard.

The preceding requirements can be met by using YYYYMM. For the requirement of distributing data from every month within two years to a table shard (that is, one table shard stores the data of one month), create at least 24 physical table shards because a year has 12 months. Create three physical table shards for each database shard because the PolarDB-X instance contains eight database shards. The data definition language (DDL) statement is as follows.

```
create table test_yyyymm_tb (
  id int,
  name varchar(30) DEFAULT NULL,
  create_time datetime DEFAULT NULL,
  primary key(id)
) ENGINE=InnoDB DEFAULT CHARSET=utf8
dbpartition by YYYYMM(create_time)
tbpartition by YYYYMM(create_time) tbpartitions 3;
```

Precautions

- YYYYMM does not support distributing data from every month in every year to a separate table shard. Instead, the number of table shards must be fixed.
- After a cycle over months (for example, a cycle exists between 2012-03 and 2013-03), data from the same month may be routed to the same database or table shard, depending on the actual number of table shards.

11.11.6.11. YYYYWEEK

Requirements

- The shard key must be of the DATE, DATETIME, or TIMESTAMP type.
- The version of the PolarDB-X instance must be 5.1.28-1320920 or later. For more information, see [View the instance version](#).

Routing method

Calculate the hash value based on the year and weeks of the year in the time value of the database shard key and then perform the remainder operation on the hash value based on the number of database shards. This completes route computing.

For example, YYYYWEEK('2012-12-31 12:12:12') is equivalent to $(2013 \times 52 + 1) \% D$, with the date 2012-12-31 falling on the first week of 2013, where D indicates the number of database shards.

Scenarios

YYYYWEEK can be used to partition databases by year and the number of weeks in a year. We recommend that you use YYYYWEEK with tbpartition YYYYWEEK(ShardKey).

Assume that a PolarDB-X database is partitioned into eight physical databases. Our customer has the following requirements:

- Perform database sharding for a service by year and by week.
- Distribute data from every week within two years to a separate table shard.
- Distribute a query with the database and table shard keys to a physical table shard of a physical database shard.

The preceding requirements can be met by using YYYYWEEK. For the requirement of distributing data from every week within two years to a table shard (that is, one table shard stores the data of one week), create at least 106 physical table shards because a year has roughly 53 weeks (rounded). Create 14 physical table shards for each database shard because the PolarDB-X instance contains eight database shards ($14 \times 8 = 112 > 106$). We recommend that the number of table shards be an integer multiple of the number of database shards. The data definition language (DDL) statement is as follows:

```
create table test_yyyymm_tb (  
  id int,  
  name varchar(30) DEFAULT NULL,  
  create_time datetime DEFAULT NULL,  
  primary key(id)  
) ENGINE=InnoDB DEFAULT CHARSET=utf8  
  dbpartition by YYYYWEEK(create_time)  
  tbpartition by YYYYWEEK(create_time) tbpartitions 14;
```

Precautions

- YYYYWEEK does not support distributing data from every week in every year to a separate table shard. Instead, the number of table shards must be fixed.
- After a cycle over weeks (for example, a cycle exists between the first week of 2012 and the first week of 2013), data from the same week after a cycle may be routed to the same database shard or table shard, depending on the actual number of table shards.

11.11.6.12. YYYYDD

Requirements

- The shard key must be of the DATE, DATETIME, or TIMESTAMP type.
- The version of the PolarDB-X instance must be 5.1.28-1320920 or later. For more information, see [View the instance version](#).

Routing method

Calculate the hash value based on the year and days of the year in the time value of the database shard key and then perform the remainder operation on the hash value based on the number of database shards. This completes route computing.

For example, YYYYDD('2012-12-31 12:12:12') is equivalent to $(2012 \times 366 + 365) \% D$, with 2012-12-31 as the 365th day of 2012, where D indicates the number of database shards.

Scenarios

Database sharding is performed by year and the number of days in a year. We recommend that you use YYYYDD with tbpartition YYYYDD(ShardKey).

Assume that a PolarDB-X database is partitioned into eight physical databases. Our customer has the following requirements:

- Perform database sharding for a service by year and day.
- Distribute data from every week within two years to a separate table shard.
- Distribute a query with the database and table shard keys to a physical table shard of a physical database shard.

The preceding requirements can be met by using YYYYDD. For the requirement of distributing data from every day within two years to a table shard (that is, one table shard stores the data of one day), create at least 732 physical table shards because a year has up to 366 days. Create 92 physical table shards for each database shard because the PolarDB-X instance contains eight database shards ($732/8 = 91.5$, rounded to 92). We recommend that the number of table shards be an integer multiple of the number of database shards. The data definition language (DDL) statement is as follows:

```
create table test_yyyydd_tb (  
  id int,  
  name varchar(30) DEFAULT NULL,  
  create_time datetime DEFAULT NULL,  
  primary key(id)  
) ENGINE=InnoDB DEFAULT CHARSET=utf8  
  dbpartition by YYYYDD(create_time)  
  tbpartition by YYYYDD(create_time) tbpartitions 92;
```

Precautions

- YYYYDD does not support distributing data from every day in every year to a separate table shard. Instead, the number of table shards must be fixed.
- After a cycle of a specific date (for example, a cycle exists between 2012-03-01 and 2013-03-01), data from the same date may be routed to the same database shard or table shard, depending on the actual number of table shards.

11.11.6.13. YYYYMM_OPT

Requirements

- The shard key must be of the DATE, DATETIME, or TIMESTAMP type.
- The year and month of user data increase naturally over time, rather than randomly.
- The version of the PolarDB-X instance must be 5.1.28-1320920 or later. For more information, see [View the instance version](#).

Optimizations

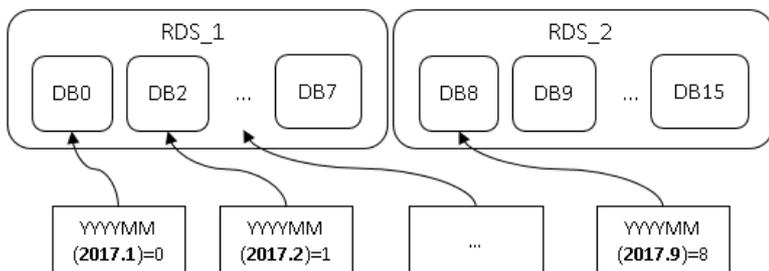
Compared with YYYYMM, YYYYMM_OPT maintains the even distribution of data among ApsaraDB RDS for MySQL instances as the timeline increases.

For example, assume that two ApsaraDB RDS for MySQL instances are attached to a PolarDB-X instance, with 16 database shards. DB0 to DB7 shards are located on one ApsaraDB RDS for MySQL instance, and DB8 to DB15 shards are located on the other ApsaraDB RDS for MySQL instance.

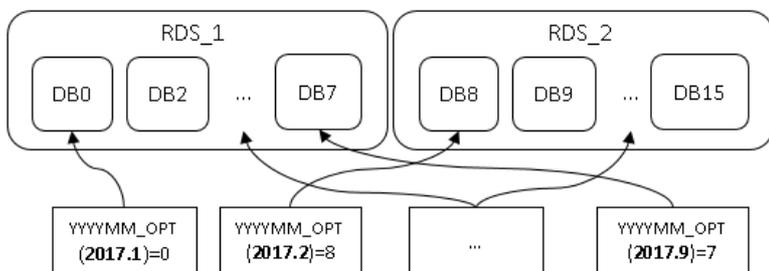
The following figure shows the mappings when YYYYMM and YYYYMM_OPT are used for database sharding, respectively.

Comparison between YYYYMM and YYYYMM_OPT

As the time goes on linearly, YYYYMM fills data in ApsaraDB for RDS instances in sequence (data is first distributed to the database shards of RDS_1, then to the database shards of RDS_2, and then to the database shards of RDS_1 again).



YYYYMM_OPT evenly distributes data between ApsaraDB for RDS instances as the time goes on (data is alternately distributed between RDS_1 and RDS_2, so that the data size of the two RDS instances is balanced).



- YYYYMM_OPT distributes data evenly to each ApsaraDB RDS for MySQL instance, helping to maximize the performance of each ApsaraDB RDS for MySQL instance.
- How to choose between YYYYMM and YYYYMM_OPT :
 - YYYYMM_OPT can be used to distribute data evenly to each ApsaraDB RDS for MySQL instance if the time of service data generation increases sequentially and the data volume does not differ much between the time points.
 - YYYYMM is applicable if the time of data generation increases randomly rather than sequentially.

Routing method

- Calculate the hash value based on the year and months of the year in the time value of the database shard key and then perform the remainder operation on the hash value based on the number of database shards. This completes route computing.
- The hash calculation based on the database and table shard key considers the data distribution among the ApsaraDB RDS for MySQL instances that connect to the PolarDB-X instances.

Scenarios

- Databases and tables need to be partitioned by year and month, respectively.
- Data must be evenly distributed to each ApsaraDB RDS for MySQL instance that connects to the PolarDB-X instance.
- The time of the shard key increases sequentially rather than randomly, and the data volume is relatively average from month to month. For example, the number of monthly journal logs increases every month, and the log data is not concentrated on the same ApsaraDB RDS for MySQL instance.

Precautions

- YYYYMM_OPT does not support distributing data from every month in every year to a separate table shard. Instead, the number of table shards must be fixed.
- After a cycle over months (for example, a cycle exists between 2012-03 and 2013-03), data from the same month may be routed to the same database or table shard, depending on the actual number of table shards.

11.11.6.14. YYYYWEEK_OPT

Requirements

- The shard key must be of the DATE, DATETIME, or TIMESTAMP type.
- The version of the PolarDB-X instance must be 5.1.28-1320920 or later. For more information, see [View the instance version](#).

Optimizations

- Compared with YYYYWEEK, YYYYWEEK_OPT maintains the even distribution of data among ApsaraDB RDS for MySQL instances as the timeline increases. The effect is similar to YYYYMM_OPT.
- YYYYWEEK_OPT distributes data evenly to each ApsaraDB RDS for MySQL instance, helping to maximize the performance of each ApsaraDB RDS for MySQL instance.
- How to choose between YYYYWEEK and YYYYWEEK_OPT:
 - YYYYWEEK_OPT can be used to distribute data evenly to each ApsaraDB RDS for MySQL instance if the time of service data increases sequentially and the data volume does not differ much between time points.
 - YYYYWEEK is applicable if the time of data generation increases randomly rather than sequentially.

Routing method

- Calculate the hash value based on the year and weeks of the year in the time value of the database shard key and then perform the remainder operation on the hash value based on the number of database shards. This completes route computing.
- The hash calculation based on the database and table shard key considers the data distribution among the ApsaraDB RDS for MySQL instances that connect to the PolarDB-X instances.

Scenarios

- Databases and tables are partitioned by year and week, respectively.
- Data must be evenly distributed to each ApsaraDB RDS for MySQL instance that connects to the PolarDB-X instance.
- The time of the shard key increases sequentially rather than randomly, and the data volume is relatively average from week to week. For example, the number of weekly journal logs increases every week, and the log data is not concentrated on the same ApsaraDB RDS for MySQL instance.

Precautions

- YYYYWEEK_OPT does not support distributing data from every week in every year to a separate table shard. Instead, the number of table shards must be fixed.
- After a cycle over weeks (for example, a cycle exists between the first week of 2012 and the first week of 2013), data from the same week after a cycle may be routed to the same database shard or table shard, depending on the actual number of table shards.

11.11.6.15. YYYYDD_OPT

Requirements

- The shard key must be of the DATE, DATETIME, or TIMESTAMP type.
- The version of the PolarDB-X instance must be 5.1.28-1320920 or later. For more information, see [View the instance version](#).

Optimizations

- Compared with YYYYDD, YYYYDD_OPT maintains the even distribution of data among ApsaraDB RDS for MySQL instances as the timeline increases. The effect is similar to YYYYMM_OPT.
- YYYYDD_OPT distributes data evenly to each ApsaraDB RDS for MySQL instance, helping to maximize the performance of each ApsaraDB RDS for MySQL instance.
- How to choose between YYYYDD and YYYYDD_OPT:
 - YYYYDD_OPT can be used to distribute data evenly to each ApsaraDB RDS for MySQL instance if the time of service data generation increases sequentially and the data volume does not differ much between time points.
 - YYYYDD is applicable if the time of data generation increases randomly rather than sequentially.

Routing method

- Calculate the hash value based on the year and days of the year in the time value of the database shard key and then perform the remainder operation on the hash value based on the number of database shards. This completes route computing.
- The hash calculation based on the database and table shard key considers the data distribution among the ApsaraDB RDS for MySQL instances that connect to the PolarDB-X instances.

Scenarios

- Databases and tables need to be partitioned by year and by day, respectively.
- Data must be evenly distributed to each ApsaraDB RDS for MySQL instance that connects to the PolarDB-X instance.
- The time of the shard key increases sequentially rather than randomly, and the data volume is relatively average from day to day. For example, the number of daily journal logs increases every day, and the log data is not concentrated on the same ApsaraDB RDS for MySQL instance.

Precautions

- YYYYDD_OPT does not support distributing data from every day in every year to a separate table shard. Instead, the number of table shards must be fixed.
- After a cycle of a specific date (for example, a cycle exists between 2012-03-01 and 2013-03-01), data from the same date may be routed to the same database shard or table shard, depending on the actual number of table shards.

11.12. Automatic protection of important SQL statements

In Distributed Relational Database Service (DRDS), the data manipulation language (DML) statements are the same as MySQL statements.

We recommend that you include the shard key in the `SELECT` and `UPDATE` statements of DRDS. The `INSERT` statement of DRDS must include the shard key and a non-empty key value.

By default, DRDS disables full-table deletion and updating to avoid misoperation.

The following statements are prohibited by default:

- A `DELETE` statement without the `WHERE` or `LIMIT` condition
- An `UPDATE` statement without the `WHERE` or `LIMIT` condition

If you need to perform full-table deletion or update, you can temporarily skip this limit by using the following hint:

```
HINT: /! TDDL:FORBID_EXECUTE_DML_ALL=false*/
```

Examples

- Full-table deletion is intercepted by default.

```
mysql> delete from tt;
ERR-CODE: [TDDL-4620][ERR_FORBID_EXECUTE_DML_ALL] Forbid execute DELETE ALL or UPDATE ALL sql. More: [http://middleware.alibaba-inc.com/faq/faqByFaqCode.html?faqCode=TDDL-4620]
```

The operation is successful if the following hint is added:

```
mysql> /*! TDDL:FORBID_EXECUTE_DML_ALL=false*/delete from tt;
Query OK, 10 rows affected (0.21 sec)
```

- Full-table update is intercepted by default.

```
mysql> update tt set id = 1;
ERR-CODE: [TDDL-4620][ERR_FORBID_EXECUTE_DML_ALL] Forbid execute DELETE ALL or UPDATE ALL sql. More: [http://middleware.alibaba-inc.com/faq/faqByFaqCode.html?faqCode=TDDL-4620]
```

The operation is successful if the following HINT is added:

```
mysql> /*! TDDL:FORBID_EXECUTE_DML_ALL=false*/update tt set id = 1;
Query OK, 10 rows affected (0.21 sec)
```

- This limit does not apply to DELETE or UPDATE statements that contain the WHERE or LIMIT condition.

```
mysql> delete from tt where id = 1;
Query OK, 1 row affected (0.21 sec)
```

11.13. PolarDB-X sequence

11.13.1. Overview

A Distributed Relational Database Service (DRDS) sequence is a 64-digit number that corresponds to the signed BIGINT type in MySQL. It is used to create a globally unique and sequentially incremental numeric sequence, such as the values of primary key columns and unique index columns.

DRDS sequences are used in the following two ways:

- Explicit sequences are created and maintained by using sequence-specific data definition language (DDL) syntax and can be used independently. The sequence value can be acquired by using `select seq.nextval;`, in which `seq` indicates the sequence name.
- Implicit sequences are used to automatically fill in primary keys with `AUTO_INCREMENT` defined and are automatically maintained by DRDS.

 **Notice** DRDS creates implicit sequences only after `AUTO_INCREMENT` is defined for partitioned tables and broadcast tables. This is not the case for non-partition tables. The `AUTO_INCREMENT` value of a non-partition table is created by ApsaraDB RDS for MySQL.

Types and features of DRDS sequences

Currently, three types of DRDS sequences are supported.

| Type (abbreviation) | Globally unique | Consecutive | Monotonically increasing | Monotonically increasing within the same connection | Non-single point | Data type | Readability |
|----------------------------|-----------------|-------------|---|---|------------------|--------------------------|-------------|
| Group sequence (GROUP) | Yes | No | No | Yes | Yes | All integer types | High |
| Time-based sequence (TIME) | Yes | No | Monotonically increasing at the macro level and non-monotonically increasing at the micro level | Yes | Yes | Only BIGINT is supported | Low |
| Simple sequence (SIMPLE) | Yes | Yes | Yes | Yes | No | All integer types | High |

Concepts:

- **Consecutive:** If the current value is n, the next value must be n + 1. If the next value is not n + 1, it is nonconsecutive.
- **Monotonically increasing:** If the current value is n, the next value must be a number greater than n.
- **Single point:** The risk of single point of failure exists.
- **Monotonically increasing at the macro level and non-monotonically increasing at the micro level:** An example of this is 1, 3, 2, 4, 5, 7, 6, 8, ...

Group sequence (GROUP, used by default)

Features

A group sequence is a globally unique sequence with natural numeric values, which are not necessarily consecutive or monotonically increasing. If the sequence type is not specified, DRDS uses the group sequence type by default.

- **Advantages:** A group sequence is globally unique and provides excellent performance, preventing single point of failure.
- **Disadvantages:** A group sequence may contain nonconsecutive values, which may not necessarily start from the initial value and do not cycle.

Implementation

The values of a group sequence are created by multiple nodes to ensure high availability. The values in a segment are nonconsecutive if the values are not all used, such as in the case of disconnection.

Time-based sequence

Features

A time-based sequence consists of a **timestamp, node ID, and serial number**. It is globally unique and automatically increments at the macro level. Value updates are database-independent and not persistently stored in databases. Only names and types are stored in databases. This delivers good performance to time-based sequences, which create values like 776668092129345536, 776668098018148352, 776668111578333184, and 776668114812141568.

 **Notice** Sequence values must be of the **BIGINT** type when used in the auto-increment columns of tables.

- **Advantages:** Time-based sequences are globally unique with good performance.
- **Disadvantages:** The values of a time-based sequence are nonconsecutive. The **START WITH**, **INCREMENT BY**, **MAXVALUE**, and **CYCLE** or **NOCYCLE** parameters are invalid for time-based sequences.

Simple sequence

Features

Only simple sequences support the **START WITH**, **INCREMENT BY**, **MAXVALUE**, and **CYCLE** or **NOCYCLE** parameters.

- **Advantages:** Simple sequences are globally unique and monotonically increasing with consecutive values.
- **Disadvantages:** Simple sequences are prone to single point of failure, poor performance, and bottlenecks. Use them with caution.

Implementation

Each sequence value must be persistently stored.

Scenarios

Group sequences, time-based sequences, and simple sequences are globally unique and can be used in **primary key columns** and **unique index columns**.

- We recommend that you use **group sequences**.
- Use only simple sequences for services that strongly depend on consecutive sequence values. Pay attention to sequence performance.
- We recommend that you use time-based sequences if you have high requirements for sequence performance, the amount of data inserted to tables is small, and large sequence values are acceptable. Time-based sequences are CPU-bound with no requirements on computing lock, database dependence, or persistent storage.

The following example shows how to create a sequence with an initial value of 100000 and a step of 1.

- A **simple sequence** creates globally unique, consecutive, and monotonically increasing values, such as 100000, 100001, 100002, 100003, 100004, ..., 200000, 200001, 200002, 200003... The values of a simple sequence are persistently stored. Even after services are restarted upon a single point of failure, values are still created consecutively from the breakpoint. However, simple sequences have poor performance because each value is persistently stored once it is created.
- A **group sequence** may create values like 200001, 200002, 200003, 200004, 100001, 100002, 100003...

Notice

- The initial value of a group sequence is not necessarily the same as the **START WITH** value (which is 100000 in this example) but is invariably greater than this value. In this example, the initial value is 200001.
- A group sequence is globally unique but may contain nonconsecutive values, which may occur when a node is faulty or the connection that only uses partial values is closed. The group sequence in this example contains nonconsecutive values because the values between 200004 and 100001 are missing.

- A **time-based sequence** may create values like 776668092129345536, 776668098018148352, 776668111578333184, 776668114812141568...

11.13.2. Explicit sequence usage

This topic describes how to use data definition language (DDL) statements to create, modify, delete, and query sequences and how to acquire the values of explicit sequences.

Create a sequence

Syntax:

```
CREATE [ GROUP | SIMPLE | TIME ] SEQUENCE <name>
[ START WITH <numeric value> ] [ INCREMENT BY <numeric value> ]
[ MAXVALUE <numeric value> ] [ CYCLE | NOCYCLE ]
```

Parameters:

| Parameter | Description | Applicable To |
|------------------|--|------------------------------------|
| START WITH | The initial sequence value. If it is not set, the default value is 1. | Simple sequence and group sequence |
| INCREMENT BY | The increment (or interval value or step) of each sequence increase. If it is not set, the default value is 1. | Simple sequence |
| MAXVALUE | The maximum sequence value. If it is not specified, the default value is the maximum value of the signed BIGINT type. | Simple sequence |
| CYCLE or NOCYCLE | Indicates whether to repeat the sequence value which starts from the value specified by START WITH after the sequence value reaches the maximum value. If it is not specified, the default value is NOCYCLE. | Simple sequence |

Note

- If the sequence type is not specified, the group sequence type is used by default.
- The INCREMENT BY, MAXVALUE, and CYCLE or NOCYCLE parameters are invalid for group sequences.
- The START WITH, INCREMENT BY, MAXVALUE, and CYCLE or NOCYCLE parameters are invalid for time-based sequences.
- Group sequences are nonconsecutive. The START WITH parameter only provides reference for group sequences. **The initial group sequence value is not necessarily the same as but is greater than the value of START WITH.**

Example 1: Create a group sequence.

- Method 1:

```
mysql> CREATE SEQUENCE seq1;
Query OK, 1 row affected (0.27 sec)
```

- Method 2:

```
mysql> CREATE GROUP SEQUENCE seq1;
Query OK, 1 row affected (0.27 sec)
```

Example 2: Create a time-based sequence.

```
mysql> CREATE TIME SEQUENCE seq1;
Query OK, 1 row affected (0.27 sec)
```

Example 3: Create a simple sequence with an initial value of 1000, step of 2, and maximum value of 9999999999, which does not repeat after increasing to the maximum value.

```
mysql> CREATE SIMPLE SEQUENCE seq2 START WITH 1000 INCREMENT BY 2 MAXVALUE 9999999999 NOCYCLE;
Query OK, 1 row affected (0.03 sec)
```

Modify a sequence

Distributed Relational Database Service (DRDS) allows you to modify sequences in the following ways:

- For simple sequences, change the values of START WITH, INCREMENT BY, MAXVALUE, and CYCLE or NOCYCLE.
- For group sequences, change the value of START WITH.
- Convert the sequence type to another.

Syntax:

```
ALTER SEQUENCE <name> [ CHANGE TO GROUP | SIMPLE | TIME ]
START WITH <numeric value> [ INCREMENT BY <numeric value> ]
[ MAXVALUE <numeric value> ] [ CYCLE | NOCYCLE ]
```

Parameters:

| Parameter | Description | Applicable To |
|------------------|--|------------------------------------|
| START WITH | The initial sequence value. If it is not set, the default value is 1. | Simple sequence and group sequence |
| INCREMENT BY | The increment (or interval value or step) of each sequence increase. If it is not set, the default value is 1. | Simple sequence |
| MAXVALUE | The maximum sequence value. If it is not specified, the default value is the maximum value of the signed BIGINT type. | Simple sequence |
| CYCLE or NOCYCLE | Indicates whether to repeat the sequence value which starts from the value specified by START WITH after the sequence value reaches the maximum value. If it is not specified, the default value is NOCYCLE. | Simple sequence |

Note

- Group sequences are nonconsecutive. The START WITH parameter only provides reference for group sequences. The initial group sequence value is not necessarily the same as but is greater than the value of START WITH.
- If you set START WITH when modifying a simple sequence, the START WITH value takes effect immediately. The following sequence value starts from the new START WITH value. For example, if you change the START WITH value to 200 when the sequence value increases to 100, the following sequence value starts from 200.
- Before changing the START WITH value, you need to analyze the existing sequence values and the speed of creating sequence values to avoid conflicts. Do not change the START WITH value unless necessary.

Example: Change the initial value, step, and maximum value of the simple sequence named seq2 to 3000, 5, and 1000000 respectively, and set CYCLE.

```
mysql> ALTER SEQUENCE seq2 START WITH 3000 INCREMENT BY 5 MAXVALUE 1000000 CYCLE;
Query OK, 1 row affected (0.01 sec)
```

Convert the sequence type to another

- Use the `CHANGE TO <sequence_type>` clause of `ALTER SEQUENCE`.
- If the `CHANGE TO` clause of `ALTER SEQUENCE` is specified, add the `START WITH` parameter to prevent duplicate values. This parameter is optional if `CHANGE TO` is not specified.

Example: Convert a group sequence to a simple sequence.

```
mysql> ALTER SEQUENCE seq1 CHANGE TO SIMPLE START WITH 1000000;
Query OK, 1 row affected (0.02 sec)
```

Delete a sequence

Syntax:

```
DROP SEQUENCE <name>
```

Example:

```
mysql> DROP SEQUENCE seq3;
Query OK, 1 row affected (0.02 sec)
```

Query sequences

Syntax:

```
SHOW SEQUENCES
```

Example: The TYPE column lists the sequence types in the abbreviated form.

```
mysql> SHOW SEQUENCES;
+-----+-----+-----+-----+-----+-----+-----+
| NAME | VALUE | INCREMENT_BY | START_WITH | MAX_VALUE | CYCLE | TYPE |
+-----+-----+-----+-----+-----+-----+-----+
| AUTO_SEQ_1 | 91820513 | 1 | 91820200 | 9223372036854775807 | N | SIMPLE |
| AUTO_SEQ_4 | 91820200 | 2 | 1000 | 9223372036854775807 | Y | SIMPLE |
| seq_test | N/A | N/A | N/A | N/A | N/A | TIME |
| AUTO_SEQ_2 | 100000 | N/A | N/A | N/A | N/A | GROUP |
| AUTO_SEQ_3 | 200000 | N/A | N/A | N/A | N/A | GROUP |
+-----+-----+-----+-----+-----+-----+-----+
5 rows in set (0.01 sec)
```

Get the sequence value

Syntax:

```
<sequence name> .NEXTVAL
```

Example:

```
SELECT sample_seq.nextVal FROM dual;
+-----+
| SAMPLE_SEQ.NEXTVAL |
+-----+
|      101001 |
+-----+
1 row in set (0.04 sec)
```

You can also write `SAMPLE_SEQ.nextVal` as a value to the SQL statement:

```
mysql> INSERT INTO some_users (name,address,gmt_create,gmt_modified,intro) VALUES ('sun',SAMPLE_SEQ.nextVal,now(),now(),'aa');
Query OK, 1 row affected (0.01 sec)
```

Note If you set the `AUTO_INCREMENT` parameter when creating a table, you do not need to specify an auto-increment column when running the `INSERT` statement. The auto-increment column is automatically maintained by DRDS.

Acquire the values of sequences in batches

Syntax:

```
SELECT <sequence name>.NEXTVAL FROM DUAL WHERE COUNT = <numeric value >
```

Example:

```
SELECT sample_seq.nextVal FROM dual WHERE count = 10;
+-----+
| SAMPLE_SEQ.NEXTVAL |
+-----+
|      101002 |
|      101003 |
|      101004 |
|      101005 |
|      101006 |
|      101007 |
|      101008 |
|      101009 |
|      101010 |
|      101011 |
+-----+
10 rows in set (0.04 sec)
```

11.13.3. Implicit sequence usage

After `AUTO_INCREMENT` is set for a primary key, the primary key is automatically filled in by using a sequence which is maintained by Distributed Relational Database Service (DRDS).

CREATE TABLE

The standard `CREATE TABLE` syntax is extended to add the sequence type for auto-increment columns. If the type keyword is not specified, the default type is `GROUP`. The sequence names that are automatically created by DRDS and associated with tables are prefixed with `AUTO_SEQ_` and suffixed with the table name.

```
CREATE TABLE <name> (  
  <column> ... AUTO_INCREMENT [ BY GROUP | SIMPLE | TIME ],  
  <column definition>,  
  ...  
) ... AUTO_INCREMENT=<start value>
```

SHOW CREATE TABLE

The sequence type is displayed for the auto-increment column of a table shard or broadcast table.

```
SHOW CREATE TABLE <name>
```

Examples

- If AUTO_INCREMENT is set but the sequence type is not specified when a table is created, the group sequence type is used by default.

Example 1

```
mysql> CREATE TABLE `xkv_shard` (  
  -> `id` bigint(20) unsigned NOT NULL AUTO_INCREMENT COMMENT ' ',  
  -> `gmt_create` timestamp NOT NULL DEFAULT '0000-00-00 00:00:00' ON UPDATE CURRENT_TIMESTAMP COMMENT ' ',  
  -> `uid` bigint(20) unsigned DEFAULT '10' COMMENT 'uid',  
  -> `msg` varchar(40) DEFAULT '127.0.0.1' COMMENT 'desc',  
  -> `val` float DEFAULT '0' COMMENT 'val',  
  -> `time` time DEFAULT NULL COMMENT 'time',  
  -> PRIMARY KEY (`id`),  
  -> UNIQUE KEY `msg` (`msg`)  
  -> ) ENGINE=InnoDB AUTO_INCREMENT=100009 DEFAULT CHARSET=utf8 dbpartition by hash(`id`);  
Query OK, 0 rows affected (1.24 sec)  
  
mysql> show create table xkv_shard;  
  
+-----+-----+  
| Table | Create Table  
+-----+-----+  
  
+-----+-----+  
| xkv_shard | CREATE TABLE `xkv_shard` (  
  `id` bigint(20) unsigned NOT NULL AUTO_INCREMENT BY GROUP COMMENT ' ',  
  `gmt_create` timestamp NOT NULL DEFAULT '0000-00-00 00:00:00' ON UPDATE CURRENT_TIMESTAMP COMMENT ' ',  
  `uid` bigint(20) unsigned DEFAULT '10' COMMENT 'uid',  
  `msg` varchar(40) DEFAULT '127.0.0.1' COMMENT 'desc',  
  `val` float DEFAULT '0' COMMENT 'val',  
  `time` time DEFAULT NULL COMMENT 'time',  
  PRIMARY KEY (`id`),  
  UNIQUE KEY `msg` (`msg`)  
  ) ENGINE=InnoDB AUTO_INCREMENT=100009 DEFAULT CHARSET=utf8 dbpartition by hash(`id`) |  
+-----+-----+  
1 row in set (0.02 sec)  
  
mysql> drop table xkv_shard;
```

- When creating a table, set AUTO_INCREMENT and specify a time-based sequence as the primary key value.

Example 2

```
mysql> CREATE TABLE `timeseq_test` (  
  -> `id` bigint(20) unsigned NOT NULL AUTO_INCREMENT BY TIME COMMENT ' ',  
  -> `gmt_create` timestamp NOT NULL DEFAULT '0000-00-00 00:00:00' ON UPDATE CURRENT_TIMESTAMP COMMENT ' ',  
  -> `uid` bigint(20) unsigned DEFAULT '10' COMMENT 'uid',  
  -> `msg` varchar(40) DEFAULT '127.0.0.1' COMMENT 'desc',  
  -> `val` float DEFAULT '0' COMMENT 'val',  
  -> `time` time DEFAULT NULL COMMENT 'time',  
  -> PRIMARY KEY (`id`),  
  -> UNIQUE KEY `msg` (`msg`)  
  -> ) ENGINE=InnoDB AUTO_INCREMENT=100009 DEFAULT CHARSET=utf8 dbpartition by hash(`id`);  
Query OK, 0 rows affected (1.27 sec)  
  
mysql> show create table timeseq_test;  
  
+-----+-----+  
| Table | Create Table  
+-----+-----+  
  
+-----+-----+  
| timeseq_test | CREATE TABLE `timeseq_test` (  
  `id` bigint(20) unsigned NOT NULL AUTO_INCREMENT BY TIME COMMENT ' ',  
  `gmt_create` timestamp NOT NULL DEFAULT '0000-00-00 00:00:00' ON UPDATE CURRENT_TIMESTAMP COMMENT ' ',  
  `uid` bigint(20) unsigned DEFAULT '10' COMMENT 'uid',  
  `msg` varchar(40) DEFAULT '127.0.0.1' COMMENT 'desc',  
  `val` float DEFAULT '0' COMMENT 'val',  
  `time` time DEFAULT NULL COMMENT 'time',  
  PRIMARY KEY (`id`),  
  UNIQUE KEY `msg` (`msg`)  
  ) ENGINE=InnoDB AUTO_INCREMENT=100009 DEFAULT CHARSET=utf8 dbpartition by hash(`id`) |  
+-----+-----+  
1 row in set (0.04 sec)
```

ALTER TABLE

Currently, **ALTER TABLE** cannot be used to change the sequence type but can be used to change the initial value. If you want to modify the data of the implicit sequence type in a table, run **SHOW SEQUENCES** to find the names and types of sequences and run **ALTER SEQUENCE** to modify the data.

```
ALTER TABLE <name> ... AUTO_INCREMENT=<start value>
```

Notice Exercise caution when changing the initial value of `AUTO_INCREMENT` after DRDS sequences are used. You need to analyze the existing sequence values and the speed of creating sequence values to avoid conflicts.

11.13.4. Limits and precautions for sequences

This topic describes the limits and precautions for sequences.

Limits and precautions

- When a time-based sequence is used in the auto-increment column of a table, the column must be of the `BIGINT` type.
- `START WITH` must be set when the sequence is changed to another type.
- When the `INSERT` statement is executed on a PolarDB-X database in non-partition mode where only one ApsaraDB RDS for MySQL database is bound or on a database in partition mode that has only one table but no broadcast table, PolarDB-X automatically optimizes and sends the statement, and bypasses the part of the optimizer that allocates the sequence value. In this case, `INSERT INTO ... VALUES (seq.nextval, ...)` is not supported. We recommend that you use the ApsaraDB RDS for MySQL auto-increment column feature instead.
- If the hint for a specific database shard is used by the `INSERT` statement such as `INSERT INTO ... VALUES ...` or `INSERT INTO ... SELECT ...` and the target table uses a sequence, PolarDB-X bypasses the optimizer and directly sends the statement so that the sequence does not take effect. The target table creates an ID by using the auto-increment feature of the ApsaraDB RDS for MySQL table.
- The auto-increment ID allocation method for the same table must be kept consistent, which may be based on PolarDB-X sequences or the auto-increment column of the ApsaraDB RDS for MySQL. If both of the two allocation methods are used for the same table, duplicate IDs may be created and making location difficult.

Troubleshoot primary key conflicts

Assume that data is directly written to ApsaraDB RDS for MySQL and that the related primary key value is not the sequence value created by PolarDB-X. If PolarDB-X automatically creates a primary key and writes it to the database, this primary key may conflict with that of the directly written data. This problem can be resolved as follows:

1. View the existing sequences by using the PolarDB-X-specified SQL statement. The sequence prefixed with `AUTO_SEQ_` is an implicit sequence. This sequence is generated when a table is created with the `AUTO_INCREMENT` parameter.

```
mysql> SHOW SEQUENCES;
+-----+-----+-----+-----+-----+-----+-----+
| NAME          | VALUE | INCREMENT_BY | START_WITH | MAX_VALUE | CYCLE | TYPE |
+-----+-----+-----+-----+-----+-----+-----+
| AUTO_SEQ_timeseq_test | N/A | N/A | N/A | N/A | N/A | TIME |
| AUTO_SEQ_xkv_shard_tbl1 | 0 | N/A | N/A | N/A | N/A | GROUP |
| AUTO_SEQ_xkv_shard | 0 | N/A | N/A | N/A | N/A | GROUP |
+-----+-----+-----+-----+-----+-----+-----+
3 rows in set (0.04 sec)
```

2. For example, if the `t_item` table contains conflicts and its primary key is ID, then retrieve the maximum primary key value of this table from PolarDB-X:

```
mysql> SELECT MAX(id) FROM t_item;
+-----+
| max(id) |
+-----+
| 8231 |
+-----+
1 row in set (0.01 sec)
```

3. Update the related value in the PolarDB-X sequence table to a value greater than 8231, such as 9000. Then, no error is returned for the auto-increment primary key created by the subsequent `INSERT` statement.

```
mysql> ALTER SEQUENCE AUTO_SEQ_USERS START WITH 9000;
Query OK, 1 row affected (0.01 sec)
```

11.14. Best practices

11.14.1. Select a shard key

A shard key is a field for database sharding and table sharding, which is used to create sharding rules during horizontal partitioning. PolarDB-X partitions a logical table horizontally into the physical database shards on each ApsaraDB RDS for MySQL instance based on the shard key.

The primary principle of sharding is to identify the business logic-specific subject of data in a table as much as possible and confirm that most (or core) database operations are performed based on this subject. Then, use the subject-related field as the shard key to perform database sharding and table sharding.

The business logic-specific subject is related to business scenarios. The following typical scenarios include business logic-specific subjects that can be used as shard keys:

- User-oriented Internet applications are operated to meet user requirements. Users are the business logic-specific subject and the user-related field can be used as the shard key.
- Seller-oriented e-commerce applications are operated to meet seller requirements. Sellers are the business logic-specific subject and the seller-related field can be used as the shard key.
- Game-oriented applications are operated to meet gamer requirements. Gamers are the business logic-specific subject and the gamer-related field can be used as the shard key.
- Internet of Vehicles (IoV) applications are operated based on vehicle information. Vehicles are the business logic-specific subject and the vehicle-related field can be used as the shard key.
- Tax-oriented applications are operated based on taxpayer information to provide front-end services. Taxpayers are the business logic-specific subject and the taxpayer-related field can be used as the shard key.

In other scenarios, you can also use the appropriate subject of business logic as the shard key.

For example, in a seller-oriented e-commerce application, the following single table must be horizontally partitioned:

```
CREATE TABLE sample_order (
  id INT(11) NOT NULL,
  sellerId INT(11) NOT NULL,
  trade_id INT(11) NOT NULL,
  buyer_id INT(11) NOT NULL,
  buyer_nick VARCHAR(64) DEFAULT NULL,
  PRIMARY KEY (id)
)
```

The `sellerId` field is used as the shard key because seller is the business logic-specific subject. In the case of database sharding but no table sharding, the distributed data definition language (DDL) statement for table creation is as follows:

```
CREATE TABLE sample_order (  
  id INT(11) NOT NULL,  
  sellerId INT(11) NOT NULL,  
  trade_id INT(11) NOT NULL,  
  buyer_id INT(11) NOT NULL,  
  buyer_nick VARCHAR(64) DEFAULT NULL,  
  PRIMARY KEY (id)  
) DBPARTITION BY HASH(sellerId)
```

If no business logic-specific subject can be used as the shard key, use the following methods to select an appropriate shard key:

- Determine the shard key based on the distribution and access of data. Distribute the data in a table to different physical database shards and table shards as evenly as possible. This method is applicable to scenarios with massive analytical queries, in which query concurrency stays at 1.
- Determine the shard key for database sharding and table sharding by combining fields of the numeric (string) type and time type. This method is applicable to log retrieval.

For example, a log system records all user operations and needs to horizontally partition the following single log table:

```
CREATE TABLE user_log (  
  userId INT(11) NOT NULL,  
  name VARCHAR(64) NOT NULL,  
  operation VARCHAR(128) DEFAULT NULL,  
  actionDate DATE DEFAULT NULL  
)
```

You can combine the user identifier with the time field to create a shard key for partitioning the table by week. The distributed DDL statement for table creation is as follows:

```
CREATE TABLE user_log (  
  userId INT(11) NOT NULL,  
  name VARCHAR(64) NOT NULL,  
  operation VARCHAR(128) DEFAULT NULL,  
  actionDate DATE DEFAULT NULL  
) DBPARTITION BY HASH(userId) TBPARTITION BY WEEK(actionDate) TBPARTITIONS 7
```

For more information about shard key selection and table shard forms, see [DDL statements](#).

 Notice Avoid using hotspot data as the shard key if possible.

11.14.2. Select the number of shards

Distributed Relational Database Service (DRDS) supports horizontal partitioning of databases and tables. By default, eight physical database shards are created on an ApsaraDB for RDS instance, and one or more physical table shards can be created on each physical database shard. The number of table shards is also called the number of shards.

Generally, we recommend that each physical table shard contain no more than 5 million rows of data. Generally, you can estimate the data growth in one to two years. Divide the estimated total data size by the total number of physical database shards, and then divide the result by the recommended maximum data size of 5 million, to obtain the number of physical table shards to be created on each physical database shard:

```
Number of physical table shards in each physical database shard = CEILING(Estimated total data size/(Number of Apsara DB for RDS instances x 8)/5,000,000)
```

Therefore, when the calculated number of physical table shards is equal to 1, only database sharding needs to be performed, that is, a physical table shard is created in each physical database shard. If the calculation result is greater than 1, we recommend that you perform both database sharding and table sharding, that is, there are multiple physical table shards in each physical database shard.

For example, if a user estimates that the total data size of a table will be about 0.1 billion rows two years later and the user has four ApsaraDB for RDS instances, then according to the preceding formula:

$$\text{Number of physical table shards in each physical database shard} = \text{CEILING}(100,000,000 / (4 \times 8) / 5,000,000) = \text{CEILING}(0.625) = 1$$

If the result is 1, only database sharding is needed, that is, one physical table shard is created in each physical database shard.

If only one ApsaraDB for RDS instance is used in the preceding example, the formula is as follows:

$$\text{Number of physical table shards in each physical database shard} = \text{CEILING}(100,000,000 / (1 \times 8) / 5,000,000) = \text{CEILING}(2.5) = 3$$

If the result is 3, we recommend that you create three physical table shards in each physical database shard.

11.14.3. Basic concepts of SQL optimization

Distributed Relational Database Service (DRDS) is an efficient and stable distributed relational database service that processes distributed relational computing. DRDS optimizes SQL statements differently from single-instance relational databases such as MySQL. DRDS focuses on the network I/O overheads in a distributed environment and pushes SQL operations down to the underlying database shards (such as databases on ApsaraDB for RDS instances) for execution, thereby reducing the network I/O overheads and improving the SQL execution efficiency.

DRDS provides commands for obtaining the SQL execution information to help SQL optimization, for example, EXPLAIN commands for obtaining SQL execution plans and TRACE commands for obtaining SQL execution processes and overheads. This topic describes the basic concepts and common commands related to SQL optimization in DRDS.

Execution plan

An SQL execution plan (or execution plan) is a set of ordered operation steps generated to access data. In DRDS, the execution plan is divided into the execution plan at the DRDS layer and the execution plan at the ApsaraDB for RDS layer. Execution plan analysis is an effective way to optimize SQL statements. Through execution plan analysis, you can know whether DRDS or ApsaraDB for RDS has generated optimal execution plans for SQL statements and whether further optimization can be made.

During SQL statement execution, based on the basic information of the SQL statement and related tables, the DRDS optimizer determines on which the database shards the SQL statement should be executed, and the specific SQL statement form, execution policy, and data merging and computing policy for each database shard. This process optimizes SQL statement execution and generates execution plans at the DRDS layer. The execution plan at the ApsaraDB for RDS layer is the native MySQL execution plan.

DRDS provides a set of EXPLAIN commands to display execution plans at different levels or with different levels of detail.

The following table briefly describes the EXPLAIN commands in DRDS.

EXPLAIN command description

| Command | Description | Example |
|---------|-------------|---------|
|---------|-------------|---------|

| Command | Description | Example |
|-------------------------|---|------------------------------------|
| EXPLAIN { SQL } | Displays the summary execution plan of SQL statements at the DRDS layer, including the database shards on which the SQL statement is run, physical statements, and general parameters. | EXPLAIN SELECT * FROM test |
| EXPLAIN DETAIL { SQL } | Displays the detailed execution plans of SQL statements at the DRDS layer, including the statement type, concurrency, returned field information, physical tables, and database groups. | EXPLAIN DETAIL SELECT * FROM test |
| EXPLAIN EXECUTE { SQL } | Displays the execution plan of the underlying ApsaraDB for RDS instance, which is equivalent to the EXPLAIN statement of MySQL. | EXPLAIN EXECUTE SELECT * FROM test |

Execution plans at the DRDS layer

The following table describes the fields in the results returned for a DRDS-layer execution plan.

Description of fields in DRDS-layer execution plans

| Field | Description |
|------------|---|
| GROUP_NAME | The name of the DRDS database shard. The suffix identifies the specific database shard. The value is consistent with the result of the SHOW NODE command. |
| SQL | The SQL statement run on this database shard. |
| PARAMS | The SQL statement parameters used when DRDS communicates with ApsaraDB for RDS over the Prepare protocol. |

The SQL field can be in two forms:

1. If an SQL statement does not contain the following parts, the execution plan is displayed as an SQL statement:
 - Aggregate function involving multiple database shards.
 - Distributed JOIN queries involving multiple shards.
 - Complex subqueries.

Example:

```
mysql> EXPLAIN SELECT * FROM test;
+-----+-----+-----+
| GROUP_NAME          | SQL                                | PARAMS |
+-----+-----+-----+
| TESTDB_1478746391548CDTCTESTDB_OXGJ_0000_RDS | select `test`.`c1`,`test`.`c2` from `test` | {} |
| TESTDB_1478746391548CDTCTESTDB_OXGJ_0001_RDS | select `test`.`c1`,`test`.`c2` from `test` | {} |
| TESTDB_1478746391548CDTCTESTDB_OXGJ_0002_RDS | select `test`.`c1`,`test`.`c2` from `test` | {} |
| TESTDB_1478746391548CDTCTESTDB_OXGJ_0003_RDS | select `test`.`c1`,`test`.`c2` from `test` | {} |
| TESTDB_1478746391548CDTCTESTDB_OXGJ_0004_RDS | select `test`.`c1`,`test`.`c2` from `test` | {} |
| TESTDB_1478746391548CDTCTESTDB_OXGJ_0005_RDS | select `test`.`c1`,`test`.`c2` from `test` | {} |
| TESTDB_1478746391548CDTCTESTDB_OXGJ_0006_RDS | select `test`.`c1`,`test`.`c2` from `test` | {} |
| TESTDB_1478746391548CDTCTESTDB_OXGJ_0007_RDS | select `test`.`c1`,`test`.`c2` from `test` | {} |
+-----+-----+-----+
8 rows in set (0.04 sec)
```

The group names displayed in the GROUP_NAME field can be found in the returned result of SHOW NODE:

```
mysql> SHOW NODE;
+-----+-----+-----+-----+-----+
| ID | NAME                                | MASTER_READ_COUNT | SLAVE_READ_COUNT | MASTER_READ_PERCENT | SLAVE_READ_PERCENT |
+-----+-----+-----+-----+-----+
| 0 | TESTDB_1478746391548CDTCTESTDB_OXGJ_0000_RDS | 69 | 0 | 100% | 0% |
| 1 | TESTDB_1478746391548CDTCTESTDB_OXGJ_0001_RDS | 45 | 0 | 100% | 0% |
| 2 | TESTDB_1478746391548CDTCTESTDB_OXGJ_0002_RDS | 30 | 0 | 100% | 0% |
| 3 | TESTDB_1478746391548CDTCTESTDB_OXGJ_0003_RDS | 29 | 0 | 100% | 0% |
| 4 | TESTDB_1478746391548CDTCTESTDB_OXGJ_0004_RDS | 11 | 0 | 100% | 0% |
| 5 | TESTDB_1478746391548CDTCTESTDB_OXGJ_0005_RDS | 1 | 0 | 100% | 0% |
| 6 | TESTDB_1478746391548CDTCTESTDB_OXGJ_0006_RDS | 8 | 0 | 100% | 0% |
| 7 | TESTDB_1478746391548CDTCTESTDB_OXGJ_0007_RDS | 50 | 0 | 100% | 0% |
+-----+-----+-----+-----+-----+
8 rows in set (0.10 sec)
```

2. Execution plans that cannot be expressed by SQL statements can be expressed by DRDS in custom format.

Example:

```
mysql> EXPLAIN DETAIL SELECT COUNT(*) FROM test;
+-----+-----+-----+
| GROUP_NAME          | SQL          | PARAMS |
+-----+-----+-----+
| TEST_DB_1478746391548CDTCTEST_DB_OXGJ_0000_RDS | Merge as test
queryConcurrency:GROUP_CONCURRENT
columns:[count(*)]
executeOn: TEST_DB_1478746391548CDTCTEST_DB_OXGJ_0000_RDS
  Query from test as test
    queryConcurrency:SEQUENTIAL
    columns:[count(*)]
    tableName:test
    executeOn: TEST_DB_1478746391548CDTCTEST_DB_OXGJ_0000_RDS
  Query from test as test
    queryConcurrency:SEQUENTIAL
    columns:[count(*)]
    tableName:test
    executeOn: TEST_DB_1478746391548CDTCTEST_DB_OXGJ_0001_RDS
  ... ..
  Query from test as test
    queryConcurrency:SEQUENTIAL
    columns:[count(*)]
    tableName:test
    executeOn: TEST_DB_1478746391548CDTCTEST_DB_OXGJ_0007_RDS
| NULL |
+-----+-----+-----+
1 row in set (0.00 sec)
```

executeOn in the SQL statement field indicates the database shard on which the SQL statement is run. DRDS finally merges the results returned by the database shards.

Execution plans at the ApsaraDB for RDS layer

The execution plans at the ApsaraDB for RDS layer are the same as the native MySQL execution plan. For more information, see [official MySQL documentation](#).

One DRDS logical table may consist of multiple shards distributed in different database shards. Therefore, you can view the execution plans at the ApsaraDB for RDS layer in multiple ways.

1. View the execution plan of an ApsaraDB for RDS database shard.

If the query condition contains a shard key, directly run the EXPLAIN EXECUTE command to display the execution plan on the corresponding database shard. Example:

```
mysql> EXPLAIN EXECUTE SELECT * FROM test WHERE c1 = 1;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | test | const | PRIMARY | PRIMARY | 4 | const | 1 | NULL |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.04 sec)
```

Notice If an SQL statement involves multiple shards (for example, its condition does not contain a shard key), the EXPLAIN EXECUTE command returns an execution plan on a random ApsaraDB for RDS database shard.

To view the execution plan of an SQL statement on a specified database shard, you can use the Hint method. Example:

```
mysql> /*! TDDL:node='TESTDB_1478746391548CDTCTESTDB_OXGJ_0000_RDS'*/EXPLAIN SELECT * FROM test;
+----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
+----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | test | ALL | NULL | NULL | NULL | NULL | 2 | NULL |
+----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.04 sec)
```

2. View the execution plans of all ApsaraDB for RDS database shards.

You can run SCAN Hint to display the execution plans of SQL statements on all database shards:

```
mysql> /*! TDDL:scan='test'*/EXPLAIN SELECT * FROM test;
+----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
+----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | test | ALL | NULL | NULL | NULL | NULL | 2 | NULL |
| 1 | SIMPLE | test | ALL | NULL | NULL | NULL | NULL | 3 | NULL |
+----+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.08 sec)
```

 **Notice**

- i. In Hint mode, DRDS only replaces the table names in case of database or table sharding, and then directly sends the logical SQL statement to ApsaraDB for RDS for execution. It will not process the result.
- ii. Execution plans obtained by using an EXPLAIN command are generated by static analysis and are not actually executed in databases.

TRACE command

The TRACE command in DRDS can track the SQL execution process and the overheads in each stage. It can be used together with the execution plan to facilitate SQL statement optimization.

The TRACE command contains two related commands: TRACE and SHOW TRACE, which must be used together.

11.14.4. SQL optimization methods

11.14.4.1. Overview

This topic describes the SQL optimization principles and methods for optimizing different types of SQL statements in PolarDB-X.

Basic principles of SQL optimization

In PolarDB-X, SQL computing that can be performed by ApsaraDB RDS for MySQL instances is called push-down computing. Push-down computing reduces data transmission, decreases overheads at the network layer and PolarDB-X layer, and improves the execution efficiency of SQL statements.

Therefore, the basic principle for SQL statement optimization in PolarDB-X is as follows: Push down as many computations as possible to ApsaraDB RDS for MySQL instances.

Push-down computations include:

- JOIN connections
- Filter conditions, such as WHERE or HAVING conditions
- Aggregate computing, such as COUNT and GROUP BY
- Sorting, such as ORDER BY

- Deduplication, such as `DISTINCT`
- Function computing, such as the `NOW()` function
- Subqueries

 **Notice** The preceding list only describes possible forms of push-down computations. It does not mean that all clauses or conditions or combinations of clauses or conditions can be pushed down for computing.

SQL statements of different types and containing different conditions can be optimized in different ways. The following uses some examples to describe how to optimize SQL statements:

- Single-table SQL optimization
 - Filter condition optimization
 - Optimization of the number of returned rows for a query
 - Grouping and sorting optimization
- JOIN query optimization
 - Optimization of push-down JOIN queries
 - Optimization of distributed JOIN queries
- Subquery optimization

11.14.4.2. Single-table SQL optimization

Single-table SQL optimization must follow the following principles:

- Make sure that the SQL statements contain the shard key.
- Use an equivalence condition for the shard key whenever possible.
- If the shard key is an IN condition, the number of values after IN should be as small as possible (far fewer than the number of shards, and remain unchanged as the business grows).
- If SQL statements do not contain a shard key, use only one of `DISTINCT`, `GROUP BY`, and `ORDER BY` in the same SQL statement.

Filter condition optimization

DRDS partitions data horizontally by the shard key. Therefore, the filter condition must contain the shard key as much as possible so that DRDS can push queries down to specific database shards based on the shard key value, to avoid scanning all tables in the DRDS instance.

For example, the shard key of the test table is c1. If the filter condition does not contain this shard key, full table scan is performed:

```
mysql> SELECT * FROM test WHERE c2 = 2;
+----+----+
| c1 | c2 |
+----+----+
| 2 | 2 |
+----+----+
1 row in set (0.05 sec)
```

The corresponding execution plan is as follows:

```
mysql> EXPLAIN SELECT * FROM test WHERE c2 = 2;
+-----+-----+
| GROUP_NAME          | SQL                                | PARAMS |
+-----+-----+
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0004_RDS | select `test`.`c1`,`c2` from `test` where (`test`.`c2` = 2) | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0007_RDS | select `test`.`c1`,`c2` from `test` where (`test`.`c2` = 2) | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0005_RDS | select `test`.`c1`,`c2` from `test` where (`test`.`c2` = 2) | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0002_RDS | select `test`.`c1`,`c2` from `test` where (`test`.`c2` = 2) | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0003_RDS | select `test`.`c1`,`c2` from `test` where (`test`.`c2` = 2) | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0006_RDS | select `test`.`c1`,`c2` from `test` where (`test`.`c2` = 2) | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0000_RDS | select `test`.`c1`,`c2` from `test` where (`test`.`c2` = 2) | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0001_RDS | select `test`.`c1`,`c2` from `test` where (`test`.`c2` = 2) | {} |
+-----+-----+
8 rows in set (0.00 sec)
```

The smaller the value range of the filter condition containing the shard key, the faster the DRDS query speed. For example, in the query on the test table, the filter condition contains the value range of the shard key c1:

```
mysql> SELECT * FROM test WHERE c1 > 1 AND c1 < 4;
+----+----+
| c1 | c2 |
+----+----+
| 2 | 2 |
| 3 | 3 |
+----+----+
2 rows in set (0.04 sec)
```

The corresponding execution plan is as follows:

```
mysql> EXPLAIN SELECT * FROM test WHERE c1 > 1 AND c1 < 4;
+-----+-----+
| GROUP_NAME          | SQL                                | PARAMS |
+-----+-----+
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0002_RDS | select `test`.`c1`,`c2` from `test` where ((`test`.`c1` > 1) AND (`test`.`c1` < 4)) | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0003_RDS | select `test`.`c1`,`c2` from `test` where ((`test`.`c1` > 1) AND (`test`.`c1` < 4)) | {} |
+-----+-----+
2 rows in set (0.00 sec)
```

The equivalence condition is executed faster than the range condition. For example:

```
mysql> SELECT * FROM test WHERE c1 = 2;
+----+----+
| c1 | c2 |
+----+----+
| 2 | 2 |
+----+----+
1 row in set (0.03 sec)
```

The corresponding execution plan is as follows:

```
mysql> EXPLAIN SELECT * FROM test WHERE c1 = 2;
+-----+-----+-----+
| GROUP_NAME          | SQL                                          | PARAMS |
+-----+-----+-----+
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0002_RDS | select `test`.`c1`,`c2` from `test` where (`test`.`c1` = 2) | {} |
+-----+-----+-----+
1 row in set (0.00 sec)
```

In addition, if you want to insert data into a table shard, the inserted field must contain a shard key.

For example, data inserted into the test table contains the shard key c1:

```
mysql> INSERT INTO test(c1,c2) VALUES(8,8);
Query OK, 1 row affected (0.07 sec)
```

Optimization of the number of returned rows for a query

When DRDS runs a query containing `LIMIT [offset,] row_count`, DRDS actually reads records before `offset` in order and directly discards them. In this way, when the value of `offset` is large, the query is slow even if the value of `row_count` is small. Take the following SQL statement as an example:

```
SELECT *
FROM sample_order
ORDER BY sample_order.id
LIMIT 10000, 2
```

Although only the 10000th and 10001st records are returned, it takes about 12 seconds to run the SQL statement because DRDS actually reads 10,002 records.

```
mysql> SELECT * FROM sample_order ORDER BY sample_order.id LIMIT 10000,2;
+-----+-----+-----+-----+-----+
| id      | sellerId | trade_id | buyer_id | buyer_nick |
+-----+-----+-----+-----+-----+
| 242012755468 | 1711939506 | 242012755467 | 244148116334 | zhangsan |
| 242012759093 | 1711939506 | 242012759092 | 244148138304 | wangwu |
+-----+-----+-----+-----+-----+
2 rows in set (11.93 sec)
```

The corresponding execution plan is as follows:

```
mysql> EXPLAIN SELECT * FROM sample_order ORDER BY sample_order.id LIMIT 10000,2;
+-----+-----+-----+
| GROUP_NAME | SQL | PARAMS |
+-----+-----+-----+
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0004_RDS | select `sample_order`.`id`,`sample_order`.`sellerId`,`sample_order`.`trade_id`,`sample_order`.`buyer_id`,`sample_order`.`buyer_nick` from `sample_order` order by `sample_order`.`id` asc limit 0,10002 | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0007_RDS | select `sample_order`.`id`,`sample_order`.`sellerId`,`sample_order`.`trade_id`,`sample_order`.`buyer_id`,`sample_order`.`buyer_nick` from `sample_order` order by `sample_order`.`id` asc limit 0,10002 | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0005_RDS | select `sample_order`.`id`,`sample_order`.`sellerId`,`sample_order`.`trade_id`,`sample_order`.`buyer_id`,`sample_order`.`buyer_nick` from `sample_order` order by `sample_order`.`id` asc limit 0,10002 | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0002_RDS | select `sample_order`.`id`,`sample_order`.`sellerId`,`sample_order`.`trade_id`,`sample_order`.`buyer_id`,`sample_order`.`buyer_nick` from `sample_order` order by `sample_order`.`id` asc limit 0,10002 | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0003_RDS | select `sample_order`.`id`,`sample_order`.`sellerId`,`sample_order`.`trade_id`,`sample_order`.`buyer_id`,`sample_order`.`buyer_nick` from `sample_order` order by `sample_order`.`id` asc limit 0,10002 | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0006_RDS | select `sample_order`.`id`,`sample_order`.`sellerId`,`sample_order`.`trade_id`,`sample_order`.`buyer_id`,`sample_order`.`buyer_nick` from `sample_order` order by `sample_order`.`id` asc limit 0,10002 | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0000_RDS | select `sample_order`.`id`,`sample_order`.`sellerId`,`sample_order`.`trade_id`,`sample_order`.`buyer_id`,`sample_order`.`buyer_nick` from `sample_order` order by `sample_order`.`id` asc limit 0,10002 | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0001_RDS | select `sample_order`.`id`,`sample_order`.`sellerId`,`sample_order`.`trade_id`,`sample_order`.`buyer_id`,`sample_order`.`buyer_nick` from `sample_order` order by `sample_order`.`id` asc limit 0,10002 | {} |
+-----+-----+-----+
8 rows in set (0.01 sec)
```

To optimize the preceding SQL statement, find the ID set, and use IN to match the actual records. The modified SQL query is as follows:

```
SELECT *
FROM sample_order o
WHERE o.id IN (
    SELECT id
    FROM sample_order
    ORDER BY id
    LIMIT 10000,2)
```

The purpose is to cache IDs in the memory first (on the premise that the number of IDs is small). If the shard key of the sample_order table is an ID, DRDS can also push down such an IN query to different database shards through rule-based calculation, avoiding full table scan and unnecessary network I/O. Check the result of the rewritten SQL query:

```
mysql> SELECT *
-> FROM sample_order o
-> WHERE o.id IN ( SELECT id FROM sample_order ORDER BY id LIMIT 10000,2);
+-----+-----+-----+-----+-----+
| id | sellerId | trade_id | buyer_id | buyer_nick |
+-----+-----+-----+-----+-----+
| 242012755468 | 1711939506 | 242012755467 | 244148116334 | zhangsan |
| 242012759093 | 1711939506 | 242012759092 | 244148138304 | wangwu |
+-----+-----+-----+-----+-----+
2 rows in set (1.08 sec)
```

The execution time is significantly reduced from 12 seconds to 1.08 seconds.

The corresponding execution plan is as follows:

```
mysql> EXPLAIN SELECT *
-> FROM sample_order o
-> WHERE o.id IN ( SELECT id FROM sample_order ORDER BY id LIMIT 10000,2);
+-----+-----+-----+-----+-----+-----+
| GROUP_NAME          | SQL                                                                 | PARAMS |
+-----+-----+-----+-----+-----+-----+
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0002_RDS | select `o`.`id`,`o`.`sellerId`,`o`.`trade_id`,`o`.`buyer_id`,`o`.`buyer_nick` from `sample_order` `o` where (`o`.`id` IN (10002)) | {} |
| SEQPERF_1478746391548CDTCSEQPERF_OXGJ_0001_RDS | select `o`.`id`,`o`.`sellerId`,`o`.`trade_id`,`o`.`buyer_id`,`o`.`buyer_nick` from `sample_order` `o` where (`o`.`id` IN (10001)) | {} |
+-----+-----+-----+-----+-----+-----+
2 rows in set (0.03 sec)
```

Grouping and sorting optimization

In DRDS, if an SQL query must use DISTINCT, GROUP BY, and ORDER BY at the same time, try to ensure that the fields after DISTINCT, GROUP BY, and ORDER BY are the same and the fields are shard keys. In this way, only a small amount of data is returned for the SQL query. This minimizes the network bandwidth consumed by distributed queries and removes the need to retrieve a large amount of data and sort the data in a temporary table, thereby maximizing the system performance.

11.14.4.3. JOIN query optimization

JOIN queries in DRDS are classified into push-down JOIN queries and non-push-down JOIN queries (distributed JOIN queries). The optimization policies for these two types of JOIN queries are different.

Optimize push-down JOIN queries

Push-down JOIN queries are classified into the following types:

- JOIN queries between single tables (non-partition tables).
- The tables involved in the JOIN query contain the shard key in the filter condition and use the same sharding algorithm (that is, the data calculated by the sharding algorithm is distributed to the same shard).
- Tables involved in the JOIN query use the shard key as the JOIN condition and use the same sharding algorithm.
- JOIN query between broadcast tables (or small table broadcast) and table shards.

In DRDS, optimize JOIN queries to push-down JOIN queries that can be executed on database shards.

Take a JOIN query between a broadcast table and table shards as an example. The broadcast table is used as the JOIN driving table (the left table in the JOIN query is called the driving table). The DRDS broadcast table stores the same data in each database shard. When the broadcast table is used as the JOIN driving table, the JOIN query between this broadcast table and table shards can be converted into single-database JOIN queries and combined for computing to improve query performance.

For example, a JOIN query is performed on the following three tables, among which the sample_area table is the broadcast table, and the sample_item and sample_buyer tables are table shards. The query execution time is about 15s:

```
mysql> SELECT sample_area.name
-> FROM sample_item i JOIN sample_buyer b ON i.sellerId = b.sellerId JOIN sample_area a ON b.province = a.id
-> WHERE a.id < 110107
-> LIMIT 0, 10;
+-----+
| name |
+-----+
| BJ |
+-----+
10 rows in set (14.88 sec)
```

If you adjust the JOIN query order and move the broadcast table to the farthest left as the JOIN driving table, the JOIN query is pushed down to a single database shard in the DRDS instance:

```
mysql> SELECT sample_area.name
-> FROM sample_area a JOIN sample_buyer b ON b.province = a.id JOIN sample_item i ON i.sellerId = b.sellerId
-> WHERE a.id < 110107
-> LIMIT 0, 10;
+-----+
| name |
+-----+
| BJ |
+-----+
10 rows in set (0.04 sec)
```

The query execution time decreases from 15 seconds to 0.04 seconds, which is a significant improvement to the query performance.

Notice The broadcast table achieves data consistency through the synchronization mechanism on database shards, with a latency of several seconds.

Optimize distributed JOIN queries

If a JOIN query cannot be pushed down (that is, the JOIN condition and filter condition do not contain the shard key), DRDS must complete part of the computing in the query. Such a query is a distributed JOIN query.

Tables in a distributed JOIN query are classified into two types based on the data size:

- **Small table:** A table that contains a small amount of data (less than 100 data records or less data than other tables) that is involved in JOIN computing after filtering.

- Large table: A table that contains a large amount of data (more than 100 data records or more data than other tables) that is involved in JOIN computing after filtering.

In most cases, Nested Loop and its derived algorithms are used in JOIN computing at the DRDS layer. If sorting is required for JOIN queries, the Sort Merge algorithm is used. When the Nested Loop algorithm is used, a smaller data size in the left table in a JOIN query indicates a smaller number of queries performed by DRDS on the right table. If the right table has indexes or contains a small amount of data, the JOIN query is even faster. Therefore, in DRDS, the left table of a distributed JOIN query is called the driving table. To optimize a distributed JOIN query, use a small table as the driving table and set as many filter conditions as possible for the driving table.

Take the following distributed JOIN query as an example. The query takes about 24 seconds:

```
mysql> SELECT t.title, t.price
-> FROM sample_order o,
-> ( SELECT * FROM sample_item i WHERE i.id = 242002396687 ) t
-> WHERE t.source_id = o.source_item_id AND o.sellerId < 1733635660;
+-----+-----+
| title          | price |
+-----+-----+
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
+-----+-----+
10 rows in set (23.79 sec)
```

The preceding JOIN query is an INNER JOIN query, with the actual size of the intermediate data involved in JOIN computing unknown. In this case, perform COUNT() on the o table and t table respectively to obtain the actual data size.

For the o table, o.sellerId < 1733635660 in the WHERE condition is only related to the o table. Then, extract and add it to the COUNT() condition of the o table. The following query result is returned:

```
mysql> SELECT COUNT(*) FROM sample_order o WHERE o.sellerId < 1733635660;
+-----+
| count(*) |
+-----+
| 504018 |
+-----+
1 row in set (0.10 sec)
```

The intermediate result of the o table contains about 500,000 records. Similarly, the t table is a subquery, which can be extracted directly for the COUNT() query:

```
mysql> SELECT COUNT(*) FROM sample_item i WHERE i.id = 242002396687;
+-----+
| count(*) |
+-----+
| 1 |
+-----+
1 row in set (0.01 sec)
```

The intermediate result of the t table contains only one record. Therefore, the o table is a large table and the t table is a small table. Use the small table as the driving table of the distributed JOIN query. The result of the adjusted JOIN query is as follows:

```
mysql> SELECT t.title, t.price
-> FROM ( SELECT * FROM sample_item i WHERE i.id = 242002396687 ) t,
-> sample_order o
-> WHERE t.source_id = o.source_item_id AND o.sellerId < 1733635660;
+-----+-----+
| title          | price |
+-----+-----+
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
| Sample Item for Distributed JOIN | 239.00 |
+-----+-----+
10 rows in set (0.15 sec)
```

The query execution time decreases from about 24 seconds to 0.15 seconds, with the query performance significantly improved.

11.14.4.4. Subquery optimization

When optimizing an SQL query that contains subqueries, push the subqueries down to database shards as much as possible to reduce the computing workload at the DRDS layer.

For this purpose, you can try two optimization methods:

- Rewrite subqueries into multi-table JOIN queries, and optimize the JOIN queries.
- Use the shard key in the JOIN condition or filter condition so that DRDS can push the query down to a specific database shard to avoid full table scan.

The following subquery is used as an example:

```
SELECT o.*
FROM sample_order o
WHERE NOT EXISTS
  (SELECT sellerId FROM sample_seller s WHERE o.sellerId = s.id)
```

Rewrite the subquery into a JOIN query query:

```
SELECT o.*
FROM sample_order o LEFT JOIN sample_seller s ON o.sellerId = s.id
WHERE s.id IS NULL
```

11.14.5. Select connection pools for an application

A database connection pool is used to manage database connections in a centralized manner, so as to improve application performance and reduce database loads.

- **Reuse resources:** Connections can be reused to avoid high performance overheads caused by frequent connection creation and release. Resource reuse can also improve the system stability.

- **Improve the system response efficiency:** After the connection initialization is complete, all requests can directly use the existing connections, which avoids the overheads of connection initialization and release and improves the system response efficiency.
- **Avoid connection leakage:** The connection pool forcibly revokes connections based on the preset de-allocation policy to avoid connection resource leakage.

We recommend that you use a connection pool to connect applications and databases for service operations. For Java programs, we recommend that you use the [Druid connection pool](#).

The following is an example of standard Druid Spring configuration:

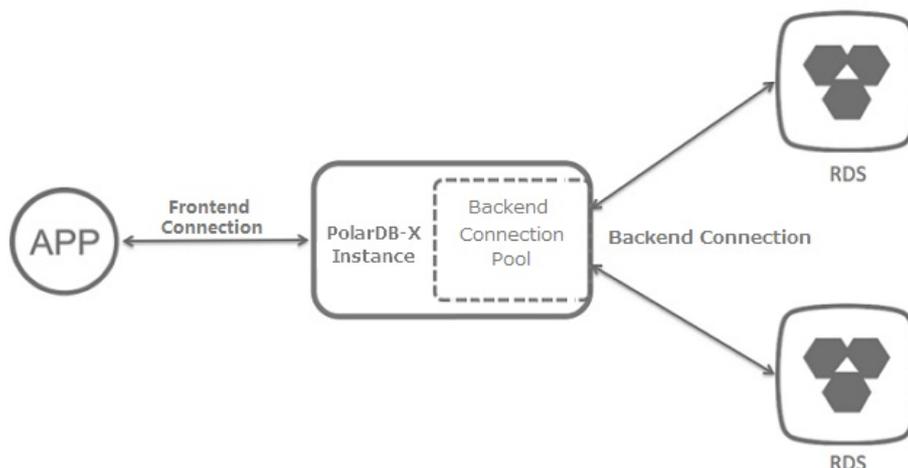
```
<bean id="dataSource" class="com.alibaba.druid.pool.DruidDataSource" init-method="init" destroy-method="close"
>
  <property name="driverClassName" value="com.mysql.jdbc.Driver" />
  <!-- Basic attributes URL, user, and password -->
  <property name="url" value="jdbc:mysql://ip:port/db? autoReconnect=true&rewriteBatchedStatements=true&
&socketTimeout=30000&connectTimeout=3000" />
  <property name="username" value="root" />
  <property name="password" value="123456" />
  <!-- Configure the initial size, minimum value, and maximum value -->
  <property name="maxActive" value="20" />
  <property name="initialSize" value="3" />
  <property name="minIdle" value="3" />
  <!-- maxWait indicates the time-out period for obtaining the connection -->
  <property name="maxWait" value="60000" />
  <!-- timeBetweenEvictionRunsMillis indicates the interval for detecting idle connections to be closed, in milliseconds
-->
  <property name="timeBetweenEvictionRunsMillis" value="60000" />
  <!-- minEvictableIdleTimeMillis indicates the minimum idle time of a connection in the connection pool, in milliseconds
ds-->
  <property name="minEvictableIdleTimeMillis" value="300000" />
  <!-- SQL statement used to check whether connections are available -->
  <property name="validationQuery" value="SELECT 'z'" />
  <!-- Whether to enable idle connection check -->
  <property name="testWhileIdle" value="true" />
  <!-- Whether to check the connection status before obtaining a connection -->
  <property name="testOnBorrow" value="false" />
  <!-- Whether to check the connection status before releasing a connection -->
  <property name="testOnReturn" value="false" />
</bean>
```

11.14.6. Connections to PolarDB-X instances

When an application connects to a PolarDB-X instance for operation, there are two types of connections from the perspective of the PolarDB-X instance:

- **Frontend connection:** a connection established by an application to the logical database in the PolarDB-X instance.
- **Backend connection:** a connection established by a node in a PolarDB-X instance to a physical database in a backend ApsaraDB RDS for MySQL instance.

PolarDB-X instance connection diagram



Frontend connection

Theoretically, the number of frontend connections is limited by the available memory size and the number of network connections to the nodes of the PolarDB-X instance. However, in actual application scenarios, when an application connects to a PolarDB-X instance, the nodes of the PolarDB-X instance usually manage a limited number of connections to perform requested operations, and do not maintain a large number of concurrent persistent connections (for example, tens of thousands of concurrent persistent connections). Therefore, the number of frontend connections that a PolarDB-X instance can accept can be considered to be unlimited.

Considering that the number of frontend connections is unlimited and a large number of idle connections are allowed, this method applies to scenarios where a large number of servers are deployed and need to maintain their connections to the PolarDB-X instance.

Note Although the number of frontend connections is considered as unlimited, operation requests obtained from frontend connections are actually executed by internal threads of the PolarDB-X instance through backend connections. Due to the limited number of internal threads and backend connections, the total number of concurrent requests that can be processed by the PolarDB-X instance is limited.

Backend connection

Each node of a PolarDB-X instance creates a backend connection pool to automatically manage and maintain the backend connections to the physical databases in the ApsaraDB RDS for MySQL instance. Therefore, the maximum number of connections in the backend connection pool of a PolarDB-X instance is directly related to the maximum number of connections supported by the ApsaraDB RDS for MySQL instance. You can use the following formula to calculate the maximum number of connections in the backend connection pool of a PolarDB-X instance:

Maximum number of connections in a backend connection pool of a PolarDB-X instance = FLOOR (Maximum number of connections in an ApsaraDB RDS for MySQL instance/Number of physical database shards in the ApsaraDB RDS for MySQL instance/Number of nodes on the PolarDB-X instance)

For example, a user has purchased an ApsaraDB RDS for MySQL instance and a PolarDB-X instance of the following types:

- The ApsaraDB RDS for MySQL instance has eight physical database shards, four cores, and 16 GB memory, supporting a maximum number of 4,000 connections.
- The PolarDB-X dedicated instance has 32 cores and 32 GB memory, with each PolarDB-X node having two cores and 2 GB memory (that is, the instance has 16 PolarDB-X nodes).

You can use the following formula to calculate the maximum number of connections in the backend connection pool of the PolarDB-X instance:

Maximum number of connections in the backend connection pool of the PolarDB-X instance = $\text{FLOOR}(4000/8/16) = \text{FLOOR}(31.25) = 31$

Note

- The calculation result of the preceding formula is the maximum number of connections in the backend connection pool of the PolarDB-X instance. In actual use, to reduce the connection pressure on the ApsaraDB RDS for MySQL instance, the PolarDB-X instance adjusts the maximum number of connections in the backend connection pool to make it smaller than the upper limit.
- We recommend that you create databases in a PolarDB-X instance on a dedicated ApsaraDB RDS for MySQL instance. Do not create databases for other applications or PolarDB-X instances on the dedicated ApsaraDB RDS for MySQL instance.

Relationship between frontend and backend connections

After an application establishes frontend connections to a PolarDB-X instance and sends SQL statement execution requests, the PolarDB-X nodes process the requests asynchronously and obtain backend connections through the internal backend connection pool, and then run optimized SQL statements on one or more physical databases.

PolarDB-X nodes process requests asynchronously and frontend connections are not bound to backend connections. Therefore, a small number of backend connections can process a large number of requests for short transactions and simple queries from many concurrent frontend connections. This is why you need to focus on the queries per second (QPS) in PolarDB-X, rather than the number of concurrent connections.

Although the number of frontend connections is considered to be unlimited, the maximum number of connections maintained in the backend connection pool of a PolarDB-X instance is limited. For more information, see "Backend connections." Therefore, note the following points in actual application scenarios:

- Avoid long or large transactions in applications. These transactions occupy many or even all backend connections when they are not committed or rolled back for a long time, which reduces the overall concurrent processing capability and increases the response time (RT).
- Monitor and optimize or remove slow SQL queries run in the PolarDB-X instance, to prevent them from occupying too many backend connections. Otherwise, the PolarDB-X instance or the ApsaraDB RDS for MySQL instance is under greater processing pressure, which may lead to reduced concurrent processing capability, increased RT, or higher SQL execution failure rate due to execution timeout. For troubleshooting and optimization of slow SQL queries, see [Troubleshoot slow SQL statements in PolarDB-X](#) and [Overview](#).
- Under normal use of connections and execution of queries, if the maximum number of connections in the backend connection pool of the PolarDB-X instance is reached, contact Customer Services for assistance.

11.14.7. Perform instance upgrade

Database performance can be measured by the response time (RT) and queries per second (QPS). RT reflects the performance of a single SQL statement. This type of performance problem can be solved through SQL optimization. PolarDB-X upgrade expands the capacity to improve performance, and is suitable for database access services with low latency and high QPS.

The performance of a PolarDB-X instance depends on the performance of PolarDB-X and ApsaraDB RDS for MySQL. Insufficient performance of any PolarDB-X or ApsaraDB RDS for MySQL node can create a bottleneck in the overall performance. This topic describes how to observe the performance metrics of a PolarDB-X instance and upgrade the PolarDB-X instance to solve the performance bottleneck. For more information about how to determine the performance of an ApsaraDB RDS for MySQL instance and upgrade the ApsaraDB RDS for MySQL instance, see the ApsaraDB RDS for MySQL documentation.

Determine the performance bottleneck of a PolarDB-X instance

The QPS and CPU performance of a PolarDB-X instance are in positive correlation. When a PolarDB-X instance encounters a performance bottleneck, the CPU utilization of the PolarDB-X instance remains high.

Observe the CPU utilization

1. On the **Basic Information** page of the PolarDB-X instance, choose **Monitoring and Alerts > Instance Monitoring** from the left-side navigation pane.
2. On the Instance Monitoring page, select a monitoring dimension and the corresponding metrics to view details.

If the CPU utilization exceeds 90% or remains above 80%, the PolarDB-X instance faces a performance bottleneck. If there is no bottleneck for the ApsaraDB RDS for MySQL instance, the current PolarDB-X instance specifications cannot meet the QPS performance requirements of the business. In this case, the PolarDB-X instance needs to be upgraded.

For more performance-related service monitoring scenarios and methods for configuring the PolarDB-X CPU utilization alert.

Upgrade PolarDB-X

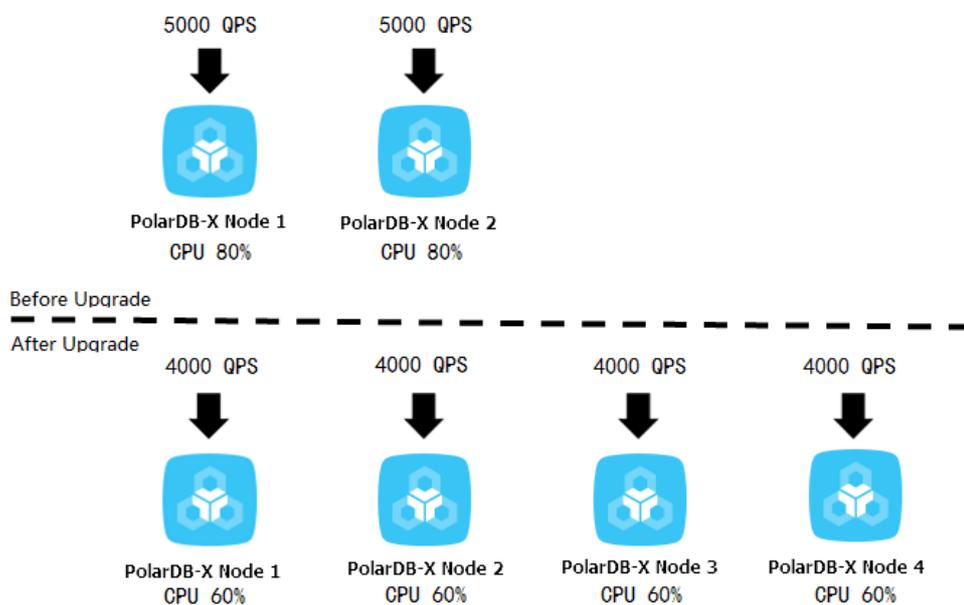
QPS is an important metric for determining whether the PolarDB-X instance specifications can meet the business requirements. Each type of instance specifications corresponds to a reference QPS value.

Some special SQL statements require more computing (such as temporary table sorting and aggregate computing) in PolarDB-X. In this case, the QPS supported by each PolarDB-X instance is lower than the standard value in its type.

PolarDB-X upgrade improves the processing performance of a PolarDB-X instance by adding nodes to share the QPS. As PolarDB-X nodes are stateless, this upgrade method linearly improves the performance of PolarDB-X instances.

For example, service A requires QPS of about 15 thousand. The current PolarDB-X instance has a 4-core virtual CPU (vCPU), 4 GB memory, and two nodes, supporting QPS of only 10 thousand. After finding that the CPU utilization of the PolarDB-X instance remains high, we upgraded the instance to 8-core vCPU and 8 GB memory, with each node handling about 4,000 QPS. Then, the performance meets service requirements, and the CPU utilization also drops to a reasonable level, as shown in the following figure.

PolarDB-X upgrade

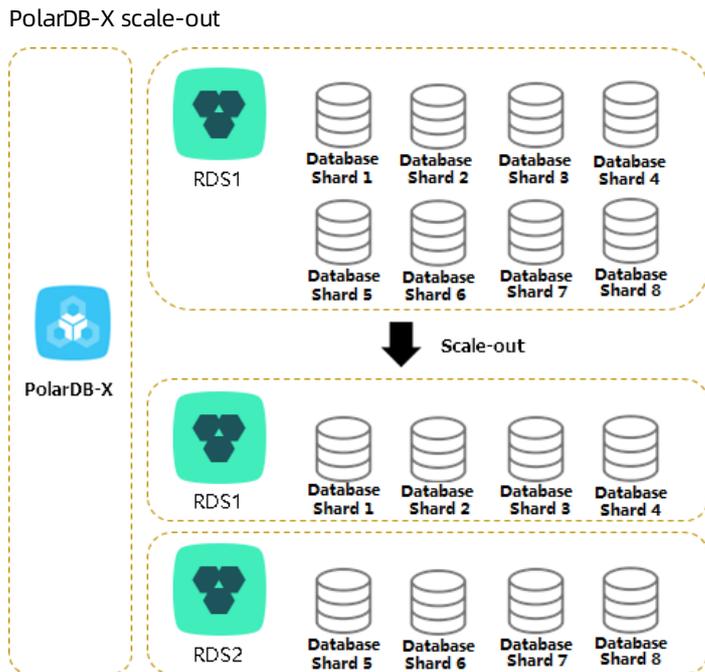


For more information about how to upgrade a PolarDB-X instance, see [Change specifications](#).

11.14.8. Perform scale-out

In PolarDB-X, smooth scale-out improves the overall performance by increasing the number of ApsaraDB RDS for MySQL instances. You can increase the number of ApsaraDB RDS for MySQL instances through PolarDB-X smooth scale-out to increase the PolarDB-X database capacity when the following conditions are met: 1. The input/output operations per second (IOPS), CPU utilization, disk space, and other metrics of the ApsaraDB RDS for MySQL instance reach their bottlenecks. 2. The bottlenecks cannot be removed through SQL optimization or ApsaraDB RDS for MySQL upgrade (for example, the disk has been upgraded to the top configuration).

PolarDB-X smooth scale-out reduces the pressure on the original ApsaraDB RDS for MySQL instance by migrating database shards to the new ApsaraDB RDS for MySQL instance. For example, before scale-out, all the eight databases are deployed in one ApsaraDB RDS for MySQL instance. After scale-out, the eight databases are deployed in two ApsaraDB RDS for MySQL instances, and the pressure on a single ApsaraDB RDS for MySQL instance is significantly reduced, as shown in the following figure.



After multiple scale-out operations, if the number of ApsaraDB RDS for MySQL instances is equal to the number of database shards, you need to create another PolarDB-X instance and ApsaraDB RDS for MySQL databases with the expected capacity, and then migrate data to further increase the data capacity. This process is complex. We recommend that you consider the data growth expected in the next two to three years and plan the number of ApsaraDB RDS for MySQL instances properly when creating a PolarDB-X database.

Determine whether scale-out is required

You can determine whether PolarDB-X smooth scale-out is required based on three ApsaraDB RDS for MySQL metrics: IOPS, CPU utilization, and disk space. You can view these metrics in the ApsaraDB RDS for MySQL console. For more information, see the ApsaraDB RDS for MySQL documentation.

IOPS and CPU utilization

If you find that the IOPS or CPU utilization remains above 80% for a long time or you frequently receive alerts, follow these steps:

1. Optimize SQL statements. Generally, you can solve the high CPU utilization problem by this method.
2. If the problem persists, upgrade the ApsaraDB RDS for MySQL instance. For more information, see the ApsaraDB RDS for MySQL documentation.
3. When the CPU utilization or IOPS exceeds the threshold, you can set read-only databases to share the load on the primary database. However, read/write splitting affects read consistency. For more information, see the [Read/write splitting](#) documentation.
4. If the problem persists, scale out the PolarDB-X instance.

Disk space

ApsaraDB RDS for MySQL has the following types of disk space:

1. Data space: the space occupied by data. The space usage continues increasing as more data is inserted. We recommend that you keep the remaining disk space above 30%.
2. System file space: the space occupied by shared tables and error log files.
3. Binary log file space: the space occupied by binary logs generated during database operation. The more update transactions there are, the larger the occupied space is.

Whether scale-out is required depends on the data space. When the data space is about to or expected to exceed the disk capacity, you can distribute the data to the databases on multiple ApsaraDB RDS for MySQL instances through scale-out.

Scale-out risks and precautions

PolarDB-X scale-out consists of four steps: **configuration** > **migration** > **switchover** > **cleanup**. For more information, see the [Perform smooth scale-out](#) documentation.

Note the following points before scale-out:

- To reduce the pressure of read operations on the source ApsaraDB RDS for MySQL instance, perform scale-out when the load on the source ApsaraDB RDS for MySQL instance is low.
- During scale-out, do not submit data definition language (DDL) tasks in the console or connect to the PolarDB-X instance to directly run DDL SQL statements. Otherwise, the scale-out task may fail.
- Scale-out requires that the source database table have a primary key. If the source database does not have a primary key, add one first.
- During scale-out, the read and write traffic is switched to the new ApsaraDB RDS for MySQL instance. The switchover process takes three to five minutes. We recommend that you perform a switchover during off-peak hours.
- Scale-out does not affect the PolarDB-X instance before the switchover. Therefore, you can cancel the scale-out through rollback before the switchover.
- Scale-out creates pressure on databases. We recommend that you perform this operation during off-peak hours.

11.14.9. Troubleshoot slow SQL statements in DRDS

11.14.9.1. Details about a low SQL statement

PolarDB-X defines an SQL statement that takes more than 1 second to run as a slow SQL statement. Slow SQL statements in PolarDB-X are classified into slow logical SQL statements and slow physical SQL statements. In PolarDB-X, an SQL statement is run step by step on PolarDB-X and ApsaraDB RDS for MySQL nodes. Large execution loss on any node will result in slow SQL statements.

- Slow logical SQL statements are slow SQL statements sent by an application to PolarDB-X.
- Slow physical SQL statements are slow SQL statements sent by PolarDB-X to ApsaraDB RDS for MySQL.

Syntax

```
SHOW FULL {SLOW | PHYSICAL_SLOW} [WHERE where_condition]
      [ORDER BY col_name [ASC | DESC], ...]
      [LIMIT {[offset,] row_count | row_count OFFSET offset}]
```

Description

The `SHOW FULL SLOW` command shows slow logical SQL statements, that is, SQL statements sent by an application to PolarDB-X.

The result set of the `SHOW FULL SLOW` command contains the following columns:

- **TRACE_ID**: the unique identifier of the SQL statement. A logical SQL statement and the physical SQL statements generated by the logical SQL statement have the same TRACE_ID. The TRACE_ID is also sent as a comment to ApsaraDB RDS for MySQL.
- **HOST**: the IP address of the client that sends the SQL statement.

 **Notice** The client IP address may not be obtained when the network type is Virtual Private Cloud (VPC).

- **START_TIME**: the time when PolarDB-X starts running the SQL statement.
- **EXECUTE_TIME**: the time consumed by PolarDB-X to run the SQL statement.
- **AFFECT_ROW**: the number of records returned or the number of rows affected by the SQL statement.
- **SQL**: the statement that is run.

The `SHOW FULL PHYSICAL_SLOW` command shows the slow physical SQL statements, that is, SQL statements sent by PolarDB-X to ApsaraDB RDS for MySQL.

The result set of `SHOW FULL PHYSICAL_SLOW` contains the following columns:

- **TRACE_ID**: the unique identifier of the SQL statement. A logical SQL statement and the physical SQL statements generated by the logical SQL statement have the same TRACE_ID. The TRACE_ID is also sent as a comment to ApsaraDB RDS for MySQL.
- **GROUP_NAME**: the name of a database group. Grouping aims to manage multiple groups of databases with identical data, such as the primary and secondary databases after data replication through ApsaraDB RDS for MySQL, which are mainly used for read/write splitting and primary/secondary switchover.
- **DBKEY_NAME**: the name of the database shard on which the SQL statement is run.
- **START_TIME**: the time when PolarDB-X starts running the SQL statement.
- **EXECUTE_TIME**: the time consumed by PolarDB-X to run the SQL statement.
- **SQL_EXECUTE_TIME**: the time consumed by PolarDB-X to call ApsaraDB RDS for MySQL to run this SQL statement.
- **GET_LOCK_CONNECTION_TIME**: the time that PolarDB-X takes to get connections from the connection pool. If the value is large, the ApsaraDB RDS for MySQL connections have been exhausted. This is typically due to a large number of slow SQL statements. You can log on to the corresponding ApsaraDB RDS for MySQL instance and run `SHOW PROCESSLIST` for troubleshooting.
- **CREATE_CONNECTION_TIME**: the time consumed by PolarDB-X to establish a connection to ApsaraDB RDS for MySQL. If the value is large, it is largely because the ApsaraDB RDS for MySQL instance is overloaded or faulty.
- **AFFECT_ROW**: the number of records returned or the number of rows affected by the SQL statement.
- **SQL**: the statement that is run.

Example 1

The following example describes how to locate the execution of a slow SQL statement on PolarDB-X and between PolarDB-X and ApsaraDB RDS for MySQL.

1. You can use certain conditions, such as the execution time and SQL string match, to obtain the specified slow SQL statement:

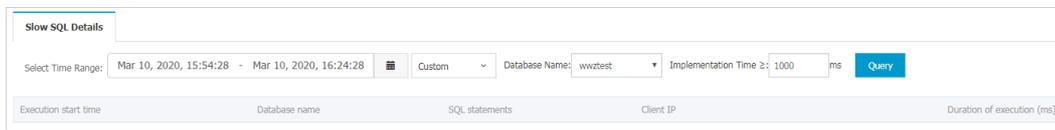
```
mysql> show full slow where `SQL` like '%select sleep(50)%';
+-----+-----+-----+-----+-----+-----+
| TRACE_ID | HOST | START_TIME | EXECUTE_TIME | AFFECT_ROW | SQL |
+-----+-----+-----+-----+-----+-----+
| ae0e565b8c00000 | 127.0.0.1 | 2017-03-29 19:28:43.028 | 50009 | 1 | select sleep(50) |
+-----+-----+-----+-----+-----+-----+
1 row in set (0.02 sec)
```

2. Based on the TRACE_ID of the slow logical SQL statement, run `SHOW FULL PHYSICAL_SLOW` to obtain the physical execution information of this SQL statement.

```
mysql> show full physical_slow where trace_id = 'ae0e565b8c00000';
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| TRACE_ID | GROUP_NAME          | DBKEY_NAME          | START_TIME | EXECUTE_TIME | SQL_EXECUTE_TIME | GETLOCK_CONNECTION_TIME | CREATE_CONNECTION_TIME | AFFECT_ROW | SQL |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| ae0e565b8c00000 | PRIV_TEST_1489167306631PJAFPRIV_TEST_APKK_0000_RDS | rdso6g5b6206sdq832ow_priv_test_apkk_0000_nfup | 2017-03-29 19:27:53.02 | 50001 | 50001 | 0 | 0 | 1 | select sleep(50) |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.01 sec)
```

- In the SQL statement details and slow SQL statement records of the ApsaraDB RDS for MySQL instance, you can query the execution information of this SQL statement on the ApsaraDB RDS for MySQL instance based on TRACE_ID.

Slow query logs



Example 2

This example describes how to locate the original SQL statement in PolarDB-X based on the slow SQL statement located in ApsaraDB RDS for MySQL.

- Based on the slow SQL query log in ApsaraDB RDS for MySQL, TRACE_ID of the slow SQL statement is ae0e55660c00000.
- Based on the TRACE_ID obtained in Step 1, run `SHOW FULL PHYSICAL_SLOW` to obtain the physical execution information of this SQL statement.

```
mysql> show full physical_slow where trace_id = 'ae0e55660c00000';
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| TRACE_ID | GROUP_NAME          | DBKEY_NAME          | START_TIME | EXECUTE_TIME | SQL_EXECUTE_TIME | GETLOCK_CONNECTION_TIME | CREATE_CONNECTION_TIME | AFFECT_ROW | SQL |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| ae0e55660c00000 | PRIV_TEST_1489167306631PJAFPRIV_TEST_APKK_0000_RDS | rdso6g5b6206sdq832ow_priv_test_apkk_0000_nfup | 2017-03-29 19:27:37.308 | 10003 | 10001 | 0 | 0 | 1 | select sleep(10) |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.02 sec)
```

11.14.9.2. Locate slow SQL statements

Generally, you can locate a slow SQL statement in two ways: Obtain historical information about slow SQL statements from slow SQL statement records, or run `SHOW PROCESSLIST` to display the real-time execution information about slow SQL statements.

You can troubleshoot slow SQL statements as follows:

- Locate slow SQL statements.
- Locate nodes with performance loss.
- Troubleshoot the performance loss.

Note During troubleshooting, we recommend that you use the MySQL command line `mysql -hIP -PPORT -uUSER -pPASSWORD -c` to create the connection. Be sure to add `-c` to prevent the MySQL client from filtering out the comments (default operation) and therefore affecting the execution of HINT.

- View slow SQL statement records

Run the following command to query top 10 slow SQL statements. This command can query logical SQL statements in PolarDB-X. One logical SQL statement corresponds to SQL statements of one or more databases or tables of the ApsaraDB RDS for MySQL instance. For more information, see [Details about a low SQL statement](#).

```
mysql> SHOW SLOW limit 10;
+-----+-----+-----+-----+-----+-----+
| TRACE_ID | HOST | START_TIME | EXECUTE_TIME | AFFECT_ROW | SQL |
+-----+-----+-----+-----+-----+-----+
| ac3133132801001 | xx.xxx.xx.97 | 2017-03-06 15:48:32.330 | 900392 | -1 | select detail_url, sum(price) from t_item group by detail_url; |
.....
+-----+-----+-----+-----+-----+-----+
10 rows in set (0.01 sec)
```

- View real-time SQL execution information

If the execution of an SQL statement is slow in the current server, run **SHOW PROCESSLIST** to view the real-time SQL execution information in the current PolarDB-X database. The value in the TIME column indicates how long the current SQL statement has been run.

```
mysql> SHOW PROCESSLIST WHERE COMMAND != 'Sleep';
+-----+-----+-----+-----+-----+-----+-----+
| ID | USER | DB | COMMAND | TIME | STATE | INFO |
| ROWS_SENT | ROWS_EXAMINED | ROWS_READ |
+-----+-----+-----+-----+-----+-----+-----+
| 0-0-352724126 | ifisibhk0 | test_123_wvvp_0000 | Query | 13 | Sending data | /*DRDS /42.120.74.88/ac47e5a72801000/ */select `t_item`.`detail_url`,SUM(`t_item`.`price`) from `t_i | NULL | NULL | NULL |
| 0-0-352864311 | cowxhthg0 | NULL | Binlog Dump | 17 | Master has sent all binlog to slave; waiting for binlog to be updated | NULL | NULL | NULL | NULL |
| 0-0-402714795 | ifisibhk0 | test_123_wvvp_0005 | Alter | 114 | Sending data | /*DRDS /42.120.74.88/ac47e5a72801000/ */ALTER TABLE `Persons` ADD `Birthday` date | NULL | NULL | NULL |
.....
+-----+-----+-----+-----+-----+-----+-----+
12 rows in set (0.03 sec)
```

The following describes each column:

- ID: the ID of the connection.
- USER: the user name of the database shard in which this SQL statement is run.
- DB: the specified database. If no database is specified, the value is NULL.
- COMMAND: the type of the command being executed. SLEEP indicates an idle connection. For more information about other commands, see [MySQL thread information documentation](#).
- TIME: the elapsed execution time of the SQL statement, in seconds.
- STATE: the current execution status. For more information, see [MySQL thread status documentation](#).
- INFO: the SQL statement being executed. The SQL statement may be too long to be displayed completely. You can derive the complete SQL statement based on information such as service parameters.

In the current example, the following slow SQL statement is identified:

```
ALTER TABLE `Persons` ADD `Birthday` date
```

11.14.9.3. Locate nodes with performance loss

When you locate a slow SQL statement in slow SQL statement records or real-time SQL execution information, you can run the TRACE command to trace the running time of the SQL statement in PolarDB-X and ApsaraDB RDS for MySQL to locate the bottleneck.

The TRACE command actually runs the SQL statement, records the time consumed on all nodes, and returns the execution result. For more information about TRACE and other control commands, see [Help statements](#).

Note The PolarDB-X TRACE command needs to maintain the context information of the connection. Some GUI clients may use connection pools, which results in command exceptions. Therefore, we recommend that you use the MySQL command line to run the TRACE command.

Run the following command for the identified slow SQL statement:

```
mysql> trace select detail_url, sum(distinct price) from t_item group by detail_url;
+-----+-----+
| detail_url | sum(price) |
+-----+-----+
| www.xxx.com | 1084326800.00 |
| www.xx1.com | 1084326800.00 |
| www.xx2.com | 1084326800.00 |
| www.xx3.com | 1084326800.00 |
| www.xx4.com | 1084326800.00 |
| www.xx5.com | 1084326800.00 |
.....
+-----+-----+
1 row in set (7 min 2.72 sec)
```

After the TRACE command is run, run SHOW TRACE to view the result. You can identify the bottleneck of the slow SQL statement based on the time consumption of each component.

```
mysql> SHOW TRACE;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+
| ID | TIMESTAMP | TYPE | GROUP_NAME | DBKEY_NAME | TIME_COST(MS) | CONNECTION_TIME_COST(MS) | ROWS | STATEMENT | CON |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+
| 0 | 0.000 | Optimize | DRDS | DRDS | 2 | 0.00 | 0 | select detail_url, sum(price) from t_item group by detail_url | NULL |
| 1 | 423507.342 | Merge Sorted | DRDS | DRDS | 411307 | 0.00 | 8 | Using Merge Sorted, Order By (`t_item`.`detail_url` asc) | NULL |
| 2 | 2.378 | Query | TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0003_RDS | rdso6g5b6206sdq832ow_test_123_wvvp_0003_hbpz | 15 | 1.59 | 1 | select `t_item`.`detail_url`,SUM(distinct `t_item`.`price`) from `t_item` group by `t_item`.`detail_url` order by `t_item`.`detail_url` asc | NULL |
| 3 | 2.731 | Query | TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0000_RDS | rdso6g5b6206sdq832ow_test_123_wvvp_0000_hbpz | 11 | 1.78 | 1 | select `t_item`.`detail_url`,SUM(distinct `t_item`.`price`) from `t_item` group by `t_item`.`detail_url` order by `t_item`.`detail_url` asc | NULL |
| 4 | 2.933 | Query | TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0004_RDS | rdso6g5b6206sdq832ow_test_123_wvvp_0004_hbpz | 15 | 1.48 | 1 | select `t_item`.`detail_url`,SUM(distinct `t_item`.`price`) from `t_item` group by `t_item`.`detail_url` order by `t_item`.`detail_url` asc | NULL |
| 5 | 3.111 | Query | TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0001_RDS | rdso6g5b6206sdq832ow_test_123_wvvp_0001_hbpz | 15 | 1.56 | 1 | select `t_item`.`detail_url`,SUM(distinct `t_item`.`price`) from `t_item` group by `t_item`.`detail_url` order by `t_item`.`detail_url` asc | NULL |
| 6 | 3.323 | Query | TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0007_RDS | rdso6g5b6206sdq832ow_test_123_wvvp_0007_hbpz | 15 | 1.54 | 1 | select `t_item`.`detail_url`,SUM(distinct `t_item`.`price`) from `t_item` group by `t_item`.`detail_url` order by `t_item`.`detail_url` asc | NULL |
| 7 | 3.496 | Query | TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0006_RDS | rdso6g5b6206sdq832ow_test_123_wvvp_0006_hbpz | 18 | 1.30 | 1 | select `t_item`.`detail_url`,SUM(distinct `t_item`.`price`) from `t_item` group by `t_item`.`detail_url` order by `t_item`.`detail_url` asc | NULL |
| 8 | 3.505 | Query | TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0005_RDS | rdso6g5b6206sdq832ow_test_123_wvvp_0005_hbpz | 423507 | 1.97 | 1 | select `t_item`.`detail_url`,SUM(distinct `t_item`.`price`) from `t_item` group by `t_item`.`detail_url` order by `t_item`.`detail_url` asc | NULL |
| 9 | 3.686 | Query | TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0002_RDS | rdso6g5b6206sdq832ow_test_123_wvvp_0002_hbpz | 14 | 1.47 | 1 | select `t_item`.`detail_url`,SUM(distinct `t_item`.`price`) from `t_item` group by `t_item`.`detail_url` order by `t_item`.`detail_url` asc | NULL |
| 10 | 423807.906 | Aggregate | DRDS | DRDS | 1413 | 0.00 | 1 | Aggregate Function (SUM(`t_item`.`price`)), Group By (`t_item`.`detail_url` asc) | NULL |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+
11 rows in set (0.01 sec)
```

In the returned results of SHOW TRACE, you can determine which node has a long execution time based on the values (in milliseconds) in the TIME_COST column. You can also see the corresponding GROUP_NAME (that is, the PolarDB-X or ApsaraDB RDS for MySQL node) and the STATEMENT column information (that is, the SQL statement being executed). By checking whether the value of GROUP_NAME is PolarDB-X, you can determine whether the slow node exists in PolarDB-X or ApsaraDB RDS for MySQL.

According to the preceding results, the Merge Sorted action on the PolarDB-X node and the TEST_123_1488766060743ACTJSANGUAN_TEST_123_WVVP_0005_RDS node of ApsaraDB RDS for MySQL take a lot of time.

11.14.9.4. Troubleshoot the performance loss

Slow nodes may exist on the PolarDB-X or ApsaraDB RDS for MySQL instance. Troubleshoot the fault accordingly after the cause is determined.

Solution for slow PolarDB-X nodes

When the `GROUP_NAME` of a slow node is in the PolarDB-X instance, check whether time-consuming computing operations such as Merge Sorted, Temp Table Merge, and Aggregate exist during SQL statement execution. If so, rectify it. For more information, see [Overview](#).

Solution for slow ApsaraDB RDS for MySQL nodes

When the slow node is on the ApsaraDB RDS for MySQL instance, check the execution plan of this SQL statement on the ApsaraDB RDS for MySQL instance.

In PolarDB-X, you can run `/*! TDDL:node={GROUP_NAME}*/EXPLAIN` to check the SQL execution plan of an ApsaraDB RDS for MySQL instance. The execution plan displays the SQL execution process information, including inter-table association and index information.

The detailed process is as follows:

1. Based on `GROUP_NAME`, assemble the HINT: `/*! TDDL:node='TEST_123_1488766060743ACTJSANGUAN_TEST_123_WWVP_0005_RDS'*/`.
2. Combine the assembled HINT and the statement prefixed by EXPLAIN to form a new SQL statement and run it. The EXPLAIN command does not actually run. It only displays the execution plan of the SQL statement.

The following example describes how to query the execution plan of the identified slow node.

```
mysql> /*! TDDL:node='TEST_123_1488766060743ACTJSANGUAN_TEST_123_WWVP_0005_RDS'*/EXPLAIN select `t_item`.`detail_url`,SUM(distinct `t_item`.`price`) from `t_item` group by `t_item`.`detail_url` order by `t_item`.`detail_url` asc;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | t_item | ALL | NULL | NULL | NULL | NULL | 1322263 | Using temporary; Using filesort |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.01 sec)
```

When the preceding SQL statement is run in ApsaraDB RDS for MySQL, the message `Using temporary; Using filesort` is returned. It indicates that slow SQL statement execution is caused by improper use of the index. In this case, you can correct the index and run the SQL statement again.

11.14.10. Handle DDL exceptions

When you run any data definition language (DDL) commands of PolarDB-X, PolarDB-X performs the corresponding DDL operation on all table shards.

Failures can be divided into two types:

1. A DDL statement fails to be executed in a database shard. DDL execution failure in any database shard may result in inconsistent table shard structures.
2. The system does not respond for a long time after a DDL statement is executed. When you perform a DDL statement on a large table, the system may make no response for a long time due to the long execution time of the DDL statement in a database shard.

Execution failures in database shards may occur for various reasons. For example, the table you want to create already exists, the column you want to add already exists, or the disk space is insufficient.

No response for a long time is generally caused by the long execution time of a DDL statement in a database shard. Taking ApsaraDB RDS for MySQL as an example, the DDL execution time depends mostly on whether the operation is an in-place (directly modifying the source table) or copy (copying data in the table) operation. An in-place operation only requires modification of metadata, while a copy operation reconstructs the whole table and also involves log and buffer operations.

To determine whether a DDL operation is an in-place or copy operation, you can view the returned value of "rows affected" after the operation is completed.

Example:

- Change the default value of a column (this operation is very fast and does not affect the table data at all):

```
Query OK, 0 rows affected (0.07 sec)
```

- Add an index (this operation takes some time, but "0 rows affected" indicates that the table data is not replicated):

```
Query OK, 0 rows affected (21.42 sec)
```

- Change the data type of column (this operation takes a long time and reconstructs all data rows in the table):

```
Query OK, 1671168 rows affected (1 min 35.54 sec)
```

Therefore, before executing a DDL operation on a large table, perform the following steps to determine whether the operation is a fast or slow operation:

1. Copy the table structure to generate a cloned table.
2. Insert some data.
3. Perform the DDL operation on the cloned table.
4. Check whether the value of "rows affected" is 0 after the operation is completed. A non-zero value means that this operation reconstructs the entire table. In this case, you need to perform this operation in off-peak hours.

Solution for failures

PolarDB-X DDL operations distribute all SQL statements to all database shards for parallel execution. Execution failure on any database shard does not affect the execution on other database shards. In addition, PolarDB-X provides the CHECK TABLE command to check the structure consistency of the table shards. Therefore, failed DDL operations can be performed again, and errors reported on database shards on which the operations have been executed do not affect the execution on other database shards. Make sure that all table shards ultimately have the same structure.

Procedure for handling DDL operation failures

1. Run the CHECK TABLE command to check the table structure. If the returned result contains only one row and the status is normal, **the table statuses are consistent**. In this case, go to Step 2. Otherwise, go to Step 3.
2. Run the SHOW CREATE TABLE command to check the table structure. If the displayed table structure is the same as the expected structure after the DDL statement is run, the DDL statement is run. Otherwise, go to Step 3.
3. Run the SHOW PROCESSLIST command to check the statuses of all SQL statements being executed. If any ongoing DDL operations are detected, wait until these operations are completed, and then perform Steps 1 and 2 to check the table structure. Otherwise, go to Step 4.
4. Perform the DDL operation again on PolarDB-X. If the Lock conflict error is reported, go to Step 5. Otherwise, go to Step 3.
5. Run the RELEASE DBLOCK command to release the DDL operation lock, and then go to Step 4.

The procedure is as follows:

1. Check the table structure consistency

Run the CHECK TABLE command to check the table structure. When the returned result contains only one row and the displayed status is OK, **the table structures are consistent**.

 **Notice** If no result is returned after you run CHECK TABLE, retry by using the CLI.

```
mysql> check table `xxxx` ;
+-----+-----+-----+-----+
| TABLE      | OP  | MSG_TYPE | MSG_TEXT |
+-----+-----+-----+-----+
| TDDL5_APP.xxxx | check | status  | OK      |
+-----+-----+-----+-----+
1 row in set (0.05 sec)
```

2. Check the table structure

Run the SHOW CREATE TABLE command to check the table structure. If table structures are consistent and correct, the DDL statement has been run.

```
mysql> show create table `xxxx` ;
+-----+-----+-----+-----+
| Table | Create Table
+-----+-----+-----+-----+
| xxxx | CREATE TABLE `xxxx` (
  `id` int(11) NOT NULL DEFAULT '0',
  `NAME` varchar(1024) NOT NULL DEFAULT '',
  PRIMARY KEY (`id`)
) ENGINE=InnoDB DEFAULT CHARSET=utf8 dbpartition by hash(`id`) tpartition by hash(`id`) tpartitions 3
+-----+-----+-----+-----+
1 row in set (0.05 sec)
```

3. Check the SQL statements being executed.

If some DDL statement executions are slow and no response is received for a long time, you can run the SHOW PROCESSLIST command to check the status of all SQL statements being executed.

```
mysql> SHOW PROCESSLIST WHERE COMMAND != 'Sleep';
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| ID      | USER      | DB      | COMMAND      | TIME | STATE      | INFO
| ROWS_SENT | ROWS_EXAMINED | ROWS_READ |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 0-0-352724126 | ifisibhk0 | test_123_wvvp_0000 | Query      | 15 | Sending data | /*DRDS/x
x.xxx.xx.88/ac47e5a72801000/*select `t_item`.`detail_url`,SUM(`t_item`.`price`) from `t_i | NULL | NULL |
NULL |
| 0-0-352864311 | cowxhthg0 | NULL      | Binlog Dump | 13 | Master has sent all binlog to slave; waiting for binlog
to be updated | NULL
| NULL | NULL | NULL |
| 0-0-402714566 | ifisibhk0 | test_123_wvvp_0005 | Query      | 14 | Sending data | /*DRDS/x
x.xxx.xx.88/ac47e5a72801000/*select `t_item`.`detail_url`,`t_item`.`price` from `t_i | NULL | NULL | N
ULL |
| 0-0-402714795 | ifisibhk0 | test_123_wvvp_0005 | Alter      | 114 | Sending data | /*DRDS/x
x.xxx.xx.88/ac47e5a72801000/*ALTER TABLE `Persons` ADD `Birthday` date | NULL | NULL | NULL |
.....
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
12 rows in set (0.03 sec)
```

The value in the TIME column indicates the number of seconds that the command has been executed. If a command execution is too slow, as shown in the figure, you can run the KILL '0-0-402714795' command to cancel the slow command.

 **Notice** In PolarDB-X, one logical SQL statement corresponds to multiple statements on database shards. Therefore, you may need to kill multiple commands to stop a logical DDL statement. You can determine the logical SQL statement to which a command belongs based on the INFO column in the SHOW PROCESSLIST result set.

4. Handle the lock conflict error

PolarDB-X adds a database lock before performing a DDL operation and releases the lock after the operation. The KILL DDL operation may not release the lock. If you perform the DDL operation again, the following error message will be returned:

```
Lock conflict , maybe last DDL is still running
```

In this case, run **RELEASE DBLOCK** to release the lock. After the command is canceled and the lock is released, run the DDL statement again during off-peak hours or when the service is stopped.

Other problems

Clients cannot display the modified table structures.

To enable some clients to obtain table structures from system tables (such as COLUMNS or TABLES), PolarDB-X creates a shadow database in database shard 0 on your ApsaraDB RDS for MySQL instance. The shadow database name must be the same as the name of your PolarDB-X logical database. It stores all table structures and other information in the user database.

The client obtains the PolarDB-X table structure from the system table of the shadow database. During the processing of DDL exceptions, the table structure may be modified normally in the user database but not in the shadow database due to some reasons. In this case, you need to connect to the shadow database and perform the DDL operation on the table again in the database.

 **Notice** The CHECK TABLE command does not check whether the table structure in the shadow database is consistent with that in the user database.

11.14.11. Efficiently scan DRDS data

Distributed Relational Database Service (DRDS) supports efficient data scanning and uses aggregate functions for statistical summary during full table scan.

The following describes common scanning scenarios:

- **Scan of table without database or table shards:** DRDS transmits the original SQL statement to the backend ApsaraDB RDS for MySQL database for execution. In this case, DRDS supports any aggregate functions.
- **Non-full table scan:** DRDS transmits the original SQL statement to each single ApsaraDB RDS for MySQL database for execution. For example, when the shard key in the WHERE clause is Equal, non-full table scan is performed. In this case, DRDS also supports any aggregate functions.
- **Full table scan:** Currently, the supported aggregate functions are COUNT, MAX, MIN, and SUM. In addition, LIKE, ORDER BY, LIMIT, and GROUP BY are also supported during full table scan.
- **Parallel scan of all table shards:** If you need to export data from all databases, you can run the SHOW command to view the table topology and scan all table shards in parallel. For more information, see the following section.

Traverse tables by using a hint

1. Run the SHOW TOPOLOGY FROM TABLE_NAME command to obtain the table topology.

```
mysql> SHOW TOPOLOGY FROM DRDS_USERS;
+-----+-----+-----+
| ID | GROUP_NAME | TABLE_NAME |
+-----+-----+-----+
| 0 | DRDS_00_RDS | drds_users |
| 1 | DRDS_01_RDS | drds_users |
+-----+-----+-----+
2 rows in set (0.06 sec)
```

By default, the non-partition table is stored in database shard 0.

2. Traverse each table for TOPOLOGY.
 - i. Run the current SQL statement in database shard 0.

```
#!/TDDL:node='DRDS_00_RDS'*/SELECT * FROM DRDS_USERS;
```

- ii. Run the current SQL statement in database shard 1.

```
#!/TDDL:node='DRDS_01_RDS'*/SELECT * FROM DRDS_USERS;
```

 **Notice** We recommend that you run `SHOW TOPOLOGY FROM TABLE_NAME` to obtain the latest table topology before each scan.

Parallel scans

DRDS allows you to run mysqldump to export data. However, if you want to scan data faster, you can enable multiple sessions for each table shard to scan tables in parallel.

```
mysql> SHOW TOPOLOGY FROM LJLTEST;
+-----+-----+-----+
| ID | GROUP_NAME | TABLE_NAME |
+-----+-----+-----+
| 0 | TDDL5_00_GROUP | ljlttest_00 |
| 1 | TDDL5_00_GROUP | ljlttest_01 |
| 2 | TDDL5_00_GROUP | ljlttest_02 |
| 3 | TDDL5_01_GROUP | ljlttest_03 |
| 4 | TDDL5_01_GROUP | ljlttest_04 |
| 5 | TDDL5_01_GROUP | ljlttest_05 |
| 6 | TDDL5_02_GROUP | ljlttest_06 |
| 7 | TDDL5_02_GROUP | ljlttest_07 |
| 8 | TDDL5_02_GROUP | ljlttest_08 |
| 9 | TDDL5_03_GROUP | ljlttest_09 |
| 10 | TDDL5_03_GROUP | ljlttest_10 |
| 11 | TDDL5_03_GROUP | ljlttest_11 |
+-----+-----+-----+
12 rows in set (0.06 sec)
```

As shown above, the table has four database shards, and each database shard has three table shards. Run the following SQL statement to operate on the table shards of the TDDL5_00_GROUP database:

```
#!/TDDL:node='TDDL5_00_GROUP'*/select * from ljlttest_00;
```

 **Note** TDDL5_00_GROUP in HINT corresponds to the GROUP_NAME column in the results of the SHOW TOPOLOGY command. In addition, the table name in the SQL statement is the table shard name.

At this time, you can enable up to 12 sessions (corresponding to 12 table shards respectively) to process data in parallel.

11.15. Appendix: PolarDB-X terms

This topic lists common terms of PolarDB-X for your reference.

| Term | Description | Remarks |
|---|--|---------|
| Cloud Native Distributed Database PolarDB-X | PolarDB-X is a distributed database service that was independently developed by Alibaba to solve the bottlenecks of single-instance database services. PolarDB-X is compatible with MySQL protocols and syntax. It supports automatic sharding, smooth scale-out, auto scaling, and transparent read/write splitting, and provides O&M capabilities for distributed databases throughout their entire lifecycle. | - |
| TDDL | Taobao Distributed Data Layer (TDDL) was developed by Alibaba and has become a preferred component for nearly 1,000 core applications of Alibaba. | - |
| PolarDB-X Console | PolarDB-X Console is designed for database administrators (DBAs) to isolate resources as required and perform operations, such as instance management, database and table management, read/write splitting configuration, smooth scale-out, monitoring data display, and IP address whitelist. | - |
| DRDS Manager | DRDS Manager is designed for global O&M personnel and DBAs to manage all PolarDB-X resources and monitor the system. | - |
| PolarDB-X Server | PolarDB-X Server is the service layer of PolarDB-X. Multiple server nodes make up a server cluster to provide distributed database services, including the read/write splitting, routed SQL execution, result merging, dynamic database configuration, and globally unique ID (GUID). | - |
| Load balancer | PolarDB-X server nodes are stateless, and therefore requests can be randomly routed to any PolarDB-X server node. The load balancer is used to complete this task. Server Load Balancer (SLB) is used for overall output by Apsara Stack. VIPServer is typically used for Alibaba middleware output. | - |
| Diamond | Diamond manages the configuration and storage of PolarDB-X. It provides the configuration functions for storage, query, and notification. In PolarDB-X, Diamond stores the source data of databases, and configuration data including the sharding rules, and switches. | - |
| Data Replication System | Data Replication System migrates and synchronizes data for PolarDB-X. Its core capabilities include full data migration and incremental data synchronization. Its derived features include smooth data import, smooth scale-out, and global secondary index. Data Replication System requires the support of ZooKeeper and PolarDB-X Rtools. | - |
| PolarDB-X instance (PolarDB-X instance) | A PolarDB-X instance consists of multiple PolarDB-X server nodes. A PolarDB-X instance can contain multiple PolarDB-X databases. | - |
| PolarDB-X instance ID (PolarDB-X instance ID) | An instance ID uniquely identifies an PolarDB-X instance. | - |
| Number of nodes on a PolarDB-X instance | The number of PolarDB-X server nodes in a PolarDB-X instance. | - |

| Term | Description | Remarks |
|------------------------|--|---|
| VIP | The virtual IP addresses (VIPs) of the load balancer can be classified as: <ul style="list-style-type: none"> 1. Public VIP, which is accessible from the Internet. It is used for testing. 2. Private VIP, which is accessible only from the Alibaba Cloud internal network. | - |
| VPC | Virtual Private Cloud (VPC) is generally used on Alibaba Cloud. | - |
| Region | A region is a geographical location, such as East China. This concept is generally used for Alibaba Cloud. | - |
| Azone | A physical area with independent power grids and networks within one region, such as Hangzhou Zone A. This concept is generally used for Alibaba Cloud. | - |
| Logical SQL statement | A logical SQL statement is an SQL statement sent from an application to PolarDB-X. | - |
| Physical SQL statement | A physical SQL statement is an SQL statement obtained after PolarDB-X parses a logical SQL statement and sends it to ApsaraDB RDS for MySQL for execution. | Logical SQL statements and physical SQL statements may be the same or different. Logical and physical SQL statements may be in a one-to-one or one-to-many mapping. |
| QPS | The queries per second (QPS) is the average number of logical SQL statements executed by PolarDB-X per second in a statistical period, | instead of the number of transactions. Most control statements, such as COMMIT and SET, are not counted in QPS. |
| RT | The response time (RT) is the average response time (in milliseconds) of logical SQL statements executed by PolarDB-X in a statistical period. The RT of an SQL statement is calculated as follows: (Time when PolarDB-X writes the last packet of the result set) - (Time when PolarDB-X receives the SQL statement) | - |
| Physical QPS | The physical QPS is the average number of physical SQL statements that PolarDB-X executes on ApsaraDB RDS for MySQL per second in a statistical period. | - |

| Term | Description | Remarks |
|--|--|---|
| Physical RT | <p>The physical RT is the average response time (in milliseconds) of physical SQL statements executed by PolarDB-X on ApsaraDB RDS for MySQL in a statistical period.</p> <p>The RT of a physical SQL statement is calculated as follows: (Time when PolarDB-X receives the result set returned by ApsaraDB RDS for MySQL) - (Time when PolarDB-X starts to obtain the connection to ApsaraDB RDS for MySQL)</p> | This includes the time of establishing a connection to ApsaraDB RDS for MySQL or obtaining a connection from the connection pool, the network transmission time, and the time of executing the SQL statement by ApsaraDB RDS for MySQL. |
| Connections | The number of connections established between the application and PolarDB-X, | instead of the number of connections established between PolarDB-X and ApsaraDB RDS for MySQL. |
| Inbound traffic | The network traffic generated when the application sends SQL statements to PolarDB-X. | This traffic is irrelevant to the traffic used for interaction between PolarDB-X and ApsaraDB RDS for MySQL. |
| Outbound traffic | The network traffic generated when PolarDB-X sends the result set to the application. | This traffic is irrelevant to the traffic used for interaction between PolarDB-X and ApsaraDB RDS for MySQL. |
| Number of active threads (ThreadRunning) | The number of threads running on a PolarDB-X instance. This parameter can be used to indicate the load of the PolarDB-X instance. | - |
| Global | The total monitoring data of all databases on a PolarDB-X instance. | - |
| Memory usage | The Java Virtual Machine (JVM) memory usage of a PolarDB-X server process. | - |
| Total memory usage | The memory usage of the machine where the PolarDB-X server node is located. | This metric is available only when PolarDB-X servers are deployed on ECS instances. Generally, this metric is used for Alibaba Cloud. |

| Term | Description | Remarks |
|-------------------|---|---|
| CPU utilization | The CPU utilization of the machine where a PolarDB-X server node is located. | This metric is available only when PolarDB-X servers are deployed on ECS instances. Generally, this metric is used for Alibaba Cloud. |
| System load | The load of the machine where a PolarDB-X server node is located. | This metric is available only when PolarDB-X servers are deployed on ECS instances. Generally, this metric is used for Alibaba Cloud. |
| Service port | The port used by PolarDB-X servers to provide MySQL-based services to external applications. | Generally, the port number is 3306. However, when multiple PolarDB-X nodes (mostly physical machines) are deployed on one machine, the port number will change accordingly. |
| Management port | The port used by PolarDB-X servers to provide management application program interfaces (APIs). | Generally, the port number is the service port number plus 100. |
| Start time | The time when PolarDB-X servers start. | - |
| Running time | The continuous running time of the PolarDB-X servers since the last startup time. | - |
| Total memory size | The maximum JVM memory size of a PolarDB-X server node. | - |
| Memory usage | The JVM memory that is already used by the PolarDB-X server nodes. | - |
| Number of nodes | Required. The number of machines. A PolarDB-X instance is essentially a PolarDB-X cluster, and the number of nodes refers to the number of machines in the cluster. | - |
| Instance type | Required. The type of the instance, including dedicated and shared instances. A dedicated instance works in the exclusive mode. A shared instance works in the multi-tenant mode, which is generally used in Alibaba Cloud. | - |
| Machine type | Required. The type of the machine where a PolarDB-X server node is deployed. Valid values are Auto-selected, PHY, and ECS. The PolarDB-X inventory is divided into physical machine inventory and virtual machine inventory according to the type of machines where the PolarDB-X servers are deployed. The two types cannot be mixed because their deployment and O&M methods are different. | - |
| AliUid | Required. The UID of the instance. In Apsara Stack, this ID is provided by the account system in the deployment environment. | - |

| Term | Description | Remarks |
|--------------------------|---|---------|
| Backend port | The backend port of the VIP. For a PolarDB-X server node, this port is the service port of machine where the PolarDB-X server node is deployed. | - |
| Frontend port | The frontend port of the VIP for user access. Each VIP has a set of frontend ports and backend ports. The VIP forwards data from frontend ports to backend ports. | - |
| Private network/Internet | The network type of the VIP. Valid values: <ul style="list-style-type: none"> Internet: the public VIP, which is accessible from the Internet. Private network: the private VIP (including VPC VIP), which is accessible from private networks. | - |
| lbld | The ID of an SLB instance, which is the unique ID of VIP. A VIP is managed based on this ID. | - |
| Forwarding mode | The port forwarding mode of the VIP. The following modes are supported: <ul style="list-style-type: none"> FNAT: This mode is recommended when the backend machine is a virtual machine or VPC needs to be supported. NAT: This mode can be selected when the backend machine is a physical machine. Currently, this mode is only used on Alibaba Cloud. Open FNAT: This mode is applicable only to Alibaba Cloud. | - |
| VPC ID | The ID of the destination VPC, that is, the VPC to be accessed. | - |
| VSwitch ID | The ID of the destination VSwitch, which determines the CIDR block where the VPC VIP of the instance is in. | - |
| APPName | The app name of the destination PolarDB-X database. Each PolarDB-X database has a corresponding app name for loading configurations. | - |
| UserName | The user name used to log on to the destination PolarDB-X database. | - |
| DBName | The name of the destination PolarDB-X database you want to log on to. | - |
| IP address whitelist | Only the IP addresses specified in the IP address whitelist can access the PolarDB-X instance. | - |
| Read-only instance | ApsaraDB RDS for MySQL instances where physical databases reside are divided into the following two types based on whether data can be written into the instances: <ul style="list-style-type: none"> Primary instance: Both read and write requests are allowed on such an instance. In Apsara Stack, ApsaraDB RDS for MySQL is supported. In Alibaba Cloud, ApsaraDB for RDS is supported. Read-only instance: Only read requests are allowed on such an instance. In Apsara Stack, ApsaraDB RDS for MySQL is supported. In Alibaba Cloud, ApsaraDB for RDS is supported. | - |

| Term | Description | Remarks |
|---------------------------|---|---------|
| Read SQL statement | A type of SQL statements used to read data, such as the SELECT statement. PolarDB-X determines whether an SQL statement is a read-only SQL statement when it is not in a transaction. If the SQL statement is in a transaction, PolarDB-X treats it as a write SQL statement during read/write splitting. | - |
| Read/write splitting | If read-only ApsaraDB RDS for MySQL instances exist, you can configure in the PolarDB-X console to allocate read SQL statements to the primary and read-only instances proportionally. PolarDB-X automatically identifies the type of SQL statements and allocates them proportionally. | - |
| Smooth scale-out | On the basis of horizontal partitioning, the data distribution on ApsaraDB RDS for MySQL instances is dynamically adjusted for scale-out. Generally, scale-out is completed asynchronously without any modification to the business code. | - |
| Broadcast of small tables | You can synchronize the data in a single table in a database to all database shards in advance, to convert the cross-database JOIN query into a JOIN query that can be completed on physical databases. | - |
| Horizontal partitioning | Horizontal partitioning distributes the data rows originally stored in one table to multiple tables based on specified rules to achieve horizontal linear scaling. | - |
| Partition mode | This mode allows you to create multiple database shards on an ApsaraDB RDS for MySQL instance. These database shards make up a PolarDB-X database. In this mode, all PolarDB-X functions can be used. | - |
| Non-partition mode | In this mode, a database that has been created on an ApsaraDB RDS for MySQL instance is used as a PolarDB-X database. In this mode, only PolarDB-X read/write splitting is allowed, while other PolarDB-X features such as database sharding and table sharding are not allowed. | - |
| Imported database | An existing database on the ApsaraDB RDS for MySQL instance selected for creating a PolarDB-X database. This is a unique concept for the creation of a PolarDB-X database. | - |
| Read policy | The ratio of read SQL statements assigned by PolarDB-X to the primary and read-only ApsaraDB RDS for MySQL instances. | - |
| Full table scan | If no shard field is specified in a SQL statement, PolarDB-X runs the SQL statement on all table shards and summarizes the results. You can disable this function because of its high overheads. | - |
| Shard key | A column in a logical table. PolarDB-X routes data and SQL statements to a physical table based on this column. | - |
| Data import | The operation of importing data from an existing ApsaraDB RDS for MySQL instance to a PolarDB-X database. | - |
| Full data migration | The operation of migrating all existing records from a database to PolarDB-X. An offset is recorded before full migration starts. | - |
| Offset | In a MySQL binary log file, each row represents a data change operation. The position of a line in the binary log file is called an offset. | - |

| Term | Description | Remarks |
|----------------------------|---|---------|
| Incremental data migration | The operation of reading all MySQL binary log records from the recorded offset, converting them into SQL statements, and then running them in PolarDB-X. Incremental migration continues before the switchover. | - |
| Switchover | A step of data import and smooth scale-out, which writes all the remaining incremental records from MySQL binary logs to PolarDB-X. | - |
| Cleanup | The last step of smooth scale-out, which cleans redundant data and configurations generated during smooth scale-out. | - |
| Heterogeneous indexing | For table shards of a PolarDB-X database, the WHERE condition of a SQL statement for query must contain the shard key whenever possible. In this way, PolarDB-X routes the query request to a specific database shard, improving the query efficiency. If the WHERE condition of the SQL statement does not contain the shard key, PolarDB-X performs a full table scan. PolarDB-X provides heterogeneous indexing to solve this problem. The data in a database shard or table shard of a PolarDB-X instance is fully or partially synchronized to another table based on different shard keys. The destination table to which the data is synchronized is called a heterogeneous index table. | - |
| PolarDB-X sequence | A PolarDB-X sequence (a 64-digit number of the BIGINT data type in MySQL) aims to ensure that the data (for example, PRIMARY KEY and UNIQUE KEY) in the defined unique field is globally unique and in ordered increments. | - |
| PolarDB-X hint | To facilitate PolarDB-X usage, PolarDB-X defines some hints to specify special actions. | - |

12. AnalyticDB for PostgreSQL

12.1. What is AnalyticDB for PostgreSQL?

AnalyticDB for PostgreSQL is a distributed cloud database service that uses multiple compute nodes to provide massively parallel processing (MPP) data warehousing.

AnalyticDB for PostgreSQL is developed based on the open source Greenplum database project and enhanced by Alibaba Cloud. This service has the following features:

- Compatible with Greenplum and all tools that support Greenplum.
- Supports Object Storage Service (OSS), JSON, and HyperLogLog, a probabilistic algorithm for cardinality estimation.
- Supports SQL:2003-compliant syntax and Online Analytical Processing (OLAP) aggregate functions to provide flexible hybrid analysis.
- Supports both row store and column store to enhance analytics performance.
- Supports data compression to reduce storage costs.
- Provides online scaling and performance monitoring to enable Database Administrators (DBAs), developers, and data analysts to focus on improving enterprise productivity and creating core business value instead of managing and maintaining large numbers of MPP clusters.

12.2. Quick start

12.2.1. Overview

This topic describes all operations that you can perform on an AnalyticDB for PostgreSQL instance, from instance creation to database logon. It provides a quick start guide to the operations on the AnalyticDB for PostgreSQL instance.

- [Log on to the AnalyticDB for PostgreSQL console](#)

You can log on to the AnalyticDB for PostgreSQL console.

- [Create an instance](#)

Before you perform other operations, you must first create an AnalyticDB for PostgreSQL instance in the console.

- [Configure a whitelist](#)

Before you use an AnalyticDB for PostgreSQL instance, add IP addresses or CIDR blocks needed to access your database to the whitelist of the instance to improve the security and stability of the database.

- [Create an initial account](#)

After you create an instance, you must create an initial account to log on to the database.

- [Connect to a database](#)

You can use a client that supports PostgreSQL or Greenplum to connect to a database.

12.2.2. Log on to the AnalyticDB for PostgreSQL console

This topic describes how to log on to the AnalyticDB for PostgreSQL console.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.

- A browser is available. We recommend that you use the Google Chrome browser.
 1. In the address bar of the browser, enter the URL of the Apsara Uni-manager Management Console and press the Enter key.
 2. Enter your username and password.

Obtain the username and password that are used to log on to the console from the operations administrator.

Note The first time you log on to the Apsara Uni-manager Management Console, you must change your password as prompted. For security reasons, your password must meet the minimum complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login** to go to the homepage of the Apsara Uni-manager Management Console.

12.2.3. Creates an instance

This topic describes how to create an AnalyticDB for PostgreSQL instance in the console. Before you perform other operations, you must first create an AnalyticDB for PostgreSQL instance.

1. [Log on to the AnalyticDB for PostgreSQL console.](#)
2. In the upper-right corner of the page, click **Create Instance**.
3. On the **Create AnalyticDB for PostgreSQL Instance** page, set the following parameters.

| Section | Parameter | Description |
|----------------|----------------|--|
| Region | Organization | The organization to which the instance belongs. |
| | Resource Set | The resource set to which the instance belongs. |
| | Region | The region in which you want to create the AnalyticDB for PostgreSQL instance. Note If you want to access the AnalyticDB for PostgreSQL instance from an Elastic Compute Service (ECS) instance over a virtual private cloud (VPC), you must create the instance in the same region and zone as those of the ECS instance. |
| | Zone | The zone in which you want to create the AnalyticDB for PostgreSQL instance. |
| Basic Settings | Engine | Only the integrated computing and storage version is supported. |
| | Engine Version | The engine version of the AnalyticDB for PostgreSQL instance. |
| | Node Type | Specifications for each compute node. The storage space and compute capability of a compute node vary based on the specified specifications. |

| Section | Parameter | Description |
|---------|--------------|--|
| | Nodes | The number of compute nodes. An instance must contain at least two compute nodes. The performance of an instance scales linearly with the number of compute nodes. |
| Network | Network Type | The network type of the AnalyticDB for PostgreSQL instance. Valid values: <ul style="list-style-type: none"> ◦ <i>Classic Network</i>: Cloud services within the classic network are not isolated from each other. Unauthorized access to a cloud service can be blocked only by the security group or whitelist policy of the service. ◦ <i>VPC</i>: A VPC helps you build an isolated network environment on Apsara Stack. You can customize route tables, CIDR blocks, and gateways in a VPC. We recommend that you select VPC for improved security. You can create a VPC in advance, or change the network type to VPC after the instance is created. |
| | VPC | The VPC in which you want to create the AnalyticDB for PostgreSQL instance. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> ? Note VPC: You can use a VPC to build an isolated network environment on Apsara Stack. You can customize route tables, CIDR blocks, and gateways in a VPC. </div> |
| | vSwitch | The vSwitch to which the AnalyticDB for PostgreSQL instance is attached. |
| | IP Whitelist | The IP addresses that are allowed to access the instance. |

4. After you set the preceding parameters, click **Submit**.

12.2.4. Configure a whitelist

To ensure a secure and stable database, you must add IP addresses or CIDR blocks that are allowed to access the database to a whitelist.

1. [Log on to the AnalyticDB for PostgreSQL console](#).
2. Find the instance that you want to manage and click its ID. The **Basic Information** page appears.
3. In the left-side navigation pane, click **Security Controls**. The **Security Controls** page appears.
4. On the **Whitelist Settings** tab, click **Modify** corresponding to the *default* whitelist. The **Modify Group** page appears.

? **Note** You can also click **Clear** corresponding to the *default* whitelist to delete the IP addresses in the default whitelist. Then, click **Add Group** to create another whitelist.

5. Delete 127.0.0.1 from the *default* whitelist and enter your IP addresses in the whitelist. The following table describes the parameters.

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|--------------|---|
| Group Name | Specify the name of the whitelist. The name must be 2 to 32 characters in length and can contain lowercase letters, digits, and underscores (_). It must start with a lowercase letter and end with a letter or digit. The default whitelist cannot be modified or deleted. |
| IP Addresses | <p>Enter the CIDR blocks or IP addresses that are allowed to access the database. Use commas (,) to separate multiple CIDR blocks or IP addresses.</p> <ul style="list-style-type: none"> ◦ A whitelist can contain IP addresses such as 10.10.10.1 and CIDR blocks such as 10.10.10.0/24. This CIDR block indicates that all IP addresses in the 10.10.10.X format have access to the database. ◦ The percent sign (%) or 0.0.0.0/0 indicates that all IP addresses are allowed to access the database. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin: 10px 0;"> <p> Notice We recommend that you do not use this configuration because it reduces the security of the database.</p> </div> <ul style="list-style-type: none"> ◦ Default whitelists of new instances contain the loopback address 127.0.0.1. This configuration allows no access from external IP addresses. ◦ You can add up to 999 IP addresses or CIDR blocks to a whitelist. |

6. Click **OK**.

What's next

- We recommend that you maintain the whitelist on a regular basis to ensure secure access for AnalyticDB for PostgreSQL.
- You can click **Modify** or **Delete** to modify or delete custom whitelists.

12.2.5. Create an initial account

After you create an instance, you must create an initial account to log on to the database.

1. [Log on to the AnalyticDB for PostgreSQL console.](#)
2. Find the instance that you want to manage and click its ID. The **Basic Information** page appears.
3. In the left-side navigation pane, click **Account Management**. The **Account Management** page appears.
4. In the upper-right corner of the page, click **Create Account**. The **Create Account** page appears.
5. Enter the database account and password, and click **OK**.

| Parameter | Description |
|--------------|---|
| Account | The name of the account must be 2 to 16 characters in length and can contain lowercase letters, digits, and underscores (_). It must start with a letter and end with a letter or digit. |
| New Password | The password must be 8 to 32 characters in length. It must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. |
| Password | Enter the password again. |

12.2.6. Obtain client tools

The interface protocol of AnalyticDB for PostgreSQL is compatible with Greenplum Community Edition and PostgreSQL 8.2. You can use the Greenplum or PostgreSQL client to connect to AnalyticDB for PostgreSQL.

 Note

Apsara Stack is an isolated environment. You must deploy software installation packages to the internal environment.

Graphical client tools

AnalyticDB for PostgreSQL users can directly use client tools that support Greenplum, such as [SQL Workbench](#), [Navicat Premium](#), [Navicat for PostgreSQL](#), and [pgadmin III v1.6.3](#).

Command-line client psql (for RHEL 6, RHEL 7, CentOS 6, and CentOS 7)

For Red Hat Enterprise Linux (RHEL) 6, RHEL 7, CentOS 6, and CentOS 7, download the tools from the following links and decompress the packages to use them:

- For RHEL 6 or CentOS 6, click [hybriddb_client_package_el6](#).
- For RHEL 7 or CentOS 7, click [hybriddb_client_package_el7](#).

Command-line client psql (for other Linux systems)

For other Linux systems, perform the following operations to compile the client tools:

1. Obtain the source code by using one of the following methods:
 - Obtain the git directory. You must first install the git tool.

```
git clone https://github.com/greenplum-db/gpdb.git
cd gpdb
git checkout 5d870156
```

- Download the code.

```
wget https://github.com/greenplum-db/gpdb/archive/5d87015609abd330c68a5402c1267fc86cbc9e1f.zip
unzip 5d87015609abd330c68a5402c1267fc86cbc9e1f.zip
cd gpdb-5d87015609abd330c68a5402c1267fc86cbc9e1f
```

2. Use GCC and other compilers.

```
./configure
make -j32
make install
```

3. Use psql and pg_dump. The two tools are located in the following paths:

```
psql: /usr/local/pgsql/bin/psql
pg_dump: /usr/local/pgsql/bin/pg_dump
```

Command-line client psql (for Windows and other systems)

For Windows and other systems, go to the Pivotal website to download [HybridDB Client](#).

12.2.7. Connect to a database

Greenplum Database and AnalyticDB for PostgreSQL are both developed based on PostgreSQL 8.2 and fully compatible with its message protocol. AnalyticDB for PostgreSQL users can use tools that support the PostgreSQL 8.2 message protocol, such as libpq, Java Database Connectivity (JDBC), Open Database Connectivity (ODBC), psycopg2, and pgAdmin III.

Context

AnalyticDB for PostgreSQL provides `psql`, a binary program of Red Hat. For more information about the download link, see [Obtain the client tool](#). The Greenplum official website provides an easy-to-install installation package that includes JDBC, ODBC, and libpq. For more information, see [Greenplum official documentation](#).

Note

- Apsara Stack is an isolated environment. To access Apsara Stack, you must prepare the necessary software installation packages in advance.
- By default, AnalyticDB for PostgreSQL instances can be accessed only by clients that are deployed on Elastic Compute Service (ECS) instances within the same region and zone.

psql

`psql` is a common tool used together with Greenplum, and provides a variety of command functions. Its binary files are located in the `bin` directory of Greenplum. To use `psql`, perform the following steps:

1. Use one of the following methods to connect to the database:

- Connection string

```
psql "host=yourgpdbaddress.gpdb.rds.aliyuncs.com port=3432 dbname=postgres user=gpdbaccount password=gpdbpassword"
```

- Specified parameters

```
psql -h yourgpdbaddress.gp.aliyun-inc.com -p 3432 -d postgres -U gpdbaccount
```

The following section describes the parameters:

- `-h`: the host address.
- `-p`: the port used to connect to the database.
- `-d`: the name of the database. The default value is `postgres`.
- `-U`: the account used to connect to the database.

You can run the `psql --help` command to view more options. You can also run the `\?` command to view the commands supported in `psql`.

2. Enter the password to go to the `psql` shell interface.

```
postgres=>
```

References

- For more information about the Greenplum `psql` usage, see [psql](#).
- AnalyticDB for PostgreSQL also supports `psql` commands of PostgreSQL. Pay attention to the differences between Greenplum `psql` and PostgreSQL `psql`. For more information, see [PostgreSQL 8.3.23 Documentation - psql](#).

pgAdmin III

pgAdmin III is a PostgreSQL graphical client that can be directly used to connect to AnalyticDB for PostgreSQL. For more information, click [here](#). For more information about other graphical clients, see [Obtain the client tool](#).

1. Download pgAdmin III 1.6.3 or earlier.

You can download pgAdmin III 1.6.3 from the [PostgreSQL website](#). pgAdmin III 1.6.3 supports various operating systems, such as Windows, macOS, and Linux.

 **Note** AnalyticDB for PostgreSQL is compatible with PostgreSQL 8.2. Therefore, you must use pgAdmin III 1.6.3 or earlier to connect to AnalyticDB for PostgreSQL. pgAdmin 4 and later are not supported.

2. Choose **File > Add Server**.
3. In the New Server Registration dialog box, set the parameters.
4. Click **OK** to connect to AnalyticDB for PostgreSQL.

JDBC

JDBC uses the interface provided by PostgreSQL. Use the following method to download the JDBC driver:

Click [here](#) to download the official JDBC of PostgreSQL. Then, add it to the environment variables.

Sample code:

```
import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.ResultSet;
import java.sql.SQLException;
import java.sql.Statement;
public class gp_conn {
    public static void main(String[] args) {
        try {
            Class.forName("org.postgresql.Driver");
            Connection db = DriverManager.getConnection("jdbc:postgresql://mygpdbpub.gpdb.rds.aliyuncs.com:3432/postgres",
                "mygpdb", "mygpdb");
            Statement st = db.createStatement();
            ResultSet rs = st.executeQuery(
                "select * from gp_segment_configuration;");
            while (rs.next()) {
                System.out.print(rs.getString(1));
                System.out.print(" | ");
                System.out.print(rs.getString(2));
                System.out.print(" | ");
                System.out.print(rs.getString(3));
                System.out.print(" | ");
                System.out.print(rs.getString(4));
                System.out.print(" | ");
                System.out.print(rs.getString(5));
                System.out.print(" | ");
                System.out.print(rs.getString(6));
                System.out.print(" | ");
                System.out.print(rs.getString(7));
                System.out.print(" | ");
                System.out.print(rs.getString(8));
                System.out.print(" | ");
                System.out.print(rs.getString(9));
                System.out.print(" | ");
                System.out.print(rs.getString(10));
                System.out.print(" | ");
                System.out.println(rs.getString(11));
            }
            rs.close();
            st.close();
        } catch (ClassNotFoundException e) {
            e.printStackTrace();
        } catch (SQLException e) {
            e.printStackTrace();
        }
    }
}
```

Python

Python uses `psycopg2` to connect to Greenplum and PostgreSQL. Perform the following operations:

1. Install `psycopg2`. Use one of the following methods to install `psycopg2` in Cent OS:
 - Method 1: Run the `yum -y install python-psycopg2` command.
 - Method 2: Run the `pip install psycopg2` command.
 - Method 3: Run the following source code:

```
yum install -y postgresql-devel*
wget http://initd.org/psycopg/tarballs/PSYCOPG-2-6/psycopg2-2.6.tar.gz
tar xf psycopg2-2.6.tar.gz
cd psycopg2-2.6
python setup.py build
sudo python setup.py install
```

2. Run the following commands to set PYTHONPATH and reference it:

```
import psycopg2
sql = 'select * from gp_segment_configuration;'
conn = psycopg2.connect(database='gpdb', user='mygpdb', password='mygpdb', host='mygpdbpub.gpdb.rds.aliyuncs.com', port=3432)
conn.autocommit = True
cursor = conn.cursor()
cursor.execute(sql)
rows = cursor.fetchall()
for row in rows:
    print row
conn.commit()
conn.close()
```

An output similar to the following one is displayed:

```
(1, -1, 'p', 'p', 's', 'u', 3022, '192.168.2.158', '192.168.2.158', None, None)(6, -1, 'm', 'm', 's', 'u', 3019, '192.168.2.47', '192.168.2.47', None, None)(2, 0, 'p', 'p', 's', 'u', 3025, '192.168.2.148', '192.168.2.148', 3525, None)(4, 0, 'm', 'm', 's', 'u', 3024, '192.168.2.158', '192.168.2.158', 3524, None)(3, 1, 'p', 'p', 's', 'u', 3023, '192.168.2.158', '192.168.2.158', 3523, None)(5, 1, 'm', 'm', 's', 'u', 3026, '192.168.2.148', '192.168.2.148', 3526, None)
```

libpq

libpq is the C language interface to AnalyticDB for PostgreSQL. You can use the libpq library to access and manage PostgreSQL databases in a C program. You can find its static and dynamic libraries in the lib directory.

For more information about example programs, see [Example Programs](#).

For more information about libpq, see [PostgreSQL 9.4.17 Documentation - Chapter 31. libpq - C Library](#).

ODBC

PostgreSQL ODBC is an open source version based on the GNU Lesser General Public License (LGPL) protocol. You can download it from the [PostgreSQL website](#).

1. Install the driver.

```
yum install -y unixODBC.x86_64
yum install -y postgresql-odbc.x86_64
```

2. View the driver configurations.

```

cat /etc/odbcinst.ini
# Example driver definitions
# Driver from the postgresql-odbc package
# Setup from the unixODBC package
[PostgreSQL]
Description = ODBC for PostgreSQL
Driver = /usr/lib/psqlodbcw.so
Setup = /usr/lib/libodbcpsqlS.so
Driver64 = /usr/lib64/psqlodbcw.so
Setup64 = /usr/lib64/libodbcpsqlS.so
FileUsage = 1
# Driver from the mysql-connector-odbc package
# Setup from the unixODBC package
[MySQL]
Description = ODBC for MySQL
Driver = /usr/lib/libmyodbc5.so
Setup = /usr/lib/libodbcmyS.so
Driver64 = /usr/lib64/libmyodbc5.so
Setup64 = /usr/lib64/libodbcmyS.so
FileUsage = 1

```

3. Configure the data source name (DSN). Replace **** in the following code with the corresponding connection information.

```

[mygpdb]
Description = Test to gp
Driver = PostgreSQL
Database = ****
Servername = ****.gpdb.rds.aliyuncs.com
UserName = ****
Password = ****
Port = ****
ReadOnly = 0

```

4. Test the connectivity.

```

echo "select count(*) from pg_class" | isql mygpdb
+-----+
| Connected!          |
|                    |
| sql-statement      |
| help [tablename]   |
| quit               |
|                    |
+-----+
SQL> select count(*) from pg_class
+-----+
| count  |
+-----+
| 388    |
+-----+
SQLRowCount returns 1
1 rows fetched

```

5. After ODBC is connected to the instance, connect the application to ODBC. For more information, see [psqlODBC - PostgreSQL ODBC driver](#) and [psqlODBC HOWTO - C#](#).

References

- [Pivotal Greenplum official documentation](#)

- [PostgreSQL psqLODBC](#)
- [Compiling psqLODBC on Unix](#)
- [Download ODBC connectors](#)
- [Download JDBC connectors](#)
- [The PostgreSQL JDBC Interface](#)

12.3. Instances

12.3.1. Reset the password

When you use AnalyticDB for PostgreSQL, you can reset the password of your database account in the AnalyticDB for PostgreSQL console if you forget your password.

 **Note** To ensure data security, we recommend that you change your password on a regular basis.

1. [Log on to the AnalyticDB for PostgreSQL console](#).
2. Find the target instance and click its ID. The **Basic Information** page appears.
3. In the left-side navigation pane, click **Account Management**. The **Account Management** page appears.
4. Click **Reset Password** in the Actions column corresponding to an account. The **Reset Account Password** page appears.
5. After you enter and confirm the new password, click **OK**.

 **Note** The password must be 8 to 32 characters in length. It must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. We recommend that you do not use a previously used password.

12.3.2. View monitoring information

You can go to the monitoring information page in the console to view the operation status of an instance.

1. [Log on to the AnalyticDB for PostgreSQL console](#).
2. Find the instance that you want to manage and click its ID. The **Basic Information** page appears.
3. In the left-side navigation pane, click **Monitoring and Alarms**. The **Monitoring and Alarms** page appears.
Specify a duration of time n up to seven days in length to view the metrics for that last n period.

12.3.3. Switch the network type of an instance

The default network type of an instance is Virtual Private Cloud (VPC). After an instance is created, you can switch its network type between classic network and VPC as needed.

Context

AnalyticDB for PostgreSQL supports two network types: classic network and VPC. Both network types use BGP connections, and are independent of the public network of your service provider. These network types only differ in functions, and you can choose a network type based on your requirements. The two network types are applicable to different scenarios:

- **Classic network:** IP addresses are allocated by Alibaba Cloud. Classic networks are easy to configure and use. This network type is suitable for users who do not need to perform complex operations, or who only require short deployment cycles.
- **VPC:** a logically isolated private network. You can customize the network topology and IP addresses and

connect through a leased line. This network type is suitable for advanced users.

 **Warning** Switching the network type will cause the database service to stop. Proceed with caution.

1. [Log on to the AnalyticDB for PostgreSQL console.](#)
2. Find the target instance and click its ID. The **Basic Information** page appears.
3. In the left-side navigation pane, click **Database Connection**. The **Database Connection** page appears.
4. In the upper-right corner of the page, click **Switch to Classic Network** or **Switch to VPC**.
5. If you click **Switch to VPC**, you must select the destination **VPC** and **VSwitch**. Click **OK**.

 **Note** To switch the network type to VPC, a VPC and a VSwitch must exist or be created in the zone where the instance is located.

6. If you click **Switch to Classic Network**, click **OK** in the displayed message.

 **Note** After you switch the network type, it takes 3 to 30 minutes for the instance to enter the running state.

12.3.4. Restart an instance

To better meet your requirements, AnalyticDB for PostgreSQL updates the minor kernel version on a regular basis. When you create an instance, the latest database kernel is used by default. After a new version is released, you can restart your instance to update the database kernel and use its extended features. This topic describes how to restart an instance.

 **Warning** The database service may be interrupted when you restart an instance.

1. [Log on to the AnalyticDB for PostgreSQL console.](#)
2. Find the instance that you want to manage and click its ID. The **Basic Information** page appears.
3. In the upper-right corner of the page, click **Restart Instance**.

 **Note** The restart process takes about 3 to 30 minutes. During the restart period, the instance cannot provide external services. We recommend that you take precautionary measures before you restart the instance. After the instance is restarted and enters the running state, you can access the database.

12.3.5. Import data

12.3.5.1. Import data from or export data to OSS in parallel

AnalyticDB for PostgreSQL allows you to import data from or export data to Object Storage Service (OSS) tables in parallel by using the OSS external table feature, `gpossex`. AnalyticDB for PostgreSQL also supports GZIP compression for OSS external tables to reduce file size and storage costs. `gpossex` can read from and write to TEXT and CSV files, even when they are compressed in GZIP packages.

- Create an OSS external table extension (`oss_ext`)

To use an OSS external table, you must first create an OSS external table extension in AnalyticDB for PostgreSQL. You must create an extension for each database that you need to access.

- To create the extension, execute the `CREATE EXTENSION IF NOT EXISTS oss_ext;` statement.
- To delete the inextension, execute the `DROP EXTENSION IF EXISTS oss_ext;` statement.

- Import data in parallel
 - i. Distribute data evenly into multiple files in OSS. We recommend that you set the number of OSS files to an integer that is the multiple of the number of compute nodes in AnalyticDB for PostgreSQL.
 - ii. Create a READABLE external table in AnalyticDB for PostgreSQL.
 - iii. Execute the following statement to import data in parallel:

```
INSERT INTO <Destination table> SELECT * FROM <External table>
```

 **Note**

- The data import performance depends on the OSS performance and resources of AnalyticDB for PostgreSQL instances, such as CPU, I/O, memory, and network resources. To ensure the best import performance, we recommend that you use column store and compression when you create a table. For example, you can specify the following clause: `WITH (APPENDONLY=true, ORIENTATION=column, COMPRESSIONTYPE=zlib, COMPRESSIONLEVEL=5, BLOCKSIZE=1048576)`. For more information, see [Greenplum Database official documentation on database table creation syntax](#).
- To ensure the best import performance, we recommend that you configure OSS and AnalyticDB for PostgreSQL instances within the same region.

- Export data in parallel
 - i. Create a WRITABLE external table in AnalyticDB for PostgreSQL.
 - ii. Execute the following statement to export data to OSS in parallel:

```
INSERT INTO <External table> SELECT * FROM <Source table>
```

- Create OSS external tables

 **Note** The syntax to create and use external tables is the same as that of Greenplum Database, except for the syntax of location-related parameters.

```

CREATE [READABLE] EXTERNAL TABLE tablename
( columnname datatype [, ...] | LIKE othertable )
LOCATION ('ossprotocol')
FORMAT 'TEXT'
  (( [HEADER]
    [DELIMITER [AS] 'delimiter' | 'OFF']
    [NULL [AS] 'null string']
    [ESCAPE [AS] 'escape' | 'OFF']
    [NEWLINE [ AS ] 'LF' | 'CR' | 'CRLF']
    [FILL MISSING FIELDS] ))
  | 'CSV'
  (( [HEADER]
    [QUOTE [AS] 'quote']
    [DELIMITER [AS] 'delimiter']
    [NULL [AS] 'null string']
    [FORCE NOT NULL column [, ...]]
    [ESCAPE [AS] 'escape']
    [NEWLINE [ AS ] 'LF' | 'CR' | 'CRLF']
    [FILL MISSING FIELDS] ))
[ ENCODING 'encoding' ]
[ [LOG ERRORS [INTO error_table]] SEGMENT REJECT LIMIT count
  [ROWS | PERCENT] ]
CREATE WRITABLE EXTERNAL TABLE table_name
( column_name data_type [, ...] | LIKE other_table )
LOCATION ('ossprotocol')
FORMAT 'TEXT'
  (( [DELIMITER [AS] 'delimiter']
    [NULL [AS] 'null string']
    [ESCAPE [AS] 'escape' | 'OFF'] ))
  | 'CSV'
  (( [QUOTE [AS] 'quote']
    [DELIMITER [AS] 'delimiter']
    [NULL [AS] 'null string']
    [FORCE QUOTE column [, ...]] ]
    [ESCAPE [AS] 'escape'] ))
[ ENCODING 'encoding' ]
[ DISTRIBUTED BY (column, [ ... ] ) | DISTRIBUTED RANDOMLY ]
ossprotocol:
  oss://oss_endpoint prefix=prefix_name
  id=userossid key=userosskey bucket=ossbucket compressiontype=[none|gzip] async=[true|false]
ossprotocol:
  oss://oss_endpoint dir=[folder/[folder/]...]/file_name
  id=userossid key=userosskey bucket=ossbucket compressiontype=[none|gzip] async=[true|false]
ossprotocol:
  oss://oss_endpoint filepath=[folder/[folder/]...]/file_name
  id=userossid key=userosskey bucket=ossbucket compressiontype=[none|gzip] async=[true|false]

```

Parameters

Common parameters

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|-----------------------|---|
| Protocol and endpoint | <p>This parameter is in the <code>protocol name://oss_endpoint</code> format. The protocol name is <code>oss</code>. <code>oss_endpoint</code> is the domain name that is used to access OSS in a region.</p> <p>Note You can access the database from a virtual private cloud (VPC) host by using an internal endpoint that contains "internal" in the name to avoid generating public traffic.</p> |
| id | The AccessKey ID of the OSS account. |
| key | The AccessKey secret of the OSS account. |
| bucket | The bucket where the data file is located. You must use OSS to create the bucket before you import data. |
| prefix | <p>The prefix of the path name corresponding to the data file. Prefixes are directly matched and cannot be controlled by regular expressions. The prefix, filepath, and dir parameters are mutually exclusive. Only one of the parameters can be specified at a time.</p> <ul style="list-style-type: none"> If you create a READABLE external table for data import, all OSS files that contain the specified prefix are imported. <ul style="list-style-type: none"> If you set prefix to <code>test/filename</code>, the following files are imported: <ul style="list-style-type: none"> <code>test/filename</code> <code>test/filenamexxx</code> <code>test/filename/aa</code> <code>test/filenameyyy/aa</code> <code>test/filenameyyy/bb/aa</code> If you set prefix to <code>test/filename/</code>, only the following file out of the preceding files is imported: <ul style="list-style-type: none"> <code>test/filename/aa</code> If you create a WRITABLE external table for data export, each exported file has a unique name based on this parameter. <p>Note One or more files can be exported for each compute node. The names of exported files are in the <code>prefix_tablename_uuid.x</code> format. <code>uuid</code> indicates a timestamp in microseconds as an int64 value. <code>x</code> indicates the node ID. You can use an external table for multiple export operations. Each export operation is assigned a <code>uuid</code> value. The files exported during each operation share a <code>uuid</code> value.</p> |

| Parameter | Description |
|-----------|--|
| dir | <p>The virtual folder path in OSS. The prefix, filepath, and dir parameters are mutually exclusive. Only one of the parameters can be specified at a time.</p> <ul style="list-style-type: none"> A folder path must end with a forward slash (/). Example: <code>test/mydir/</code>. If you use this parameter when you create an external table for data import, all files in the specified virtual directory (except for its subdirectories and contained files) are imported. Unlike filepath, dir does not require you to specify the names of files in the directory. If you use this parameter when you create an external table for data export, all data is exported as multiple files within the specified directory. The names of exported files are in the <code>filename.x</code> format, where x is a number. The values of x may not be consecutive. |
| filepath | <p>The file name that contains a path in OSS. The prefix, filepath, and dir parameters are mutually exclusive. Only one of the parameters can be specified at a time. You can specify only the filepath parameter when you create a READABLE external table for data import.</p> <ul style="list-style-type: none"> The file name includes the file path, but not the bucket name. The file name that is specified for data import must be in the <code>filename</code> or <code>filename.x</code> format. The values of x must be consecutive numbers that start from 1. <p>For example, if filepath is set to filename and OSS contains the following files, the imported files include filename, filename.1, and filename.2, but filename.4 is not imported because filename.3 does not exist.</p> <pre>filename filename.1 filename.2 filename.4</pre> |

Import mode parameters

| Parameter | Description |
|------------------|---|
| async | <p>Specifies whether to enable asynchronous data import.</p> <ul style="list-style-type: none"> By default, asynchronous data import is enabled. You can set <code>async</code> to <code>false</code> or <code>f</code> to disable asynchronous data import. You can enable the worker thread to load data from OSS to accelerate the import performance. By default, asynchronous data import is used. Asynchronous data import consumes more hardware resources than normal data import. |
| compressiontype | <p>The compression format of the imported files. Valid values:</p> <ul style="list-style-type: none"> <code>none</code>: The import files are not compressed. This is the default value. <code>gzip</code>: The imported files are compressed in the GZIP format. Only the GZIP format is supported. |
| compressionlevel | <p>The compression level of the files that are written to OSS. Valid values: 1 to 9. Default value: 6.</p> |

Export mode parameters

| Parameter | Description |
|----------------------|--|
| oss_flush_block_size | The size of each data block that is written to OSS. Valid values: 1 to 128. Default value: 32. Unit: MB. |
| oss_file_max_size | The maximum size for each file that is written to OSS. If the limit is exceeded, subsequent data is written to another file. Valid values: 8 to 4000. Default value: 1024. Unit: MB. |
| num_parallel_worker | The number of parallel compression threads for data that is written to OSS. Valid values: 1 to 8. Default value: 3. |

For data export, take note of the following items:

- WRITABLE is the keyword of the external table for data export. You must specify this keyword when you create an external table.
- Only the prefix and dir parameters are supported for data export. The filepath parameter is not supported.
- You can use the DISTRIBUTED BY clause to write data from compute nodes to OSS based on the specified distribution keys.

Other common parameters

The following table describes the fault-tolerance parameters that can be used for data import and export.

Fault-tolerance parameters

| Parameter | Description |
|-----------------------|---|
| oss_connect_timeout | The connection timeout period. Default value: 10. Unit: seconds. |
| oss_dns_cache_timeout | The Alibaba Cloud DNS (DNS) timeout period. Default value: 60. Unit: seconds. |
| oss_speed_limit | The minimum rate tolerated. Default value: 1024 bit/s (1 Kbit/s). |
| oss_speed_time | The maximum amount of time tolerated. Default value: 15. Unit: seconds. |

When the default values are used for the preceding parameters, a timeout occurs if the transmission rate is lower than 1 Kbit/s for 15 consecutive seconds. For more information, see **Troubleshooting** in *OSS SDK reference*.

The other parameters are compatible with the external table syntax of Greenplum Database. For more information about the syntax, see [Greenplum Database official documentation on external table syntax](#). The following parameters are included:

- FORMAT: the supported file format, such as TEXT and CSV.
- ENCODING: the data encoding format of a file, such as UTF-8.
- LOG ERRORS refers to improperly imported data that can be ignored and is instead written to error_table. You can also use the count parameter to specify the error reporting threshold.

Examples

```
# Create a READABLE external table of OSS.
create readable external table ossexample
  (date text, time text, open float, high float,
  low float, volume int)
  location('oss://oss-cn-hangzhou.aliyuncs.com
  prefix=osstest/example id=XXX
  key=XXX bucket=testbucket compressiontype=gzip')
  FORMAT 'csv' (QUOTE '' DELIMITER E'\t')
  ENCODING 'utf8'
```

```

LOG ERRORS INTO my_error_rows SEGMENT REJECT LIMIT 5;
create readable external table ossexample
(date text, time text, open float, high float,
 low float, volume int)
location('oss://oss-cn-hangzhou.aliyuncs.com
dir=osstest/id=XXX
key=XXX bucket=testbucket')
FORMAT 'csv'
LOG ERRORS SEGMENT REJECT LIMIT 5;
create readable external table ossexample
(date text, time text, open float, high float,
 low float, volume int)
location('oss://oss-cn-hangzhou.aliyuncs.com
filepath=osstest/example.csv id=XXX
key=XXX bucket=testbucket')
FORMAT 'csv'
LOG ERRORS SEGMENT REJECT LIMIT 5;
# Create a WRITABLE external table of OSS.
create WRITABLE external table ossexample_exp
(date text, time text, open float, high float,
 low float, volume int)
location('oss://oss-cn-hangzhou.aliyuncs.com
prefix=osstest/exp/outfromhdb id=XXX
key=XXX bucket=testbucket') FORMAT 'csv'
DISTRIBUTED BY (date);
create WRITABLE external table ossexample_exp
(date text, time text, open float, high float,
 low float, volume int)
location('oss://oss-cn-hangzhou.aliyuncs.com
dir=osstest/exp/id=XXX
key=XXX bucket=testbucket') FORMAT 'csv'
DISTRIBUTED BY (date);
# Create a heap table named example to which you want to import data.
create table example
(date text, time text, open float,
 high float, low float, volume int)
DISTRIBUTED BY (date);
# Import data to the example heap table from the ossexample table in parallel.
insert into example select * from ossexample;
# Export data from the example heap table to OSS in parallel.
insert into ossexample_exp select * from example;
# The following execution plan shows that all compute nodes are involved in the task.
# All compute nodes read data from OSS in parallel. AnalyticDB for PostgreSQL performs a redistribution motion operati
on to compute the data by using a hash algorithm, and then distributes the data to its compute nodes after computing.
After a compute node receives data, it performs an insert operation to add the data to AnalyticDB for PostgreSQL.
explain insert into example select * from ossexample;
          QUERY PLAN
-----
Insert (slice0; segments: 4) (rows=250000 width=92)
-> Redistribute Motion 4:4 (slice1; segments: 4) (cost=0.00..11000.00 rows=250000 width=92)
    Hash Key: ossexample.date
      -> External Scan on ossexample (cost=0.00..11000.00 rows=250000 width=92)
(4 rows)
# The following query plan shows that each compute node directly exports data to OSS without redistributing the data.
explain insert into ossexample_exp select * from example;
          QUERY PLAN
-----
Insert (slice0; segments: 3) (rows=1 width=92)
-> Seq Scan on example (cost=0.00..0.00 rows=1 width=92)
(2 rows)

```

TEXT and CSV format description

The following parameters specify the formats of files read from and written to OSS. You can specify the parameters in the external DDL parameters.

- \n: the line feed for TEXT and CSV files.
- DELIMITER: the delimiter of columns.
 - If the DELIMITER parameter is specified, the QUOTE parameter must also be specified.
 - Recommended column delimiters include commas (,), vertical bars (|), and special characters such as \t.
- QUOTE: encloses user data that contains special characters by column.
 - Strings that contain special characters must be enclosed by QUOTE to differentiate user data from control characters.
 - To optimize the efficiency, it is unnecessary to enclose data such as integers in QUOTE characters.
 - QUOTE cannot be the same string as specified in DELIMITER. The default value of QUOTE is double quotation marks (").
 - User data that contains QUOTE characters must also contain ESCAPE characters to differentiate user data from machine code.
- ESCAPE: the escape character.
 - Place an escape character before a special character that needs to be escaped to indicate that it is not a special character.
 - If ESCAPE is not specified, the default value is the same as QUOTE.
 - You can also use other characters such as backslashes (\) as ESCAPE characters, which is used by MySQL.

Default control characters for TEXT and CSV files

Default control characters for TEXT and CSV files

| Control character | TEXT | CSV |
|-------------------|---------------------------|--------------------------------------|
| DELIMITER | Tab (\t) | Comma (,) |
| QUOTE | double quotation mark (") | double quotation mark (") |
| ESCAPE | N/A | Same as QUOTE |
| NULL | Backslash plus N (\N) | Empty string without quotation marks |

 **Note** All control characters must be single-byte characters.

SDK troubleshooting

The following **Error log information** table lists the error logs generated when an error occurs during the import or export process.

Error log information

| Keyword | Description |
|------------|--|
| code | The HTTP status code of the error request. |
| error_code | The error code returned by OSS. |

| Keyword | Description |
|-----------|---|
| error_msg | The error message returned by OSS. |
| req_id | The UUID used to identify the request. If you require assistance from OSS developers, you can submit a ticket that contains the req_id parameter of the failed request. |

You can handle limitout-related errors by using parameters related to oss_ext.

References

- [Greenplum Database official documentation on external table syntax](#)
- [Greenplum Database official documentation on table creation syntax](#)

12.3.5.2. Import data from MySQL

You can use the mysql2pgsql tool to migrate tables from MySQL to AnalyticDB for PostgreSQL, Greenplum Database, PostgreSQL, or Postgres Plus Advanced Server (PPAS).

Background information

mysql2pgsql connects to both the source MySQL database and the destination AnalyticDB for PostgreSQL database. The tool retrieves the data that you want to export from the source MySQL database, and uses a COPY statement to import the data to the destination database. This tool supports simultaneous data import over multiple threads. Each worker thread imports data from some database tables.

To download the binary installation package of mysql2pgsql, click [here](#).

To view instructions on source code compilation of mysql2pgsql, click [here](#).

Procedure

1. Modify the my.cfg configuration file to configure the connection information of source and destination databases.
 - i. Modify the connection information of the source MySQL database.

 **Note** You must have the read permissions on all user tables.

```
[src.mysql]
host = "192.168.1.1"
port = "3306"
user = "test"
password = "test"
db = "test"
encodingdir = "share"
encoding = "utf8"
```

- ii. Modify the connection information of the destination PostgreSQL, PPAS, or AnalyticDB for PostgreSQL database.

 **Note** You must have the write permissions on the destination table.

```
[desc.pgsql]
connect_string = "host=192.168.1.2 dbname=test port=3432 user=test password=pgsql"
```

2. Use mysql2pgsql to import data.

```
./mysql2pgsql -l <tables_list_file> -d -n -j <number of threads> -s <schema of target table>
```

Parameters

| Parameter | Description |
|-----------|---|
| -l | Optional. This parameter is used to specify a text file that contains tables to be synchronized. If you do not specify this parameter, all the tables in the database that is specified in the configuration file are synchronized. <tables_list_file> specifies the name of a file that contains a collection of tables to be synchronized and conditions for table queries. You can specify the file content in the following format: <pre>table1 : select * from table_big where column1 < '2016-08-05'</pre> <pre>table2 :</pre> <pre>table3</pre> <pre>table4: select column1, column2 from tableX where column1 != 10</pre> <pre>table5: select * from table_big where column1 >= '2016-08-05'</pre> |
| -d | Optional. This parameter indicates the table creation DDL statement that creates the destination table but does not synchronize data. |
| -n | Optional. This parameter must be used along with -d to specify that the table partition definition is not included in the DDL statement. |
| -j | Optional. This parameter is used to specify the number of threads that are used for data synchronization. If you do not specify this parameter, five concurrent threads are used. |
| -s | Optional. This parameter is used to specify the schema of the destination table. Only one schema at a time can be specified by the command. If you do not specify this parameter, the data is imported into the table under the public schema. |

Typical usage

Full database migration

1. Run the following command to obtain the DDL statements of the corresponding destination table:

```
./mysql2pgsql -d
```

2. Create a table in the destination database based on these DDL statements with the distribution key information added.

3. Run the following command to synchronize all tables:

```
./mysql2pgsql
```

This command migrates the data from all MySQL tables in the database that is specified in the configuration file to the destination database. By default, five concurrent threads are used to read and import data from involved tables.

Partial table migration

1. Create a file named tab_list.txt and enter the following content:

```
t1
t2 : select * from t2 where c1 > 138888
```

2. Run the following command to synchronize the specified t1 and t2 tables. For the t2 table, only data that meets the c1 > 138888 condition is migrated.

```
./mysql2pgsql -l tab_list.txt
```

12.3.5.3. Import data from PostgreSQL

You can use the `pgsql2pgsql` tool to migrate tables across AnalyticDB for PostgreSQL, Greenplum Database, PostgreSQL, and Postgres Plus Advanced Server (PPAS).

Context

`pgsql2pgsql` supports the following features:

- Full migration across PostgreSQL, PPAS, Greenplum Database, and AnalyticDB for PostgreSQL.
- Full migration and incremental migration from PostgreSQL or PPAS (version 9.4 or later) to AnalyticDB for PostgreSQL or ApsaraDB RDS for PPAS.

You can download the software packages from the [dbsync project](#) library.

- To download the binary installation package of `pgsql2pgsql`, click [here](#).
- To view instructions on source code compilation of `pgsql2pgsql`, click [here](#).

Procedure

1. Modify the `my.cfg` configuration file to configure the connection information of source and destination databases.
 - i. Modify the connection information of the source PostgreSQL database.

 **Note** In the connection information of the source PostgreSQL database, we recommend that you set the user to the owner of the source database.

```
[src.pgsql]
connect_string = "host=192.168.0.1 dbname=test port=3432 user=test password=pgsql"
```

- ii. Modify the connection information of the on-premises temporary PostgreSQL database.

```
[local.pgsql]
connect_string = "host=192.168.0.2 dbname=test port=3432 user=test2 password=pgsql"
```

- iii. Modify the connection information of the destination PostgreSQL database.

 **Note** You must have the write permissions on the destination table.

```
[desc.pgsql]
connect_string = "host=192.168.0.2 dbname=test port=3432 user=test3 password=pgsql"
```

-  **Note**
- If you want to synchronize incremental data, you must have the permissions to create replication slots in the source database.
 - PostgreSQL 9.4 and later support logic flow replication. Therefore, source databases of these versions support incremental data migration. A kernel supports logic flow replication only after you configure the following parameters:


```
wal_level = logical
max_wal_senders = 6
max_replication_slots = 6
```

2. Use `pgsql2pgsql` to perform full database migration.

```
./pgsql2pgsql
```

By default, the migration program migrates the table data of all users from the source PostgreSQL database to the destination PostgreSQL database.

3. View the status information.

You can view the status information in a single migration process by connecting to the on-premises temporary database. The status information is stored in the `db_sync_status` table and includes the start and end time of full data migration, the start time of incremental data migration, and the status of incremental data synchronization.

12.3.5.4. Use the `\COPY` statement to import data

You can use the `\COPY` statement to import data from local text files to AnalyticDB for PostgreSQL instances. Local text must be properly formatted and include necessary delimiters such as commas (,), semicolons (;), or special characters.

Context

- Parallel writing of large amounts of data is unavailable because the `\COPY` statement writes data in series by using the coordinator node. If you need to import a large amount of data in parallel, you can use the data import method based on Object Storage Service (OSS).
- The `\COPY` statement is a psql instruction. If you use the database statement `COPY` instead of the `\COPY` statement, you must note that only stdin is supported. This `COPY` statement does not support files because the root user does not have the superuser permissions to perform operations on files.
- AnalyticDB for PostgreSQL also allows you to use JDBC to execute the `COPY` statement. The `CopyIn` method is encapsulated within JDBC. For more information, see [Interface CopyIn](#).
- For more information about how to use the `COPY` statement, see [COPY](#).

Procedure

- Import data by using the following sample code:

```
\COPY table [(column [, ...])] FROM {'file' | STDIN}
[ [WITH]
  [OIDS]
  [HEADER]
  [DELIMITER [ AS ] 'delimiter']
  [NULL [ AS ] 'null string']
  [ESCAPE [ AS ] 'escape' | 'OFF']
  [NEWLINE [ AS ] 'LF' | 'CR' | 'CRLF']
  [CSV [QUOTE [ AS ] 'quote']
    [FORCE NOT NULL column [, ...]]
  [FILL MISSING FIELDS]
  [[LOG ERRORS [INTO error_table] [KEEP]
  SEGMENT REJECT LIMIT count [ROWS | PERCENT] ]
\COPY {table [(column [, ...])] | (query)} TO {'file' | STDOUT}
[ [WITH]
  [OIDS]
  [HEADER]
  [DELIMITER [ AS ] 'delimiter']
  [NULL [ AS ] 'null string']
  [ESCAPE [ AS ] 'escape' | 'OFF']
  [CSV [QUOTE [ AS ] 'quote']
    [FORCE QUOTE column [, ...]] ]
  [IGNORE EXTERNAL PARTITIONS ]
```

12.4. Databases

12.4.1. Overview

The operations based on Greenplum Database in AnalyticDB for PostgreSQL are the same as those in Greenplum Database, including their schemas, supported data types, and user permissions. Except for specific operations that are exclusive to Greenplum Database such as those on distribution keys and append-optimized (AO) tables, you can refer to PostgreSQL for other operations.

References

- [Pivotal Greenplum Official Documentation](#)
- [GP 4.3 Best Practice](#)

12.4.2. Create a database

After you log on to the AnalyticDB for PostgreSQL instance, you can execute SQL statements to create databases.

As in PostgreSQL, you can execute SQL statements to create databases in AnalyticDB for PostgreSQL. For example, after psql is connected to Greenplum, execute the following statements:

```
=> create database mygpdb;  
CREATE DATABASE  
=> \c mygpdb  
psql (9.4.4, server 8.3devel)  
You are now connected to database "mygpdb" as user "mygpdb".
```

12.4.3. Create a distribution key

AnalyticDB for PostgreSQL is a distributed database where data is distributed across all the data nodes. You must create distribution keys to distribute the data in AnalyticDB for PostgreSQL. Distribution keys are vital to query performance. Distribution keys are used to ensure **even data distribution**. Appropriate selection of keys can significantly improve query performance.

Specify a distribution key

In AnalyticDB for PostgreSQL, tables can be distributed across all compute nodes in hash or random mode. You must specify a distribution key when you create a table. Imported data is distributed to the specific compute node based on the hash value calculated by the distribution key.

```
=> create table vtbl(id serial, key integer, value text, shape cuboid, location geometry, comment text) distributed by (key  
);  
CREATE TABLE
```

If you do not specify the distribution key, AnalyticDB for PostgreSQL randomly allocates the id field by using the round-robin algorithm. You can specify the distribution key by adding `distributed by (key)` to the table creation statement.

Rules for selecting the distribution key

- Select evenly distributed columns or multiple columns to prevent data skew.
- Select fields that are commonly used for connection operations, especially for highly concurrent statements.
- Select condition columns that feature high concurrency queries and high filterability.
- Do not use random distribution.

12.4.4. Construct data

In some test scenarios, you must construct data to fill the database.

1. Create a function that generates random strings.

```
CREATE OR REPLACE FUNCTION random_string(integer) RETURNS text AS $body$
SELECT array_to_string(array
    (SELECT substring('0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz'
        FROM (ceil(random()*62))::int
        FOR 1)
    FROM generate_series(1, $1), '');
$body$
LANGUAGE SQL VOLATILE;
```

2. Create a partition key.

```
CREATE TABLE tbl(id serial, KEY integer, locate geometry, COMMENT text) distributed by (key);
```

3. Construct data.

```
INSERT INTO tbl(KEY, COMMENT, locate)
SELECT
    KEY,
    COMMENT,
    ST_GeomFromText(locate) AS locate
FROM
    (SELECT
        (a + 1) AS KEY,
        random_string(ceil(random() * 24)::integer) AS COMMENT,
        'POINT(' || ceil(random() * 36 + 99) || ' ' || ceil(random() * 24 + 50) || ') ' AS locate
    FROM
        generate_series(0, 99999) AS a)
AS t;
```

12.4.5. Query data

This topic describes the query statements and how to view the query plans.

Query statement sample

```
=> select * from tbl where key = 751;
| id | key | value | shape | locate | comment |
+----+----+-----+-----+-----+-----+
| 751 | 751 | red | 010100000000000000000000C05B40000000000004A40 | B9hPhjeNWPqV |
(1 row)
Time: 513.101 ms
```

View a query plan

```
=> explain select * from tbl where key = 751;
Gather Motion 1:1 (slice1; segments: 1) (cost=0.00..1519.28 rows=1 width=53)
-> Seq Scan on tbl (cost=0.00..1519.28 rows=1 width=53)
    Filter: key = 751
Settings: effective_cache_size=8GB; gp_statistics_use_fkeys=on
Optimizer status: legacy query optimizer
```

12.4.6. Manage extensions

You can use extensions to expand database features. AnalyticDB for PostgreSQL enables you to manage extensions.

Supported extensions

AnalyticDB for PostgreSQL supports the following extensions:

- PostGIS: processes geographic data.
- MADlib: provides a machine learning function library.
- fuzzystmatch: implements fuzzy match of strings.
- orafunc: provides compatibility with some Oracle functions.
- oss_ext: reads data from Object Storage Service (OSS).
- HyperLogLog: collects statistics.
- PL/Java: compiles user-defined functions (UDFs) in PL/Java.
- pgcrypto: provides cryptographic functions.
- IntArray: provides integer array-related functions, operators, and indexes.

Create an extension

Execute the following statements to create an extension:

```
CREATE EXTENSION <extension name>;  
CREATE SCHEMA <schema name>;  
CREATE EXTENSION IF NOT EXISTS <extension name> WITH SCHEMA <schema name>;
```

Note

Before you create a MADlib extension, you must create a ppythonu extension.

```
CREATE EXTENSION ppythonu;  
CREATE EXTENSION madlib;
```

Delete an extension

Execute the following statements to delete an extension.

```
DROP EXTENSION <extension name>;  
DROP EXTENSION IF EXISTS <extension name> CASCADE;
```

 **Note** If an object depends on an extension that you want to delete, you must add the CASCADE keyword to delete the object.

12.4.7. Manage users and permissions

This topic describes how to manage users and permissions in AnalyticDB for PostgreSQL.

Manage users

The system prompts you to specify an initial username and password when you create an instance. This initial user is the root user. After the instance is created, you can use the root user account to connect to the database. The system also creates superusers such as aurora and replicator for internal management.

You can run the `\du+` command to view the information of all the users after you connect to the database by using the client tool of PostgreSQL or Greenplum. Example:

```
postgres=> \du+
                List of roles
 Role name | Attributes | Member of | Description
-----+-----+-----+-----
 root_user |           | rds_superuser
 ...
```

AnalyticDB for PostgreSQL does not provide superuser permissions, but offers a similar role, RDS_SUPERUSER, which is consistent with the permission system of ApsaraDB RDS for PostgreSQL. The root user, such as root_user in the preceding example, has the permissions of the RDS_SUPERUSER role. This permission attribute can only be identified by viewing the user description.

The root user has the following permissions:

- Create databases and users and perform actions such as LOGIN, excluding the SUPERUSER permissions.
- View and modify the data tables of other users and perform actions such as SELECT, UPDATE, DELETE, and changing owners.
- View the connection information of other users, cancel their SQL statements, and kill their connections.
- Create and delete extensions.
- Create other users with RDS_SUPERUSER permissions. Example:

```
CREATE ROLE root_user2 RDS_SUPERUSER LOGIN PASSWORD 'xyz';
```

Manage permissions

You can manage permissions at the database, schema, and table levels. For example, if you want to grant read permissions on a table to a user and revoke their write permissions, you can execute the following statements:

```
GRANT SELECT ON TABLE t1 TO normal_user1;
REVOKE UPDATE ON TABLE t1 FROM normal_user1;
REVOKE DELETE ON TABLE t1 FROM normal_user1;
```

12.4.8. Manage JSON data

JSON has become a basic data type in the Internet and Internet of things (IoT) fields. For more information about JSON, visit [JSON official website](#). PostgreSQL provides full support for JSON. AnalyticDB for PostgreSQL is optimized by Alibaba Cloud to support the JSON type based on the PostgreSQL syntax.

Check whether the current version supports JSON

Execute the following statement to check whether the current version supports JSON:

```
=> SELECT ''::json;
```

If the following output is displayed, the JSON type is supported and the instance is ready for use. If the operation fails, restart the instance.

```
json
-----
 ""
(1 row)
```

If the following output is displayed, the JSON type is not supported.

```
ERROR: type "json" does not exist
LINE 1: SELECT ''::json;
          ^
```

The preceding command converts data from the string type to the JSON type. PostgreSQL supports operations on JSON data based on this conversion.

JSON conversion in the database

Database operations include reading and writing. The written data is typically converted from the string type to the JSON type. The contents of a string must meet the JSON standard, such as strings, digits, arrays, and objects. Example:

String

```
=> SELECT "'hijson'::json;
      json
-----
'hijson'
(1 row)
```

:: is used for explicit type conversion in PostgreSQL, Greenplum, and AnalyticDB for PostgreSQL. The database calls the input function of the JSON type during the conversion. Therefore, JSON format check is performed.

Example:

```
=> SELECT '{hijson:1024}'::json;
ERROR: invalid input syntax for type json
LINE 1: SELECT '{hijson:1024}'::json;
          ^
DETAIL: Token "hijson" is invalid.
CONTEXT: JSON data, line 1: {hijson...
=>
```

In the preceding example, `hijson` must be enclosed in double quotation marks (`"`) because JSON requires the KEY value to be a string. A syntax error is returned when `{hijson:1024}` is entered.

Apart from explicit type conversion, database records can also be converted to JSON.

Typically, JSON is not used for a string or a digit, but an object that contains one or more key-value pairs. AnalyticDB for PostgreSQL can support most JSON scenarios after data is converted from the string type to the object type. Example:

```
=> select row_to_json(row('{"a":"a"}', 'b'));
      row_to_json
-----
{"f1":{"a":"a"},"f2":"b"}
(1 row)
=> select row_to_json(row('{"a":"a"}::json, 'b'));
      row_to_json
-----
{"f1":{"a":"a"},"f2":"b"}
(1 row)
```

You can see the differences between the string and JSON. The entire record is converted into the JSON type.

JSON data types

- Object

The object is the most frequently used data type in JSON. Example:

```
=> select '{"key":"value"}::json;
      json
-----
{"key":"value"}
(1 row)
```

- Integer and floating point number

JSON supports only three data types for numeric values: integer, floating-point number, and constant expression. AnalyticDB for PostgreSQL supports the three types.

```
=> SELECT '1024'::json;
      json
-----
1024
(1 row)
=> SELECT '0.1'::json;
      json
-----
0.1
(1 row)
```

The following information is required in some special situations:

```
=> SELECT '1e100'::json;
      json
-----
1e100
(1 row)
=> SELECT '{"f":1e100}'::json;
      json
-----
{"f":1e100}
(1 row)
```

Extra-long numbers are also supported. Example:

```
=> SELECT '9223372036854775808'::json;
      json
-----
9223372036854775808
(1 row)
```

- Array

```
=> SELECT '[[1,2],[3,4,5]]::json;
      json
-----
[[1,2],[3,4,5]]
(1 row)
```

Operators

Operators supported by JSON

```

=> select oprname,opcode from pg_operator where oprleft = 3114;
oprname |      opcode
-----+-----
-> | json_object_field
->> | json_object_field_text
-> | json_array_element
->> | json_array_element_text
#> | json_extract_path_op
#>> | json_extract_path_text_op
(6 rows)

```

Basic usage

```

=> SELECT '{"f":"1e100"}::json -> 'f';
?column?
-----
"1e100"
(1 row)
=> SELECT '{"f":"1e100"}::json ->> 'f';
?column?
-----
1e100
(1 row)
=> select '{"f2":{"f3":1},"f4":{"f5":99,"f6":"stringy"}}::json#>array['f4','f6'];
?column?
-----
"stringy"
(1 row)
=> select '{"f2":{"f3":1},"f4":{"f5":99,"f6":"stringy"}}::json#>{f4,f6}';
?column?
-----
"stringy"
(1 row)
=> select '{"f2":["f3",1],"f4":{"f5":99,"f6":"stringy"}}::json#>>'f2,0';
?column?
-----
f3
(1 row)

```

JSON functions

Supported JSON functions

postgres=# \df *json*

| List of functions | | | | |
|-------------------|---------------------------|------------------|---|--------|
| Schema | Name | Result data type | Argument data types | Type |
| pg_catalog | array_to_json | json | anyarray | normal |
| pg_catalog | array_to_json | json | anyarray, boolean | normal |
| pg_catalog | json_array_element | json | from_json json, element_index integer | normal |
| pg_catalog | json_array_element_text | text | from_json json, element_index integer | normal |
| pg_catalog | json_array_elements | SETOF json | from_json json, OUT value json | normal |
| pg_catalog | json_array_length | integer | json | normal |
| pg_catalog | json_each | SETOF record | from_json json, OUT key text, OUT value json | normal |
| pg_catalog | json_each_text | SETOF record | from_json json, OUT key text, OUT value text | normal |
| pg_catalog | json_extract_path | json | from_json json, VARIADIC path_elems text[] | normal |
| pg_catalog | json_extract_path_op | json | from_json json, path_elems text[] | normal |
| pg_catalog | json_extract_path_text | text | from_json json, VARIADIC path_elems text[] | normal |
| pg_catalog | json_extract_path_text_op | text | from_json json, path_elems text[] | normal |
| pg_catalog | json_in | json | cstring | normal |
| pg_catalog | json_object_field | json | from_json json, field_name text | normal |
| pg_catalog | json_object_field_text | text | from_json json, field_name text | normal |
| pg_catalog | json_object_keys | SETOF text | json | normal |
| pg_catalog | json_out | cstring | json | normal |
| pg_catalog | json_populate_record | anyelement | base anyelement, from_json json, use_json_as_text boolean | normal |
| pg_catalog | json_populate_recordset | SETOF anyelement | base anyelement, from_json json, use_json_as_text boolean | normal |
| pg_catalog | json_recv | json | internal | normal |
| pg_catalog | json_send | bytea | json | normal |
| pg_catalog | row_to_json | json | record | normal |
| pg_catalog | row_to_json | json | record, boolean | normal |
| pg_catalog | to_json | json | anyelement | normal |

(24 rows)

Basic usage

```

=> SELECT array_to_json('{{1,5},{99,100}}'::int[]);
 array_to_json
-----
[[1,5],[99,100]]
(1 row)
=> SELECT row_to_json(row(1,'foo'));
 row_to_json
-----
{"f1":1,"f2":"foo"}
(1 row)
=> SELECT json_array_length('[1,2,3,{"f1":1,"f2":[5,6]},4]');
 json_array_length
-----
5
(1 row)
=> select * from json_each('{"f1":[1,2,3],"f2":{"f3":1},"f4":null,"f5":99,"f6":"stringy"}') q;
 key | value
-----+-----
 f1 | [1,2,3]
 f2 | {"f3":1}
 f4 | null
 f5 | 99
 f6 | "stringy"
(5 rows)
=> select json_each_text('{"f1":[1,2,3],"f2":{"f3":1},"f4":null,"f5":null}');
 json_each_text
-----
(f1,"[1,2,3]")
(f2,"{"f3":1}")
(f4,)
(f5,null)
(4 rows)
=> select json_array_elements('[1,true,[1,[2,3]],null,{"f1":1,"f2":[7,8,9]},false]');
 json_array_elements
-----
1
true
[1,[2,3]]
null
{"f1":1,"f2":[7,8,9]}
false
(6 rows)
create type jpop as (a text, b int, c timestamp);
=> select * from json_populate_record(null::jpop,'{"a":"blurfl","x":43.2}', false) q;
 a | b | c
-----+---+---
 blurfl | |
(1 row)
=> select * from json_populate_recordset(null::jpop,'[{"a":"blurfl","x":43.2},{"b":3,"c":"2012-01-20 10:42:53"}]',false) q;
 a | b | c
-----+---+---
 blurfl | |
 | 3 | Fri Jan 20 10:42:53 2012
(2 rows)

```

Code examples

Create a table

```

create table tj(id serial, ary int[], obj json, num integer);
=> insert into tj(ary, obj, num) values('{{1,5}}::int[], '{"obj":1}', 5);
INSERT 0 1
=> select row_to_json(q) from (select id, ary, obj, num from tj) as q;
      row_to_json
-----
{"f1":1,"f2":[1,5],"f3":{"obj":1},"f4":5}
(1 row)
=> insert into tj(ary, obj, num) values('{{2,5}}::int[], '{"obj":2}', 5);
INSERT 0 1
=> select row_to_json(q) from (select id, ary, obj, num from tj) as q;
      row_to_json
-----
{"f1":1,"f2":[1,5],"f3":{"obj":1},"f4":5}
{"f1":2,"f2":[2,5],"f3":{"obj":2},"f4":5}
(2 rows)
    
```

Join multiple tables

```

create table tj2(id serial, ary int[], obj json, num integer);
=> insert into tj2(ary, obj, num) values('{{2,5}}::int[], '{"obj":2}', 5);
INSERT 0 1
=> select * from tj, tj2 where tj.obj->>'obj' = tj2.obj->>'obj';
 id | ary | obj | num | id | ary | obj | num
-----+-----+-----+-----+-----+-----+-----+-----
  2 | {2,5} | {"obj":2} | 5 | 1 | {2,5} | {"obj":2} | 5
(1 row)
=> select * from tj, tj2 where json_object_field_text(tj.obj, 'obj') = json_object_field_text(tj2.obj, 'obj');
 id | ary | obj | num | id | ary | obj | num
-----+-----+-----+-----+-----+-----+-----+-----
  2 | {2,5} | {"obj":2} | 5 | 1 | {2,5} | {"obj":2} | 5
(1 row)
    
```

Use the JSON function index

```

CREATE TEMP TABLE test_json (
  json_type text,
  obj json
);
=> insert into test_json values('aa', '{"f2":{"f3":1},"f4":{"f5":99,"f6":"foo"}}');
INSERT 0 1
=> insert into test_json values('cc', '{"f7":{"f3":1},"f8":{"f5":99,"f6":"foo"}}');
INSERT 0 1
=> select obj->'f2' from test_json where json_type = 'aa';
?column?
-----
{"f3":1}
(1 row)
=> create index i on test_json (json_extract_path_text(obj, '{f4}'));
CREATE INDEX
=> select * from test_json where json_extract_path_text(obj, '{f4}') = '{"f5":99,"f6":"foo"}';
 json_type |      obj
-----+-----
 aa | {"f2":{"f3":1},"f4":{"f5":99,"f6":"foo"}}
(1 row)
    
```

 Note

JSON data cannot be used as the distribution key and does not support JSON aggregate functions.

The following example describes how to use Python to access the database:

```
#!/bin/env python
import time
import json
import psycopg2
def gpquery(sql):
    conn = None
    try:
        conn = psycopg2.connect("dbname=sanity1x2")
        conn.autocommit = True
        cur = conn.cursor()
        cur.execute(sql)
        return cur.fetchall()
    except Exception as e:
        if conn:
            try:
                conn.close()
            except:
                pass
        time.sleep(10)
        print e
    return None
def main():
    sql = "select obj from tj;"
    #rows = Connection(host, port, user, pwd, dbname).query(sql)
    rows = gpquery(sql)
    for row in rows:
        print json.loads(row[0])
if __name__ == "__main__":
    main()
```

12.4.9. Use HyperLogLog

AnalyticDB for PostgreSQL is highly optimized by Alibaba Cloud. It has the features of Greenplum Database and supports HyperLogLog. AnalyticDB for PostgreSQL is suited for industries such as Internet advertising and estimation analysis that require quick estimation of business metrics such as page views (PVs) and unique visitors (UVs).

Create a HyperLogLog extension

Execute the following statement to create a HyperLogLog extension:

```
CREATE EXTENSION hll;
```

Basic types

- Execute the following statement to create a table containing the hll field:

```
create table agg (id int primary key, userids hll);
```

- Execute the following statement to convert int to hll_hashval:

```
select 1::hll_hashval;
```

Basic operators

- The hll type supports =, !=, <>, ||, and #.

```
select hll_add_agg(1::hll_hashval) = hll_add_agg(2::hll_hashval);
select hll_add_agg(1::hll_hashval) || hll_add_agg(2::hll_hashval);
select #hll_add_agg(1::hll_hashval);
```

- The hll_hashval type supports =, !=, and <>.

```
select 1::hll_hashval = 2::hll_hashval;
select 1::hll_hashval <> 2::hll_hashval;
```

Basic functions

- Hash functions such as hll_hash_boolean, hll_hash_smallint, and hll_hash_bigint.

```
select hll_hash_boolean(true);
select hll_hash_integer(1);
```

- hll_add_agg: converts the int format to the hll format.

```
select hll_add_agg(1::hll_hashval);
```

- hll_union: aggregates the hll fields.

```
select hll_union(hll_add_agg(1::hll_hashval),hll_add_agg(2::hll_hashval));
```

- hll_set_defaults: sets the precision.

```
select hll_set_defaults(15,5,-1,1);
```

- hll_print: displays debugging information.

```
select hll_print(hll_add_agg(1::hll_hashval));
```

Examples

```
create table access_date (acc_date date unique, userids hll);
insert into access_date select current_date, hll_add_agg(hll_hash_integer(user_id)) from generate_series(1,10000) t(user_id);
insert into access_date select current_date-1, hll_add_agg(hll_hash_integer(user_id)) from generate_series(5000,20000) t(user_id);
insert into access_date select current_date-2, hll_add_agg(hll_hash_integer(user_id)) from generate_series(9000,40000) t(user_id);
postgres=# select #userids from access_date where acc_date=current_date;
?column?
-----
9725.85273370708
(1 row)
postgres=# select #userids from access_date where acc_date=current_date-1;
?column?
-----
14968.6596883279
(1 row)
postgres=# select #userids from access_date where acc_date=current_date-2;
?column?
-----
29361.5209149911
(1 row)
```

12.4.10. Use the CREATE LIBRARY statement

AnalyticDB for PostgreSQL introduces the CREATE LIBRARY and DROP LIBRARY statements to help you import custom software packages.

Syntax

```
CREATE LIBRARY library_name LANGUAGE [JAVA] FROM oss_location OWNER ownername
CREATE LIBRARY library_name LANGUAGE [JAVA] VALUES file_content_hex OWNER ownername
DROP LIBRARY library_name
```

Parameters

| Parameter | Description |
|------------------|---|
| library_name | The name of the library that you want to install. If the library that you want to install has the same name as an existing library, you must delete the existing library before you install the new one. |
| LANGUAGE [JAVA] | The programming language that you want to use. Only PL/Java is supported. |
| oss_location | The location of the package. You can specify the Object Storage Service (OSS) bucket and object names. Only one object can be specified and the specified object cannot be a compressed file. Sample format: <pre>oss://oss_endpoint filepath=[folder/[folder/]...]/file_name id=userossid key=userosskey bucket=ossbucket</pre> |
| file_content_hex | The content of the file. The byte stream is in hexadecimal notation. For example, 73656c6563742031 indicates the hexadecimal byte stream of "select 1". You can use this syntax to import packages without the need to use OSS. |
| ownername | The user. |
| DROP LIBRARY | Deletes a library. |

Examples

- Example 1: Install a JAR package named analytics.jar.

```
create library example language java from 'oss://oss-cn-hangzhou.aliyuncs.com filepath=analytics.jar id=xxx key=yyy
bucket=zzz';
```

- Example 2: Import the file content with the byte stream in hexadecimal notation.

```
create library pglib LANGUAGE java VALUES '73656c6563742031' OWNER "myuser";
```

- Example 3: Delete a library.

```
drop library example;
```

- Example 4: View installed libraries.

```
select name, lanname from pg_library;
```

12.4.11. Create and use a PL/Java UDF

AnalyticDB for PostgreSQL allows you to compile and upload JAR software packages that are written in PL/Java language, and use these JAR packages to create user-defined functions (UDFs). AnalyticDB for PostgreSQL supports PL/Java 1.5.0 and JVM 1.8. This topic describes how to create a PL/Java UDF. For more information about PL/Java examples, see [PL/Java code](#). For more information about the compiling method, see [PL/Java documentation](#).

Procedure

1. In AnalyticDB for PostgreSQL, execute the following statement to create a PL/Java extension. You need only to execute the statement once for each database.

```
create extension pljava;
```

2. Compile the UDF based on your business requirements. For example, you can use the following code to compile the Test.java file:

```
public class Test
{
    public static String substring(String text, int beginIndex,
        int endIndex)
    {
        try {
            Process process = null;
            process = Runtime.getRuntime().exec("echo Test running");
        } catch (Exception e) {
            return "" + e;
        }
        return text.substring(beginIndex, endIndex);
    }
}
```

3. Compile the manifest.txt file:

```
Manifest-Version: 1.0
Main-Class: Test
Specification-Title: "Test"
Specification-Version: "1.0"
Created-By: 1.7.0_99
Build-Date: 01/20/2016 21:00 AM
```

4. Run the following commands to compile and package the program:

```
javac Test.java
jar cfm analytics.jar manifest.txt Test.class
```

5. Upload the analytics.jar file generated in Step 4 to Object Storage Service (OSS) by using the following OSS console command:

```
osscmd put analytics.jar oss://zzz
```

6. In AnalyticDB for PostgreSQL, execute the CREATE LIBRARY statement to import the file to AnalyticDB for PostgreSQL.

```
create library example language java from 'oss://oss-cn-hangzhou.aliyuncs.com filepath=analytics.jar id=xxx key=yyy bucket=zzz';
```

 **Note** You can only use the filepath variable in the CREATE LIBRARY statement to import files one at a time. The CREATE LIBRARY statement also supports byte streams to import files without the need to use OSS. For more information, see [Use the CREATE LIBRARY statement](#).

7. In AnalyticDB for PostgreSQL, execute the following statements to create and use the UDF:

```
create table temp (a varchar) distributed randomly;
insert into temp values ('my string');
create or replace function java_substring(varchar, int, int) returns varchar as 'Test.substring' language java;
select java_substring(a, 1, 5) from temp;
```

12.5. Table

12.5.1. Create a table

You can create tables within your databases.

Syntax

The following statement shows how to create a table. Not all clauses are required. Use the clauses that can fulfill your business needs.

```
CREATE [[GLOBAL | LOCAL] {TEMPORARY | TEMP}] TABLE table_name (
  [ { column_name data_type [ DEFAULT default_expr ]
    [ column_constraint [ ... ]
  [ ENCODING ( storage_directive [,...] ) ]
  ]
  | table_constraint
  | LIKE other_table [{INCLUDING | EXCLUDING}
    {DEFAULTS | CONSTRAINTS}] ...}
  [, ... ]
)
[ INHERITS ( parent_table [, ... ] ) ]
[ WITH ( storage_parameter=value [, ... ] ) ]
[ ON COMMIT {PRESERVE ROWS | DELETE ROWS | DROP} ]
[ TABLESPACE tablespace ]
[ DISTRIBUTED BY (column, [ ... ] ) | DISTRIBUTED RANDOMLY ]
[ PARTITION BY partition_type (column)
  [ SUBPARTITION BY partition_type (column) ]
  [ SUBPARTITION TEMPLATE ( template_spec ) ]
  [...]
  ( partition_spec
    | [ SUBPARTITION BY partition_type (column) ]
    [...]
  ( partition_spec
    [ ( subpartition_spec
      [ (...) ]
    ) ]
  ) ]
)
```

Definition of the column_constraint clause:

```
[CONSTRAINT constraint_name]
  NOT NULL | NULL
  | UNIQUE [USING INDEX TABLESPACE tablespace]
    [WITH ( FILLFACTOR = value )]
  | PRIMARY KEY [USING INDEX TABLESPACE tablespace]
    [WITH ( FILLFACTOR = value )]
  | CHECK ( expression )
  | REFERENCES table_name [ ( column_name [, ... ] ) ]
    [ key_match_type ]
    [ key_action ]
```

Definition of the storage_directive clause of columns:

```
COMPRESSTYPE={ZLIB|QUICKLZ|RLE_TYPE|NONE}  
[COMPRESLEVEL={0-9}]  
[BLOCKSIZE={8192-2097152}]
```

Definition of the storage_parameter clause of tables:

```
APPENDONLY={TRUE|FALSE}  
BLOCKSIZE={8192-2097152}  
ORIENTATION={COLUMN|ROW}  
CHECKSUM={TRUE|FALSE}  
COMPRESSTYPE={ZLIB|QUICKLZ|RLE_TYPE|NONE}  
COMPRESLEVEL={0-9}  
FILLFACTOR={10-100}  
OIDS={TRUE|FALSE}
```

Definition of the table_constraint clause:

```
[CONSTRAINT constraint_name]  
UNIQUE ( column_name [, ... ] )  
    [USING INDEX TABLESPACE tablespace]  
    [WITH ( FILLFACTOR=value )]  
| PRIMARY KEY ( column_name [, ... ] )  
    [USING INDEX TABLESPACE tablespace]  
    [WITH ( FILLFACTOR=value )]  
| CHECK ( expression )  
| FOREIGN KEY ( column_name [, ... ] )  
    REFERENCES table_name [ ( column_name [, ... ] ) ]  
    [ key_match_type ]  
    [ key_action ]  
    [ key_checking_mode ]
```

Valid values of key_match_type:

```
MATCH FULL  
| SIMPLE
```

Valid values of key_action:

```
ON DELETE  
| ON UPDATE  
| NO ACTION  
| RESTRICT  
| CASCADE  
| SET NULL  
| SET DEFAULT
```

Valid values of key_checking_mode:

```
DEFERRABLE  
| NOT DEFERRABLE  
| INITIALLY DEFERRED  
| INITIALLY IMMEDIATE
```

Valid values of partition_type:

```
LIST
|RANGE
```

Definition of the `partition_specification` clause:

```
partition_element [, ...]
```

Definition of the `partition_element` clause:

```
DEFAULT PARTITION name
|[PARTITION name] VALUES (list_value [,...])
|[PARTITION name]
  START ((datatype) 'start_value') [INCLUSIVE | EXCLUSIVE]
  [ END ((datatype) 'end_value') [INCLUSIVE | EXCLUSIVE] ]
  [ EVERY ((datatype) [number | INTERVAL] 'interval_value') ]
|[PARTITION name]
  END ((datatype) 'end_value') [INCLUSIVE | EXCLUSIVE]
  [ EVERY ((datatype) [number | INTERVAL] 'interval_value') ]
[ WITH ( partition_storage_parameter=value [, ... ] ) ]
[ TABLESPACE tablespace ]
```

Definition of the `subpartition_spec` or `template_spec` clause:

```
subpartition_element [, ...]
```

Definition of the `subpartition_element` clause:

```
DEFAULT SUBPARTITION name
|[SUBPARTITION name] VALUES (list_value [,...])
|[SUBPARTITION name]
  START ((datatype) 'start_value') [INCLUSIVE | EXCLUSIVE]
  [ END ((datatype) 'end_value') [INCLUSIVE | EXCLUSIVE] ]
  [ EVERY ((datatype) [number | INTERVAL] 'interval_value') ]
|[SUBPARTITION name]
  END ((datatype) 'end_value') [INCLUSIVE | EXCLUSIVE]
  [ EVERY ((datatype) [number | INTERVAL] 'interval_value') ]
[ WITH ( partition_storage_parameter=value [, ... ] ) ]
[ TABLESPACE tablespace ]
```

Definition of the `storage_parameter` clause:

```
APPENDONLY={TRUE|FALSE}
BLOCKSIZE={8192-2097152}
ORIENTATION={COLUMN|ROW}
CHECKSUM={TRUE|FALSE}
COMPRESSTYPE={ZLIB|QUICKLZ|RLE_TYPE|NONE}
COMPRESSLEVEL={1-9}
FILLFACTOR={10-100}
OIDS={TRUE|FALSE}
```

Parameters

The [Table creation parameters](#) table describes the key parameters for creating a table.

Table creation parameters

| Parameter | Description |
|-------------------------------|---|
| TABLE_NAME | The name of the table that you want to create. |
| column_name | The name of a column that you want to create in the table. |
| data_type | The data type of the column. For columns that contain textual data, set the data type to VARCHAR or TEXT. We recommend that you do not use the CHAR type. |
| DEFAULT default_expr | Specifies a default value for the column. The system assigns this default value to all columns that do not have a value. The default value can be a variable-free expression. Subqueries or cross-references to other columns in the table are not allowed. The data type of the default expression must match the data type of the column. If a column does not have a default value, this parameter is left empty. |
| ENCODING storage_directive | Specifies the compression type and block size for the column data. This clause is valid only for append-optimized, column-oriented tables. Column compression settings are inherited from the table level to the partition level and then to the sub-partition level. The lowest-level settings have the highest priority. |
| INHERITS | Configures all columns in the new table to automatically inherit a parent table. You can use INHERITS to create a persistent relationship between the new child table and its parent table. Schema modifications to the parent table are applied to the child table as well. When the parent table is also scanned, the data of the child table is scanned as well. |
| LIKE other_table | Specifies a table from which the new table automatically copies all column names, data types, NOT NULL constraints, and distribution policies. Storage properties such as append-optimized or partition structure are not copied. Unlike INHERITS, the new table is completely decoupled from the original table after the new table is created. |
| CONSTRAINT constraint_name | Configures a column or table constraint. If a constraint is violated, the constraint name is displayed in the error message. Constraint names can be used to communicate helpful information to client applications. Constraint names that contain spaces must be enclosed by double quotation marks (''). |
| WITH (storage_option=value) | Configures storage options for the table or its indexes. |
| ON COMMIT | The operation that the system performs on the temporary tables at the end of a transaction. Valid values: <ul style="list-style-type: none"> • PRESERVE ROWS: No special action is taken. The data is retained after the transaction is complete. The data is released only when the session is disconnected. • DELETE ROWS: All rows in the temporary table are deleted. • DROP: The temporary table is deleted. |
| TABLESPACE tablespace | Specifies the name of the tablespace in which you want to create the new table. If this parameter is not specified, the default tablespace of the database is used. |

| Parameter | Description |
|-----------------------|--|
| DISTRIBUTED BY | <p>Configures the distribution policy for the database.</p> <ul style="list-style-type: none"> DISTRIBUTED BY (column, [...]): specifies the distribution key. The system distributes data based on the distribution key. <p>To evenly distribute data, you must set the distribution key to the primary key of the table or a unique column or a set of columns.</p> <ul style="list-style-type: none"> DISTRIBUTED RANDOMLY: randomly distributes data. <div style="background-color: #e0f2f1; padding: 5px; border: 1px solid #ccc;"> <p> Note We recommend that you do not use random distribution.</p> </div> |
| PARTITION BY | <p>Configures a partition key to partition the table. Partitioning large tables improves data access efficiency.</p> <p>To partition a table is to create a top-level (parent) table and multiple lower-level (child) tables. After a partition table is created, its parent table is always empty. Data is stored in the lowest-level child tables. In a multi-level partition table, data is stored only in the lowest-level sub-partitions.</p> <p>Valid values: RANGE, LIST, and a combination of the two.</p> |
| SUBPARTITION BY | Configures a multi-level partition table. |
| SUBPARTITION TEMPLATE | Specifies a sub-partition template to create sub-partitions (lower-level child tables). This ensures that all parent partitions have the same sub-partition structure. |

Examples

Create a table and configure a distribution key. By default, a primary key is used as a distribution key in AnalyticDB for PostgreSQL.

```
CREATE TABLE films (
code  char(5) CONSTRAINT firstkey PRIMARY KEY,
title  varchar(40) NOT NULL,
did    integer NOT NULL,
date_prod date,
kind   varchar(10),
len    interval hour to minute
);
CREATE TABLE distributors (
did    integer PRIMARY KEY DEFAULT nextval('serial'),
name   varchar(40) NOT NULL CHECK (name <> '')
);
```

Create a compressed table and configure a distribution key.

```
CREATE TABLE sales (txn_id int, qty int, date date)
WITH (appendonly=true, compresslevel=5)
DISTRIBUTED BY (txn_id);
```

Use sub-partition templates of each level and the default partition to create a three-level partition table.

```
CREATE TABLE sales (id int, year int, month int, day int,
region text)
DISTRIBUTED BY (id)
PARTITION BY RANGE (year)
SUBPARTITION BY RANGE (month)
SUBPARTITION TEMPLATE (
START (1) END (13) EVERY (1),
DEFAULT SUBPARTITION other_months )
SUBPARTITION BY LIST (region)
SUBPARTITION TEMPLATE (
SUBPARTITION usa VALUES ('usa'),
SUBPARTITION europe VALUES ('europe'),
SUBPARTITION asia VALUES ('asia'),
DEFAULT SUBPARTITION other_regions)
( START (2008) END (2016) EVERY (1),
DEFAULT PARTITION outlying_years);
```

12.5.2. Principles and scenarios of row store, column store, heap tables, and AO tables

AnalyticDB for PostgreSQL supports row store, column store, heap tables, and append-optimized (AO) tables. This topic describes their principles and scenarios.

Row store and column store

Comparison

| Dimension | Row store | Column store |
|---------------------------------------|--|--|
| Definition | Row store stores data in the form of rows. Each row is a tuple. To read a column, you must deform all of the columns that precede it. Therefore, the costs for accessing the first and the last columns are different. | Column store stores data as columns. Each column corresponds to a file or a batch of files. The cost of reading each column is the same. However, if you want to read multiple columns, you must access multiple files. The more columns you access, the higher the overheads are. |
| Compression ratio | Low. | High. |
| Cost of reading a column | The higher the column number, the higher the cost of reading the column. | Same. |
| Vector computing and JIT architecture | Not suitable. Not suitable for batch computing. | Suitable. More efficient when you access and collect statistics of a batch of data. |

| Dimension | Row store | Column store |
|-----------|---|---|
| Scenarios | <p>If you need to perform a large number of update and delete operations due to online transaction processing (OLTP) requirements such as when you query table details where multiple columns are returned, you can use row store.</p> <p>You can use partition tables if you have complex requirements. For example, if you want to partition data based on time, you can use row store to query the details of recent data and use column store to obtain more statistics from historical data.</p> | <p>You can use column store if you need data statistics because of online analytical processing (OLAP) requirements.</p> <p>If you want a higher compression ratio, you can use column store.</p> |

Heap tables

A heap table is heap storage. All changes to the heap table generate redo logs that can be used to restore data by point in time. However, heap tables cannot implement logical incremental backup because all data blocks in the tables can be changed and it is inconvenient to record the position by using heap storage.

Commit and redo logs are used to ensure reliability when transactions are finished. You can also implement redundancy by using secondary nodes through redo logs.

AO tables

AO tables are used to append data for storage. When you delete the updated data, you can use another bitmap file to mark the row that you want to delete and then use the offset of the bit to determine whether the row was deleted.

When the transaction is finished, you must call the `fsync()` function to record the offset of the data block that performs the last write operation. Even if the data block contains only a single record, a new data block is appended for the next transaction. The data block is synchronized to the secondary node for data redundancy.

AO tables are not suitable for small transactions because the `fsync()` function is called at the end of each transaction. This data block is not reused even if space is left.

AO tables are suitable for OLAP scenarios, batch data writing, high compression ratio, and logical backup that supports incremental backup. During backup, you need only to record the offset from the backup and the bitmap deletion mark for each full backup.

Use scenarios of heap tables

- When multiple small transactions are handled, use a heap table.
- When you need to restore data by point in time, use a heap table.

Use scenarios of AO tables

- When you want to use column store, use an AO table.
- When data is written in batches, use an AO table.

12.5.3. Enable the column store and compression features

If you want to improve performance, speed up data import, or reduce costs for tables that have infrequent updates and multiple fields, we recommend that you use column store and compression. These features increase the compression ratio by up to threefold to ensure high performance and speed up data import.

To enable the column store and compression features, you must specify the column store and compression options when you create a table. For example, you can add the following clause to the CREATE statement to enable these features. For more information about the table creation syntax, see [Create a table](#).

```
with (APPENDONLY=true, ORIENTATION=column, COMPRESSTYPE=zlib, COMPRESSLEVEL=5, BLOCKSIZE=1048576, OIDS=false)
```

 **Note** AnalyticDB for PostgreSQL supports only zlib and RLE_TYPE compression algorithms. If you specify the quicklz algorithm, it is automatically converted to zlib.

12.5.4. Add a field to a column store table and set the default value

This topic describes how to add a field to a column store table and set the default value for the field, and how to use the ANALYZE statement to view the impact of updated data on the size of the column store table.

Context

In a column store table, each column is stored as a file, and two columns in the same row correspond to each other by using the offset. For example, if you add two fields of the INT8 type, you can quickly locate column B from column A by using the offset.

When you add the field, AO tables are not rewritten. If an AO table contains the records of deleted data, the added field must be filled with the deleted records before using the offset.

Procedure

1. Create three AO column store tables.

```
postgres=# create table tbl1 (id int, info text) with (appendonly=true, blocksize=8192, compresstype=none, orientation=column);
NOTICE: Table doesn't have 'DISTRIBUTED BY' clause -- Using column named 'id' as the Greenplum Database data distribution key for this table.
HINT: The 'DISTRIBUTED BY' clause determines the distribution of data. Make sure column(s) chosen are the optimal data distribution key to minimize skew.
CREATE TABLE
postgres=# create table tbl2 (id int, info text) with (appendonly=true, blocksize=8192, compresstype=none, orientation=column);
NOTICE: Table doesn't have 'DISTRIBUTED BY' clause -- Using column named 'id' as the Greenplum Database data distribution key for this table.
HINT: The 'DISTRIBUTED BY' clause determines the distribution of data. Make sure column(s) chosen are the optimal data distribution key to minimize skew.
CREATE TABLE
postgres=# create table tbl3 (id int, info text) with (appendonly=true, blocksize=8192, compresstype=none, orientation=column);
NOTICE: Table doesn't have 'DISTRIBUTED BY' clause -- Using column named 'id' as the Greenplum Database data distribution key for this table.
HINT: The 'DISTRIBUTED BY' clause determines the distribution of data. Make sure column(s) chosen are the optimal data distribution key to minimize skew.
CREATE TABLE
```

2. Insert 10 million entries to the first two tables and 20 million entries to the third one.

```
postgres=# insert into tbl1 select generate_series(1,10000000),'test';
INSERT 0 10000000
postgres=# insert into tbl2 select generate_series(1,10000000),'test';
INSERT 0 10000000
postgres=# insert into tbl3 select generate_series(1,20000000),'test';
INSERT 0 20000000
```

3. Analyze the tables and display their sizes.

```
postgres=# analyze tbl1;
ANALYZE
postgres=# analyze tbl2;
ANALYZE
postgres=# analyze tbl3;
ANALYZE
postgres=# select pg_size_pretty(pg_relation_size('tbl1'));
pg_size_pretty
-----
88 MB
(1 row)
postgres=# select pg_size_pretty(pg_relation_size('tbl2'));
pg_size_pretty
-----
88 MB
(1 row)
postgres=# select pg_size_pretty(pg_relation_size('tbl3'));
pg_size_pretty
-----
173 MB
(1 row)
```

4. Update all the data in the first table. Display the table size after the update. The size is twice as large as the size before the update.

```
postgres=# update tbl1 set info='test';
UPDATE 10000000
postgres=# analyze tbl1;
ANALYZE
postgres=# select pg_size_pretty(pg_relation_size('tbl1'));
pg_size_pretty
-----
173 MB
(1 row)
```

5. Add fields to the three tables and set the default values.

```
postgres=# alter table tbl1 add column c1 int8 default 1;
ALTER TABLE
postgres=# alter table tbl2 add column c1 int8 default 1;
ALTER TABLE
postgres=# alter table tbl3 add column c1 int8 default 1;
ALTER TABLE
```

6. Analyze the tables and view the table sizes.

```
postgres=# analyze tbl1;
ANALYZE
postgres=# analyze tbl2;
ANALYZE
postgres=# analyze tbl3;
ANALYZE
postgres=# select pg_size_pretty(pg_relation_size('tbl1'));
pg_size_pretty
-----
325 MB
(1 row)
postgres=# select pg_size_pretty(pg_relation_size('tbl2'));
pg_size_pretty
-----
163 MB
(1 row)
postgres=# select pg_size_pretty(pg_relation_size('tbl3'));
pg_size_pretty
-----
325 MB
(1 row)
```

When you add fields to the AO tables, the number of entries in the existing files will prevail. Even if all the entries are deleted, you must initialize the original data in the newly added fields.

12.5.5. Configure table partitions

For fact tables and large-sized tables in a database, we recommend that you configure table partitions.

Configure table partitions

You can use the `table partitioning` feature to add and remove data based on table partitions on a regular basis. You can use the `ALTER TABLE DROP PARTITION` statement to remove all the data in a partition, and use the `ALTER TABLE EXCHANGE PARTITION` statement to import data to a new data partition.

AnalyticDB for PostgreSQL supports range partitioning, list partitioning, and composite partitioning. Range partitioning supports only the numeric or date time data types of fields.

The following example shows how to use range partitioning in a table:

```
CREATE TABLE LINEITEM (  
  L_ORDERKEY      BIGINT NOT NULL,  
  L_PARTKEY       BIGINT NOT NULL,  
  L_SUPPKEY       BIGINT NOT NULL,  
  L_LINENUMBER   INTEGER,  
  L_QUANTITY      FLOAT8,  
  L_EXTENDEDPRICE FLOAT8,  
  L_DISCOUNT    FLOAT8,  
  L_TAX          FLOAT8,  
  L_RETURNFLAG   CHAR(1),  
  L_LINESTATUS   CHAR(1),  
  L_SHIPDATE     DATE,  
  L_COMMITDATE   DATE,  
  L_RECEIPTDATE  DATE,  
  L_SHIPINSTRUCT CHAR(25),  
  L_SHIPMODE     CHAR(10),  
  L_COMMENT      VARCHAR(44)  
) WITH (APPENDONLY=true, ORIENTATION=column, COMPRESSTYPE=zlib, COMPRESSLEVEL=5, BLOCKSIZE=1048576, OIDS=false) DISTRIBUTED BY (L_orderkey)  
PARTITION BY RANGE (L_SHIPDATE) (START (date '1992-01-01') INCLUSIVE END (date '2000-01-01') EXCLUSIVE EVERY (INTERVAL '1 month' ));
```

Principles of table partitioning

The purpose of partitioning is to minimize the amount of data to be scanned during a query. Therefore, partitions must be associated with the query conditions.

- Principle 1: Select the fields that are related to the query conditions to configure partitions and reduce the amount of data to be scanned.
- Principle 2: If multiple query conditions exist, configure sub-partitions to further reduce the amount of data to be scanned.

12.5.6. Configure the sort key

A sort key is an attribute of a table. Data on disks is stored in the order of the sort key.

Context

Sort keys have two major advantages:

- Speed up and optimize column-store operations. The min and max meta information the system collects seldom overlaps with each other, which features good filterability.
- Eliminate the need to perform ORDER BY and GROUP BY operations. The data directly read from the disk is ordered as required by the sorting conditions.

Create a table

```
Command: CREATE TABLE
Description: define a new table
Syntax:
CREATE [[GLOBAL | LOCAL] {TEMPORARY | TEMP}] TABLE table_name (
  [ { column_name data_type [ DEFAULT default_expr ] [ column_constraint [ ... ]
  [ ENCODING ( storage_directive [,... ] )
  ]
  | table_constraint
  | LIKE other_table [{INCLUDING | EXCLUDING}
    {DEFAULTS | CONSTRAINTS}] ...}
  [, ... ]
  [ column_reference_storage_directive [, ] ]
  )
  [ INHERITS ( parent_table [, ... ] ) ]
  [ WITH ( storage_parameter=value [, ... ] )
  [ ON COMMIT {PRESERVE ROWS | DELETE ROWS | DROP} ]
  [ TABLESPACE tablespace ]
  [ DISTRIBUTED BY (column, [ ... ] ) | DISTRIBUTED RANDOMLY ]
  [ SORTKEY (column, [ ... ] ) ]
  [ PARTITION BY partition_type (column)
    [ SUBPARTITION BY partition_type (column) ]
    [ SUBPARTITION TEMPLATE ( template_spec ) ]
    [...]
  ( partition_spec
    | [ SUBPARTITION BY partition_type (column) ]
    [...]
  ( partition_spec
    [ ( subpartition_spec
      [(...)]
    ) ]
  )
  )
  )
```

Examples:

```
create table test(date text, time text, open float, high float, low float, volume int) with(APPENDONLY=true,ORIENTATION=column) sortkey (volume);
```

Sort the table

```
VACUUM SORT ONLY [tablename]
```

Modify the sort key

This statement only modifies the catalog and does not sort data. You must execute the `VACUUM SORT ONLY` statement to sort the table.

```
ALTER [[GLOBAL | LOCAL] {TEMPORARY | TEMP}] TABLE table_name SET SORTKEY (column, [ ... ] )
```

Examples:

```
alter table test set sortkey (high,low);
```

12.6. Best practices

12.6.1. Configure memory and load parameters

To improve database stability, you must configure memory and load parameters.

Background information

AnalyticDB for PostgreSQL is a massively parallel processing (MPP) database service with high computational and resource requirements. AnalyticDB for PostgreSQL consumes all of its resources. This allows AnalyticDB for PostgreSQL to reach high processing speeds but makes it very easy to reach its limits.

When CPU, network, or hard disk resources are exhausted, the worst that can occur is a hardware bottleneck. However, if all memory resources are completely consumed, the database may not respond.

How to avoid OOM errors

Out of memory (OOM) indicates that the system is unable to provide sufficient memory requested by a process. The following prompt appears when OOM errors occur:

```
Out of memory (seg27 host.example.com pid=47093) VM Protect failed to allocate 4096 bytes, 0 MB available
```

Causes

OOM errors occur due to the following causes:

- Database nodes do not have sufficient memory.
- Kernel parameters related to the memory of the operating system are incorrectly configured.
- Data skew has occurred. This causes a compute node to request a large amount of memory.
- Query skew has occurred. For example, if the grouping fields of some aggregate or window functions are not distribution keys, the data must be redistributed. After the data is redistributed, data is skewed on a specific computer node and results in the node requesting a large amount of memory.

Solutions

1. Modify the queries to request less memory.
2. Use a resource queue of AnalyticDB for PostgreSQL to limit the number of concurrent queries. Reduce the number of queries that are executed within the cluster at the same time to reduce the overall memory requested by the system.
3. Reduce the number of compute nodes deployed on a host. For example, you can deploy 8 compute nodes instead of 16 on a host that has 128 GB of memory to allow each compute node to utilize twice as much memory.
4. Increase the memory of a host.
5. Set the `gp_vmem_protect_limit` parameter to limit the maximum VMEM that can be used by a single compute node. The memory size of a single host and the number of compute nodes deployed on the host determine the maximum memory size that a single compute node can use on average.
6. For SQL statements that have unpredictable memory usage, you can set the `statement_mem` parameter in the session to limit the memory usage of a single SQL statement. This prevents a single SQL statement from consuming all available memory.
7. Set the `statement_mem` parameter at the database level to apply to all the sessions in the database.
8. Use the resource queue of AnalyticDB for PostgreSQL to limit the maximum memory usage of the resource group. Add database users to the resource group to limit the total memory used by these users.

Configure memory-related parameters

You can properly configure the operating system, database parameters, and resource queue to effectively reduce the probability of OOM.

When you calculate the average memory usage of a single compute node on a single host, you must consider both the primary and secondary compute nodes. If the cluster encounters a host failure, the system switches the service from primary compute nodes to the corresponding secondary compute nodes. At this time, the number of compute nodes on the host is greater than usual. Therefore, you must consider the amount of resources occupied by the secondary compute nodes during failover.

The following tables describe how to configure parameters for the operating system kernel and database to avoid OOM.

The following [Operating system kernel parameters](#) table describes the parameter configuration of the operating system kernel.

Operating system kernel parameters

| Parameter | Description |
|----------------------|---|
| huge page | Do not configure the huge page parameter of the system. AnalyticDB for PostgreSQL does not support the latest version of PostgreSQL and therefore does not support the huge page feature. The huge page parameter locks a part of the allocated memory. Database nodes are not able to use this part of the memory. |
| vm.overcommit_memory | <p>If you use the swap space, set this parameter to 2. If you do not use the swap space, set this parameter to 0.</p> <p>Valid values:</p> <ul style="list-style-type: none"> • 0: The requested memory space cannot exceed the difference between the total memory and the resident set size (RSS). An error is returned only when the memory has been exceeded. • 1: Most processes use the malloc function to request memory, but do not use all of the requested memory. When this parameter is set to 1, the memory requested by the malloc function is allocated in all circumstances unless the memory is not sufficient. • 2: The swap space is also considered when the system calculates the memory space that can be requested. You can request a large amount of memory even if the swap space is triggered. |
| overcommit_ratio | <p>The larger the value is, the more memory the process can request. However, less space is reserved for the operating system. For more information about the formula that is used to calculate the memory parameters, see Examples to calculate the memory parameters.</p> <p>When this parameter is set to 2, the memory that can be requested cannot exceed $\text{swap} + \text{memory} \times \text{overcommit_ratio}$.</p> |

The following [Database parameters](#) table describes the parameter configuration of the database.

Database parameters

| Parameter | Description |
|-----------------------|---|
| gp_vmem_protect_limit | The maximum percentage of memory that all processes can request on each node. If the value is too large, it may result in a system OOM error or even more serious problems. If the value is too small, SQL statements may not be executed even when the system has enough memory. |

| Parameter | Description |
|-------------------------------------|---|
| runaway_detector_activation_percent | <p>Default value: 90. Unit: %. When the amount of memory used by a compute node exceeds $\text{runaway_detector_activation_percent} \times \text{gp_vmem_protect_limit}/100$, the query is terminated to prevent OOM.</p> <p>The node remains stopped until the memory usage reaches a value lower than $\text{runaway_detector_activation_percent} \times \text{gp_vmem_protect_limit}/100$.</p> <p>You can use the <code>gp_toolkit.session_level_memory_consumption</code> view to observe the memory usage of each session and runaway information.</p> |
| statement_mem | <p>The maximum amount of memory that a single SQL statement can request. When the maximum memory is exceeded, spill files are created. Default value: 125. Unit: MB.</p> <p>We recommend that you set this parameter based on the following formula:</p> <div style="background-color: #f0f0f0; padding: 5px; margin: 10px 0;"> $(\text{gp_vmem_protect_limit} * 0.9) / \text{max_expected_concurrent_queries}$ </div> <div style="background-color: #e0f2f7; padding: 10px; margin: 10px 0;"> <p> Note</p> <ul style="list-style-type: none"> You can specify the <code>statement_mem</code> parameter in a session. If few concurrent queries exist and if the session needs to execute a query that requires a large amount of memory, you must specify this parameter in the session. The <code>statement_mem</code> parameter is suitable for limiting memory usage in low concurrency scenarios. If you use <code>statement_mem</code> to limit the memory for high concurrency scenarios, each query is allocated with a very small amount of memory. As a result, the performance of a small number of queries with high memory requirements in high concurrency scenarios is affected. We recommend that you use the resource queue to limit the maximum memory usage in high concurrency scenarios. </div> |
| gp_workfile_limit_files_per_query | <p>The maximum number of spill files that can be created by each query. When the amount of memory requested by the query exceeds the <code>statement_mem</code> limit, spill files (also known as work files) are created. Spill files are similar to the swap space of the operating system. When the number of spill files used exceeds the limit, the query is terminated.</p> <p>The default value is 0, which indicates that an unlimited number of spill files can be created.</p> |
| gp_workfile_compress_algorithm | <p>The compression algorithm for spill files. Valid values: NONE and ZLIB.</p> <p>This parameter can optimize storage space or I/O by sacrificing CPU performance. You can set this parameter when the disk space is insufficient or when the spill files encounter a write bottleneck.</p> |

Examples to calculate the memory parameters

The following environment configurations are used in the examples:

- Host configuration:

Total RAM = 256GB
 SWAP = 64GB

- Four hosts, each deployed with eight primary compute nodes and eight secondary compute nodes.

When a host fails, the eight primary compute nodes are distributed to the remaining three hosts. A single host can be deployed with at most three extra primary compute nodes from the failed host. A single host can be deployed with at most 11 primary compute nodes.

1. Calculate the total memory allocated to AnalyticDB for PostgreSQL by the operating system.

Reserve 7.5 GB and 5% of memory for the operating system and calculate the available memory for all applications. Then, divide the available memory by an empirical coefficient of 1.7.

```
gp_vmem = ((SWAP + RAM) - (7.5 GB + 0.05 × RAM)) / 1.7
         = ((64 + 256) - (7.5 + 0.05 * 256)) / 1.7
         = 176
```

2. Use an empirical coefficient of 0.026 to calculate overcommit_ratio.

```
vm.overcommit_ratio = (RAM - (0.026 * gp_vmem)) / RAM
                    = (256 - (0.026 * 176)) / 256
                    = .982
```

Set vm.overcommit_ratio to 98.

3. Calculate the gp_vmem_protect_limit parameter by dividing gp_vmem by maximum_acting_primary_segments. The maximum_acting_primary_segments parameter indicates the number of primary compute nodes to be run on each other host after one host fails.

```
gp_vmem_protect_limit calculation
gp_vmem_protect_limit = gp_vmem / maximum_acting_primary_segments
                    = 176 / 11
                    = 16GB
                    = 16384MB
```

Configure resource queues

You can use resource queues of AnalyticDB for PostgreSQL to limit the number of concurrent queries and the total memory usage. When a query is being executed, it is added to the corresponding queue and the resources used are recorded in the queue. The resource limit of the queue is applied to all sessions in the queue.

The resource queue in AnalyticDB for PostgreSQL is similar to cgroup in Linux.

The following syntax demonstrates how to create a resource queue:

```
Command: CREATE RESOURCE QUEUE
Description: create a new resource queue for workload management
Syntax:
CREATE RESOURCE QUEUE name WITH (queue_attribute=value [, ... ])
where queue_attribute is:
    ACTIVE_STATEMENTS=integer
    [ MAX_COST=float [ COST_OVERCOMMIT={TRUE|FALSE} ]
    [ MIN_COST=float ]
    [ PRIORITY={MIN|LOW|MEDIUM|HIGH|MAX} ]
    [ MEMORY_LIMIT='memory_units' ]
| MAX_COST=float [ COST_OVERCOMMIT={TRUE|FALSE} ]
  [ ACTIVE_STATEMENTS=integer ]
  [ MIN_COST=float ]
  [ PRIORITY={MIN|LOW|MEDIUM|HIGH|MAX} ]
  [ MEMORY_LIMIT='memory_units' ]
```

The [Resource queue creation parameters](#) table describes the parameters for creating the resource queue.

Resource queue creation parameters

| Parameter | Description |
|---|---|
| ACTIVE_STATEMENTS | <p>The number of SQL statements that can be concurrently executed (in the active state). A value of -1 indicates that an unlimited number of SQL statements can be concurrently executed.</p> |
| MEMORY_LIMIT 'memory_units kB, MB or GB' | <p>The maximum memory usage allowed by all SQL statements in the resource queue. A value of -1 indicates unlimited memory usage. However, it is easy to trigger OOM errors because it is limited by the database or system parameters mentioned in the preceding sections.</p> <p>The memory usage of SQL statements is limited by resource queues and parameters.</p> <ul style="list-style-type: none"> When the <code>gp_resqueue_memory_policy</code> parameter is set to <code>none</code>, the limit is the same as that in Greenplum Database earlier than version 4.1. When the <code>gp_resqueue_memory_policy</code> parameter is set to <code>auto</code> and you have specified the <code>statement_mem</code> parameter for a session or at the database level, the allowed memory of a single query exceeds the <code>MEMORY_LIMIT</code> value of the resource queue. <p>Example:</p> <pre>=> SET statement_mem='2GB'; => SELECT * FROM my_big_table WHERE column='value' ORDER BY id; => RESET statement_mem;</pre> <ul style="list-style-type: none"> The system parameter <code>max_statement_mem</code> can limit the maximum memory usage at the compute node level. The memory requested by a single query cannot exceed <code>max_statement_mem</code>. <p>You can modify the <code>statement_mem</code> parameter at the session level, but do not modify the <code>max_statement_mem</code> parameter. We recommend that you specify <code>max_statement_mem</code> in the following formula:</p> <pre>(segghost_physical_memory) / (average_number_concurrent_queries)</pre> <ul style="list-style-type: none"> When the <code>gp_resqueue_memory_policy</code> parameter is set to <code>eager_free</code>, the database allocates the memory by stages. For example, if a query requests 1 GB of memory in total but needs only 100 MB during each stage, the database allocates 100 MB of memory to the query. You can use <code>eager_free</code> to reduce the possibility of insufficient memory for the query. |
| MAX_COST float | <p>The maximum cost of the queries that can be concurrently executed by the resource group. The cost is the estimated total cost in the SQL execution plan.</p> <p>The value of the parameter can be specified as a floating-point number such as 100.0 or an exponent such as 1e+2. A value of -1 indicates that the cost is unlimited.</p> |
| COST_OVERCOMMIT boolean | <p>Specifies whether the limit of <code>max_cost</code> can be exceeded when the system is idle. A value of <code>TRUE</code> indicates that the limit can be exceeded.</p> |
| MIN_COST float | <p>When the resources requested exceed the limit, the queries are queued. However, if the cost of a query is lower than the <code>min_cost</code> value, the query can be executed without queuing.</p> |

| Parameter | Description |
|------------------------------------|--|
| PRIORITY={MIN LOW MEDIUM HIGH MAX} | <p>The priority of the current resource queue. When resources are insufficient, CPU resources are allocated to the resource queue that has a higher priority. The SQL statements in the resource queue that has a higher priority can obtain CPU resources first. We recommend that you allocate users that initiate queries that have high real-time requirements to resource queues that have a higher priority.</p> <p>This parameter is similar to the time slice policy that is used in CPU resource groups in a Linux cgroup to schedule real-time and common tasks.</p> |

Example of modifying resource queue limits:

```
ALTER RESOURCE QUEUE myqueue WITH (MAX_COST=-1.0, MIN_COST= -1.0);
```

Example of putting a user in a resource queue:

```
ALTER ROLE sammy RESOURCE QUEUE poweruser;
```

The following table describes the parameters of resource queues.

Resource queue parameters

| Parameter | Description |
|--|---|
| gp_resqueue_memory_policy | The memory management policy of the resource queue. |
| gp_resqueue_priority | <p>Specifies whether to enable query prioritization. Valid values:</p> <ul style="list-style-type: none"> On Off If this parameter is disabled, the existing priority settings are not evaluated. |
| gp_resqueue_priority_cpucore_per_segment | <p>The number of CPU cores allocated to each compute node. For example, if an 8-core host is configured with two primary compute nodes, you can set the parameter to 4. If no other nodes exist on the primary node, set the parameter to 8.</p> <p>When the CPU is preempted, the SQL statements that are executed in a higher-priority resource group are allocated with CPU resources first.</p> |
| gp_resqueue_priority_sweeper_interval | <p>The interval at which CPU utilization is recalculated for all active statements. The share value is calculated when the SQL statement is executed. You can calculate the share value based on the priority and gp_resqueue_priority_cpucore_per_segment.</p> <p>The smaller the value and the more frequent the calculation, the better the result brought by the priority settings and the larger the overhead.</p> |

Tips for configuring resource queues:

- We recommend that you create a resource queue for each user.

The default resource queue of AnalyticDB for PostgreSQL is pg_default. If no queue is created, all users are assigned to pg_default. This operation is not recommended. We recommend that you create a resource queue for each user. Typically, a database user corresponds to a business. Different database users may correspond to different businesses or users, such as business users, analysts, developers, and DBAs.

- We recommend that you do not use superusers to execute queries.

Queries initiated by superusers are limited only by the preceding parameters and not by the resource queue. We recommend that you do not use superusers to execute queries if you want to use resource queues to limit the use of resources.

- `ACTIVE_STATEMENTS` indicates the SQL statements that can be concurrently executed within a resource queue. When the cost of a query is lower than the minimum cost specified by `min_cost`, the query can be executed without being queued.
- You can specify the `MEMORY_LIMIT` parameter to set the allowed maximum memory usage of all the SQL statements in a resource queue. The `statement_mem` parameter has a higher priority that can break through the limit of resource queues.

 **Note** The memory of all resource queues cannot exceed `gp_vmem_protect_limit`.

- You can distinguish businesses by configuring the priorities of resource queues.

For example, the report-related business has a top priority, while common businesses and analysts have lower priorities. In this case, you can create three resource queues that have the max, high, and medium priorities, respectively.

- If the number of resources requested at different periods of time varies, you can run the crontab command to periodically adjust the limits of resource queues based on usage patterns.

For example, the queue of analysts has a top priority during the day, while the queue of the report-related business has a lower priority at night. AnalyticDB for PostgreSQL does not allow you to configure resource limits by period of time. Therefore, you can only externally deploy tasks by using the `ALTER RESOURCE QUEUE` statement.

- You can use the view provided by `gp_toolkit` to view the resource usage of the resource queues.

```
gp_toolkit.gp_resq_activity
gp_toolkit.gp_resq_activity_by_queue
gp_toolkit.gp_resq_priority_backend
gp_toolkit.gp_resq_priority_statement
gp_toolkit.gp_resq_role
gp_toolkit.gp_resqueue_status
```

13.KVStore for Redis

13.1. What is KVStore for Redis?

KVStore for Redis is a database service that is compatible with Redis-native protocols. KVStore for Redis provides a highly available hot standby architecture and can scale to meet the requirements of high-performance and low-latency read/write operations.

Features

- KVStore for Redis supports various data types, such as strings, lists, sets, sorted sets, hash tables, and streams. This service also supports advanced features, such as transactions, message subscription, and message publishing.
- KVStore for Redis Enhanced Edition (Tair), which is a key-value pair cloud caching service, is an advanced version of KVStore for Redis Community Edition.

Instance type

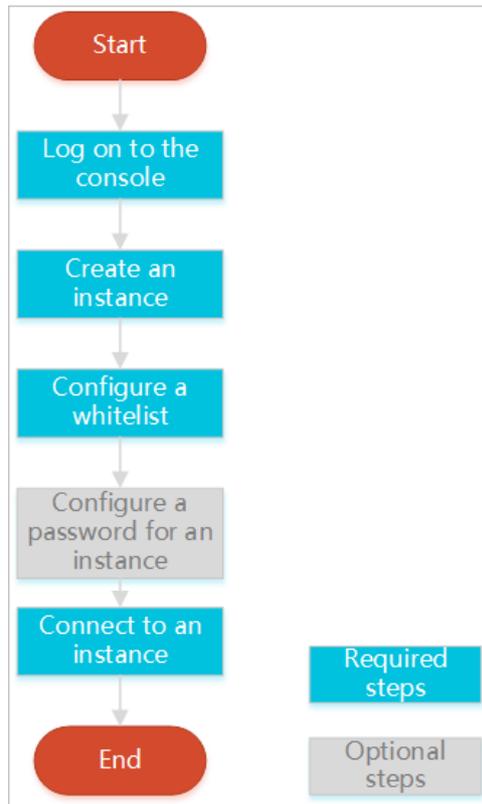
| Instance type | Overview |
|---|---|
| Community Edition | KVStore for Redis Community Edition is compatible with the data caching service of Redis-native engines. This service supports master-replica instances and cluster instances. |
| Performance-enhanced instances of KVStore for Redis Enhanced Edition (Tair) | KVStore for Redis Enhanced Edition (Tair) provides a multi-threading model and integrates some features of Alibaba Tair. KVStore for Redis Enhanced Edition (Tair) supports multiple data structures of Tair and is suitable for diverse scenarios. |

13.2. Quick Start

13.2.1. Get started with KVStore for Redis

This topic describes all operations that you can perform on an instance from instance creation to database logon. This topic helps you understand how to create and manage an instance.

The following figure shows how to manage a KVStore for Redis instance.



- **Log on to the KVStore for Redis console**
This topic describes how to log on to the KVStore for Redis console.
- **Create an instance**
KVStore for Redis supports classic networks and virtual private clouds (VPCs). You can create KVStore for Redis instances in these networks.
- **Configure a whitelist**
Before you use a KVStore for Redis instance, add IP addresses or CIDR blocks that are used to access the database to the whitelist of the instance to improve the security and stability of the database.
- If you do not specify a password when you create the instance, specify a password on the **Instance Information** page.
- **Connect to a KVStore for Redis instance**
To connect to the KVStore for Redis instance, you can use a client that supports Redis protocols or use the Redis command-line interface (redis-cli) tool.

13.2.2. Log on to the KVStore for Redis console

This topic describes how to log on to the KVStore for Redis console.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

Note When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top navigation bar, choose **Products > KVStore for Redis**.

13.2.3. Create a KVStore for Redis instance

This topic describes how to create a KVStore for Redis instance in the KVStore for Redis console.

Procedure

1. [Log on to the KVStore for Redis console](#).
2. Click **Create Instance** in the upper-right corner of the page.
3. Set the following parameters.

KVStore for Redis instance parameters

| Section | Parameter | Description |
|----------------|-------------------|--|
| Basic Settings | Organization | Select the organization to which the KVStore for Redis instance that you want to create belongs. |
| | Resource Set | Select the resource set to which the instance belongs. Notice After you select a resource set, the instance is accessible only to the members of the specified resource set. |
| Region | Region | Select the region to which the instance belongs. |
| | Zone | Select the zone to which the instance belongs. |
| Specifications | Engine Version | Select the engine version of the KVStore for Redis instance that you want to create. |
| | Architecture Type | Select the architecture type of the instance. KVStore for Redis provides cluster and standard architectures. The cluster architecture is applicable to scenarios that require a large capacity or high performance. Redis-native databases run in a single-threading model. If your database does not require high performance, we recommend that you use a standard instance. If your database requires high performance, select a cluster architecture. |
| | Node Type | Master-replica is automatically selected. The node type has one master node and one replica node. |

| Section | Parameter | Description |
|-----------------|-------------------------|--|
| | Instance Type | Select the specification of the instance. The maximum number of connections and maximum internal bandwidth vary based on the instance specification. For more information, see <i>Instance specifications</i> in <i>Product Introduction</i> . |
| Network | Network Type | <ul style="list-style-type: none"> Classic network: Cloud services in a classic network are not isolated. Unauthorized access to a cloud service is blocked by security groups or the service whitelist. VPC: A virtual private cloud (VPC) helps you build an isolated network environment on Apsara Stack. You can customize the route table, CIDR block, and gateway of a VPC. You can also migrate applications to the cloud without service interruption. You can use an Express Connect circuit or VPN to connect data centers to cloud resources in a VPC. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p> Note Before you can select VPC, you must create a VPC. For more information, see <i>Create a VPC</i> and <i>Create a vSwitch</i> in <i>VPC User Guide</i>.</p> </div> |
| Password | Instance Name | Enter a name that can help you identify the instance. <ul style="list-style-type: none"> The name must be 2 to 128 characters in length. The name can contain letters, digits, underscores (_), and hyphens (-). The name must start with a letter. |
| | Password Setting | You can select Set Now or Set after Purchase . |
| | Logon Password | Enter a password to access the instance. The password must meet the following requirements: <ul style="list-style-type: none"> The password must be 8 to 30 characters in length. The password must contain uppercase letters, lowercase letters, and digits. The password must not contain special characters. |
| | Confirm Password | Enter the specified password again. |

4. After you configure the parameters, click **Submit**.

13.2.4. Configure a whitelist

Before you use a KVStore for Redis instance, add IP addresses or CIDR blocks that are used to access the database to the whitelist of the instance to improve the security and stability of the database.

Context

 **Note** A properly configured whitelist can ensure a higher level of security protection for your KVStore for Redis instance. We recommend that you maintain the whitelist on a regular basis.

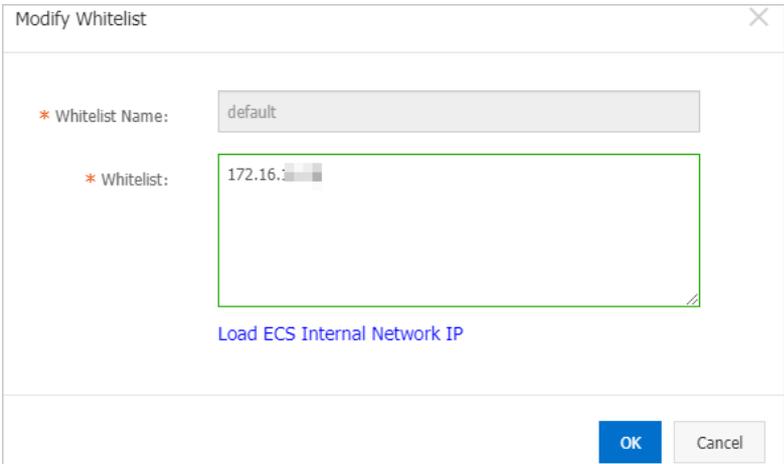
Procedure

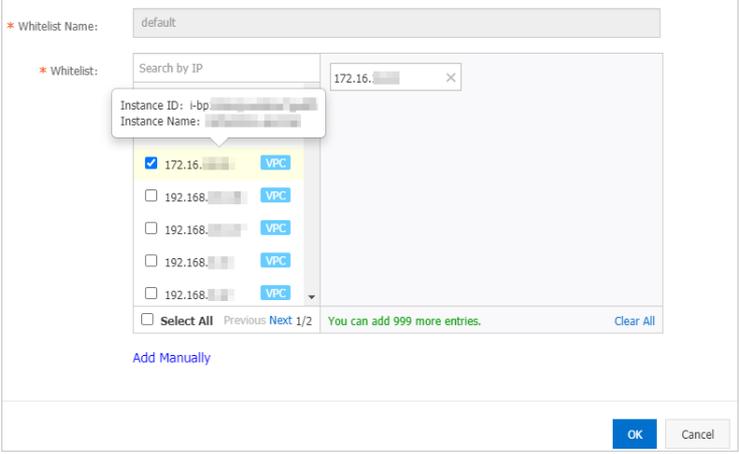
1. [Log on to the KVStore for Redis console](#).
2. On the **Instance List** page, click the ID of the instance.

- In the left-side navigation pane, click **Whitelist Settings**.
- Find the IP address whitelist that you want to manage and click **Modify**.

Note You can also click **Add Whitelist** to create an IP address whitelist. The name of the IP address whitelist must be 2 to 32 characters in length and can contain lowercase letters, digits, and underscores (_). The name of the whitelist must start with a lowercase letter and end with a lowercase letter or digit.

- In the dialog box that appears, perform one of the following operations:

| Action | Procedure |
|---|--|
| <p>Manually modify the IP address whitelist</p> | <p>Enter IP addresses or CIDR blocks.</p> <p>Manually modify the IP address whitelist</p>  <p>Note</p> <ul style="list-style-type: none"> Separate multiple IP addresses with commas (,). You can add up to 1,000 unique IP addresses. Supported formats are specific IP addresses such as 10.23.12.24 and CIDR blocks such as 10.23.12.24/24. /24 indicates the length of the IP address prefix. An IP address prefix can be 1 to 32 bits in length. If you set the prefix length to 0, for example, 0.0.0.0/0 or 127.0.0.1/0, all IP addresses are allowed to access the instance. In this case, the security risk of your instance is high. Proceed with caution. |

| Action | Procedure |
|---|--|
| <p>Add private IP addresses of ECS instances to an IP address whitelist</p> | <p>i. Click Load ECS Internal Network IP.</p> <p>ii. Select IP addresses based on your business requirements.</p> <p>Add private IP addresses of ECS instances</p>  <p>Note To find the ECS instance to which a specific IP address is assigned, you can move the pointer over the IP address. Then, the system displays the ID and name of the ECS instance to which the IP address is assigned.</p> |
| <p>Remove IP addresses from the IP address whitelist</p> | <p>To remove all IP addresses from the IP address whitelist and retain the IP address whitelist, click Delete.</p> |

6. Then, click **OK**.

13.2.5. Connect to an instance

13.2.5.1. Use a Redis client

KVStore for Redis is compatible with open source Redis. You can connect to KVStore for Redis and open source Redis in a similar manner. Therefore, you can use a client that is compatible with the Redis protocols to connect to KVStore for Redis. You can connect to a KVStore for Redis instance by using clients of different programming languages.

Prerequisites

The private IP address of an Elastic Compute Service (ECS) instance or the public IP address of an on-premises machine is added to a whitelist of the instance. For more information, see [Configure a whitelist](#).

Obtain connection information

When you use a client to connect to a KVStore for Redis instance, you must obtain the following information and specify the information in the code:

| Information | Description |
|-------------|-------------|
|-------------|-------------|

| Information | Description |
|--|--|
| Endpoint | <p>You can find the endpoint in the Connection Information section on the Instance Information page.</p> <div style="background-color: #e0f2f1; padding: 5px;"> <p> Note KVStore for Redis instances support multiple types of endpoints. We recommend that you use endpoints in a VPC for higher security and lower network latency.</p> </div> |
| Port number | The default port number is 6379. |
| The account of the instance. This parameter is not required by some clients. | By default, a KVStore for Redis instance provides a database account that is named after the instance ID, such as, r-bp10noxlhcoim2****. |
| The password of the account. | If you forget your password, you can reset the password. For more information, see Change the password . |

Commonly used clients

For the list of clients supported by Redis, see [Redis clients](#).

- [Jedis client](#)
- [PhpRedis client](#)
- [Redis-py client](#)
- [C or C++ client](#)
- [.NET client](#)
- [Node-redis client](#)
- [C# client](#) [StackExchange.Redis](#)

Jedis client

1. Download and install the Jedis client. For more information, see [Jedis](#).
2. Select a connection method to meet your business requirements.
 - o JedisPool-based connection: This connection method is recommended.
 - a. Launch the Eclipse client, create a project, and then configure the following pom file:

```

<dependency>
<groupId>redis.clients</groupId>
<artifactId>jedis</artifactId>
<version>2.7.2</version>
<type>jar</type>
<scope>compile</scope>
</dependency>
```

- b. Enter the following code in the project to add the relevant applications:

```
import org.apache.commons.pool2.PooledObject;
import org.apache.commons.pool2.PooledObjectFactory;
import org.apache.commons.pool2.impl.DefaultPooledObject;
import org.apache.commons.pool2.impl.GenericObjectPoolConfig;
import redis.clients.jedis.HostAndPort;
import redis.clients.jedis.Jedis;
import redis.clients.jedis.JedisPool;
import redis.clients.jedis.JedisPoolConfig;
```

- c. Enter the following code in the project based on the Jedis client version and modify the code based on the comments.

 **Note** For more information about how to obtain the connection endpoint and password of the KVStore for Redis instance, see [Obtain connection information](#).

■ Jedis 2.7.2

```
JedisPoolConfig config = new JedisPoolConfig();
//Maximum number of idle connections. You can configure this parameter. Make sure that the specified value does not exceed the maximum number of connections that the KVStore for Redis instance supports.
config.setMaxIdle(200);
//Maximum number of connections. You can configure this parameter. Make sure that the specified value does not exceed the maximum number of connections that the KVStore for Redis instance supports.
config.setMaxTotal(300);
config.setTestOnBorrow(false);
config.setTestOnReturn(false);
String host = "*.aliyuncs.com";
String password = "Password";
JedisPool pool = new JedisPool(config, host, 6379, 3000, password);
Jedis jedis = null;
try {
jedis = pool.getResource();
/// ... do stuff here ... for example
jedis.set("foo", "bar");
String foobar = jedis.get("foo");
jedis.zadd("sose", 0, "car");
jedis.zadd("sose", 0, "bike");
Set<String> sose = jedis.zrange("sose", 0, -1);
} finally {
if (jedis != null) {
jedis.close();
}
}
/// ... when closing your application:
pool.destroy();
```

■ Jedis 2.6 or Jedis 2.5

```
JedisPoolConfig config = new JedisPoolConfig();
//Maximum number of idle connections. You can configure this parameter. Make sure that the specified value does not exceed the maximum number of connections that the KVStore for Redis instance supports.
config.setMaxIdle(200);
//Maximum number of connections. You can configure this parameter. Make sure that the specified value does not exceed the maximum number of connections that the KVStore for Redis instance supports.
config.setMaxTotal(300);
config.setTestOnBorrow(false);
config.setTestOnReturn(false);
String host = "*.aliyuncs.com";
String password = "Password";
JedisPool pool = new JedisPool(config, host, 6379, 3000, password);
Jedis jedis = null;
boolean broken = false;
try {
    jedis = pool.getResource();
    /// ... do stuff here ... for example
    jedis.set("foo", "bar");
    String foobar = jedis.get("foo");
    jedis.zadd("sose", 0, "car");
    jedis.zadd("sose", 0, "bike");
    Set<String> sose = jedis.zrange("sose", 0, -1);
}
catch(Exception e)
{
    broken = true;
} finally {
    if (broken) {
        pool.returnBrokenResource(jedis);
    } else if (jedis != null) {
        pool.returnResource(jedis);
    }
}
```

- Single Jedis connection: This connection method does not allow a client to automatically reconnect to the KVStore for Redis instance after a connection times out. Therefore, this connection is not recommended.

Launch the Eclipse client, create a project, enter the following code, and then modify the code based on the comments.

 **Note** For more information about how to obtain the connection endpoint and password of the KVStore for Redis instance, see [Obtain connection information](#).

```
import redis.clients.jedis.Jedis;
public class jedistest {
public static void main(String[] args) {
try {
String host = "xx.kvstore.aliyuncs.com";//You can find the connection address in the console.
int port = 6379;
Jedis jedis = new Jedis(host, port);
//Authentication information.
jedis.auth("password");//password
String key = "redis";
String value = "aliyun-redis";
//Select a database. Default value: 0.
jedis.select(1);
//Specify a key.
jedis.set(key, value);
System.out.println("Set Key " + key + " Value: " + value);
//Obtain the configured key and value.
String getvalue = jedis.get(key);
System.out.println("Get Key " + key + " ReturnValue: " + getvalue);
jedis.quit();
jedis.close();
}
catch (Exception e) {
e.printStackTrace();
}
}
}
```

3. Run the project. If Eclipse returns the following result, it indicates that the client is connected to the KVStore for Redis instance.

```
Set Key redis Value aliyun-redis
Get Key redis ReturnValue aliyun-redis
```

PhpRedis client

1. Download and install the PhpRedis client. For more information, see [PhpRedis](#).
2. Enter the following code in a PHP editor and modify the code based on the comments.

 **Note** For more information about how to obtain the connection address, account, and password of the KVStore for Redis instance, see [Obtain connection information](#).

```
<?php
/* Replace the parameter values with the endpoint and port number of the instance. */
$host = "r-bp10noxlhcoim2****.redis.rds.aliyuncs.com";
$port = 6379;
/* Replace the following parameter values with the ID and password of the instance. */
$user = "test_username";
$password = "test_password";
$redis = new Redis();
if($redis->connect($host, $port) == false) {
    die($redis->getLastError());
}
if($redis->auth($password) == false) {
    die($redis->getLastError());
}
/* You can manage the database after you pass the authentication. For more information, visit https://github.com/phpRedis/phpredis. */
if($redis->set("foo", "bar") == false) {
    die($redis->getLastError());
}
$value = $redis->get("foo");
echo $value;
?>
```

3. Run the preceding code to connect to the instance.

For more information, see [PhpRedis](#).

Redis-py client

1. Download and install the redis-py client. For more information, see [redis-py](#).
2. Enter the following code in a Python editor and modify the code based on the comments.

 **Note** For more information about how to obtain the connection endpoint and password of the KVStore for Redis instance, see [Obtain connection information](#).

```
#!/usr/bin/env python
#-*- coding: utf-8 -*-
import redis
#Replace the value of the host parameter with the endpoint of the instance and replace the value of the port parameter with the port number.
host = 'localhost'
port = 6379
#Replace the following parameter value with the password of the instance.
pwd = 'test_password'
r = redis.StrictRedis(host=host, port=port, password=pwd)
#You can perform database operations after you establish a connection. For more information, visit https://github.com/andymccurdy/redis-py.
r.set('foo', 'bar');
print r.get('foo')
```

3. Run the preceding code to connect to the instance.

C or C++ client

1. Run the following commands to download, compile, and install the C client:

```
git clone https://github.com/redis/hiredis.git
cd hiredis
make
sudo make install
```

2. Enter the following code in a C or C++ editor and modify the code based on the comments.

 **Note** For more information about how to obtain the connection endpoint and password of the KVStore for Redis instance, see [Obtain connection information](#).

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <hiredis.h>
int main(int argc, char **argv) {
    unsigned int j;
    redisContext *c;
    redisReply *reply;
    if (argc < 4) {
        printf("Usage: example xxx.kvstore.aliyuncs.com 6379 instance_id password\n");
        exit(0);
    }
    const char *hostname = argv[1];
    const int port = atoi(argv[2]);
    const char *instance_id = argv[3];
    const char *password = argv[4];
    struct timeval timeout = { 1, 500000 }; // 1.5 seconds
    c = redisConnectWithTimeout(hostname, port, timeout);
    if (c == NULL || c->err) {
        if (c) {
            printf("Connection error: %s\n", c->errstr);
            redisFree(c);
        } else {
            printf("Connection error: can't allocate redis context\n");
        }
        exit(1);
    }
    /* AUTH */
    reply = redisCommand(c, "AUTH %s", password);
    printf("AUTH: %s\n", reply->str);
    freeReplyObject(reply);
    /* PING server */
    reply = redisCommand(c, "PING");
    printf("PING: %s\n", reply->str);
    freeReplyObject(reply);
    /* Set a key */
    reply = redisCommand(c, "SET %s %s", "foo", "hello world");
    printf("SET: %s\n", reply->str);
    freeReplyObject(reply);
    /* Set a key using binary safe API */
    reply = redisCommand(c, "SET %b %b", "bar", (size_t) 3, "hello", (size_t) 5);
    printf("SET (binary API): %s\n", reply->str);
    freeReplyObject(reply);
    /* Try a GET and two INCR */
    reply = redisCommand(c, "GET foo");
    printf("GET foo: %s\n", reply->str);
    freeReplyObject(reply);
    reply = redisCommand(c, "INCR counter");
    printf("INCR counter: %lld\n", reply->integer);
```

```

freeReplyObject(reply);
/* again ... */
reply = redisCommand(c,"INCR counter");
printf("INCR counter: %lld\n", reply->integer);
freeReplyObject(reply);
/* Create a list of numbers, from 0 to 9 */
reply = redisCommand(c,"DEL mylist");
freeReplyObject(reply);
for (j = 0; j < 10; j++) {
    char buf[64];
    snprintf(buf,64,"%d",j);
    reply = redisCommand(c,"LPUSH mylist element-%s", buf);
    freeReplyObject(reply);
}
/* Let's check what we have inside the list */
reply = redisCommand(c,"LRANGE mylist 0 -1");
if (reply->type == REDIS_REPLY_ARRAY) {
    for (j = 0; j < reply->elements; j++) {
        printf("%0u %s\n", j, reply->element[j]->str);
    }
}
freeReplyObject(reply);
/* Disconnects and frees the context */
redisFree(c);
return 0;
}

```

3. Compile the code.

```
gcc -o example -g example.c -I /usr/local/include/hiredis -lhiredis
```

4. Perform a test run and connect to the instance.

```
example xxx.kvstore.aliyuncs.com 6379 instance_id password
```

.NET client

 **Warning** If you need to switch or select a database from multiple databases in a cluster instance, you must set the `cluster_compat_enable` parameter to `0` and restart the client application. This disables the support for the cluster syntax of open source Redis. Otherwise, the system sends the following error message: `Multiple databases are not supported on this server; cannot switch to database`. For more information, see [Parameter configuration](#).

1. Run the following command to download the .NET client.

```
git clone https://github.com/ServiceStack/ServiceStack.Redis
```

2. Create a .NET project on the .NET client.
3. Add a reference. The reference file is stored in the library file directory `ServiceStack.Redis/lib/tests`.
4. Enter the following code in the .NET project and modify the code based on the comments. For more information, see [ServiceStack.Redis](#).

 **Note** For more information about how to obtain the connection endpoint and password of the KVStore for Redis instance, see [Obtain connection information](#).

```

using System;
using System.Collections.Generic;
using System.Linq;

```

```

using System.Text;
using System.Threading.Tasks;
using ServiceStack.Redis;
namespace ServiceStack.Redis.Tests
{
    class Program
    {
        public static void RedisClientTest()
        {
            string host = "127.0.0.1"; /*The endpoint of the host.*/
            string password = "password"; /*The password*/
            RedisClient redisClient = new RedisClient(host, 6379, password);
            string key = "test-aliyun";
            string value = "test-aliyun-value";
            redisClient.Set(key, value);
            string listKey = "test-aliyun-list";
            System.Console.WriteLine("set key " + key + " value " + value);
            string getValue = System.Text.Encoding.Default.GetString(redisClient.Get(key));
            System.Console.WriteLine("get key " + key);
            System.Console.Read();
        }
        public static void RedisPoolClientTest()
        {
            string[] testReadWriteHosts = new[] {
                "redis://password@127.0.0.1:6379" /* redis://password@endpoint:port */
            };
            RedisConfig.VerifyMasterConnections = false; /*Required.*/
            PooledRedisClientManager redisPoolManager = new PooledRedisClientManager(10 /*Number of connection pools*/,
                10 /*Connection pool timeout value*/, testReadWriteHosts);
            for (int i = 0; i < 100; i++){
                IRedisClient redisClient = redisPoolManager.GetClient(); /*Obtain the connection.*/
                RedisNativeClient redisNativeClient = (RedisNativeClient)redisClient;
                redisNativeClient.Client = null; /*KVStore for Redis does not support the CLIENT SETNAME command. Set Client to
                null.*/
                try
                {
                    string key = "test-aliyun1111";
                    string value = "test-aliyun-value1111";
                    redisClient.Set(key, value);
                    string listKey = "test-aliyun-list";
                    redisClient.AddItemToList(listKey, value);
                    System.Console.WriteLine("set key " + key + " value " + value);
                    string getValue = redisClient.GetValue(key);
                    System.Console.WriteLine("get key " + key);
                    redisClient.Dispose(); /*
                }
                catch (Exception e)
                {
                    System.Console.WriteLine(e.Message);
                }
            }
            System.Console.Read();
        }
        static void Main(string[] args)
        {
            /*Single-connection mode.*/
            RedisClientTest();
            /*Connection-pool mode.*/
            RedisPoolClientTest();
        }
    }
}

```

```
}
```

Node-redis client

1. Download and install the node-redis client.

```
npm install hiredis redis
```

2. Enter the following code in the node-redis client and modify the code based on the comments.

 **Note** For more information about how to obtain the connection endpoint and password of the KVStore for Redis instance, see [Obtain connection information](#).

```
var redis = require("redis"),
    client = redis.createClient(<port>, <"host">, {detect_buffers: true});
client.auth("password", redis.print)
```

Parameters:

- o <port>: the service port number of the KVStore for Redis database. The default port number is 6379.
- o <"host">: the endpoint of the KVStore for Redis instance.

Configuration examples:

```
var redis = require("redis"),
    client = redis.createClient(6379, "r-abcdefg.redis.rds.aliyuncs.com", {detect_buffers: true});
client.auth("password", redis.print)
```

3. Run the preceding code to connect to the KVStore for Redis instance.
4. Use KVStore for Redis.

```
//Write data to the instance.
client.set("key", "OK");
//Query data on the instance. The returned data is of the STRING type.
client.get("key", function (err, reply) {
  console.log(reply.toString()); // print `OK`
});
//If the input parameter is a buffer, the returned value is also a buffer.
client.get(new Buffer("key"), function (err, reply) {
  console.log(reply.toString()); // print `<Buffer 4f 4b>`
});
client.quit();
```

C# client StackExchange.Redis

 **Warning** If you need to switch or select a database from multiple databases in a cluster instance, you must set the cluster_compat_enable parameter to 0 and restart the client application. This disables the support for the cluster syntax of open source Redis. Otherwise, the system sends the following error message: RedisCommandException: Multiple databases are not supported on this server; cannot switch to database: 1. For more information, see [Parameter configuration](#).

1. Download and install the [StackExchange.Redis](#) client.
2. Add a reference.

```
using StackExchange.Redis;
```

3. Initialize ConnectionMultiplexer.

ConnectionMultiplexer is the core of StackExchange.Redis and is shared and reused in the entire application. You must use ConnectionMultiplexer as a singleton. ConnectionMultiplexer is initialized in the following method:

 **Note**

- For more information about how to obtain the connection endpoint and password of the KVStore for Redis instance, see [.](#)
- ConfigurationOptions contains multiple options, such as keepAlive, connectRetry, and name. For more information, see [ConfigurationOptions](#).

```
// redis config
private static ConfigurationOptions configurationOptions = ConfigurationOptions.Parse("127.0.0.1:6379,password=xxx,connectTimeout=2000");
//the lock for singleton
private static readonly object Locker = new object();
//singleton
private static ConnectionMultiplexer redisConn;
//singleton
public static ConnectionMultiplexer getRedisConn()
{
    if (redisConn == null)
    {
        lock (Locker)
        {
            if (redisConn == null || !redisConn.IsConnected)
            {
                redisConn = ConnectionMultiplexer.Connect(configurationOptions);
            }
        }
    }
    return redisConn;
}
```

4. GetDatabase() returns a light weight object. You can obtain this object from the object of ConnectionMultiplexer.

```
redisConn = getRedisConn();
var db = redisConn.GetDatabase();
```

5. You can use the client to perform database operations.

 **Note** The following examples describe the commands for common data types. These commands are slightly different from the Redis-native commands.

- String

```
//set get
string strKey = "hello";
string strValue = "world";
bool setResult = db.StringSet(strKey, strValue);
Console.WriteLine("set " + strKey + " " + strValue + ", result is " + setResult);
//incr
string counterKey = "counter";
long counterValue = db.StringIncrement(counterKey);
Console.WriteLine("incr " + counterKey + ", result is " + counterValue);
//expire
db.KeyExpire(strKey, new TimeSpan(0, 0, 5));
Thread.Sleep(5 * 1000);
Console.WriteLine("expire " + strKey + ", after 5 seconds, value is " + db.StringGet(strKey));
//mset mget
KeyValuePair<RedisKey, RedisValue> kv1 = new KeyValuePair<RedisKey, RedisValue>("key1", "value1");
KeyValuePair<RedisKey, RedisValue> kv2 = new KeyValuePair<RedisKey, RedisValue>("key2", "value2");
db.StringSet(new KeyValuePair<RedisKey, RedisValue>[] {kv1, kv2});
RedisValue[] values = db.StringGet(new RedisKey[] {kv1.Key, kv2.Key});
Console.WriteLine("mget " + kv1.Key.ToString() + " " + kv2.Key.ToString() + ", result is " + values[0] + "&&" + values[1]);
```

- o Hash

```
string hashKey = "myhash";
//hset
db.HashSet(hashKey, "f1", "v1");
db.HashSet(hashKey, "f2", "v2");
HashEntry[] values = db.HashGetAll(hashKey);
//hgetall
Console.WriteLine("hgetall " + hashKey + ", result is");
for (int i = 0; i < values.Length; i++)
{
    HashEntry hashEntry = values[i];
    Console.WriteLine(" " + hashEntry.Name.ToString() + " " + hashEntry.Value.ToString());
}
Console.WriteLine();
```

- o List

```
//list key
string listKey = "myList";
//rpush
db.ListRightPush(listKey, "a");
db.ListRightPush(listKey, "b");
db.ListRightPush(listKey, "c");
//lrange
RedisValue[] values = db.ListRange(listKey, 0, -1);
Console.WriteLine("lrange " + listKey + " 0 -1, result is ");
for (int i = 0; i < values.Length; i++)
{
    Console.WriteLine(values[i] + " ");
}
Console.WriteLine();
```

- o Set

```
//set key
string setKey = "mySet";
//sadd
db.SetAdd(setKey, "a");
db.SetAdd(setKey, "b");
db.SetAdd(setKey, "c");
//sismember
bool isContains = db.SetContains(setKey, "a");
Console.WriteLine("set " + setKey + " contains a is " + isContains );
```

- Sorted Set

```
string sortedSetKey = "myZset";
//sadd
db.SortedSetAdd(sortedSetKey, "xiaoming", 85);
db.SortedSetAdd(sortedSetKey, "xiaohong", 100);
db.SortedSetAdd(sortedSetKey, "xiaofei", 62);
db.SortedSetAdd(sortedSetKey, "xiaotang", 73);
//zrevrangebyscore
RedisValue[] names = db.SortedSetRangeByRank(sortedSetKey, 0, 2, Order.Ascending);
Console.WriteLine("zrevrangebyscore " + sortedSetKey + " 0 2, result is ");
for (int i = 0; i < names.Length; i++)
{
    Console.WriteLine(names[i] + " ");
}
Console.WriteLine();
```

13.2.5.2. Use redis-cli

The redis-cli tool is a command-line interface (CLI) of Redis. You can use redis-cli to connect to a KVStore for Redis instance from an Elastic Compute Service (ECS) instance or on-premises machine and manage data.

Prerequisites

- The ECS instance and KVStore for Redis instance are deployed in the same classic network or virtual private cloud (VPC).
- The private IP address of the ECS instance is added to the whitelist of the KVStore for Redis instance. For more information, see [Configure a whitelist](#).
- The ECS instance runs a Linux operating system and open source Redis is installed. For more information, see [Redis official website](#).

Procedure

- Log on to the CLI of the ECS instance and run the following command to connect to the Redis instance:

```
src/redis-cli -h <hostname> -p <port>
```

Parameters

| Parameter | Description |
|------------|--|
| <hostname> | The internal endpoint of the KVStore for Redis instance. You can find the endpoint in the Connection Information section on the Instance Information page. |
| <port> | The service port number of the KVStore for Redis instance. Default value: 6379. |

Example

```
src/redis-cli -h r-bp1zxszhcgatnx****.redis.rds.aliyuncs.com -p 6379
```

2. Run the following command to verify the password:

```
AUTH <password>
```

<password>: the password that you specify when you create the instance. If you forget your password, you can reset the password. For more information, see [Change the password](#).

Example:

```
AUTH testaccount:Rp829dlwa
```

If the password verification is successful, the following result is returned:

```
OK
```

13.3. Instance management

13.3.1. Change a password

If you forget your password, need to change your password, or have not set a password for an instance, you can set a new password for the instance.

Procedure

1. [Log on to the KVStore for Redis console](#).
2. On the **Instance List** page, click the ID of the instance.
3. In the upper-right corner of the **Basic Information** page, click **Modify Password**.
4. In the dialog box that appears, enter the current password and a new password.

Note

- If you forget your password, you can click **Forgot password** in the Change Password dialog box and enter a new password.
- The password must be 8 to 32 characters in length.
- The password must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. Special characters include `!@#$$%^&*()_+ = .`

5. Click **OK**.

13.3.2. Configure a whitelist

Before you use a KVStore for Redis instance, add IP addresses or CIDR blocks that are used to access the database to the whitelist of the instance to improve the security and stability of the database.

Context

 **Note** A properly configured whitelist can ensure a higher level of security protection for your KVStore for Redis instance. We recommend that you maintain the whitelist on a regular basis.

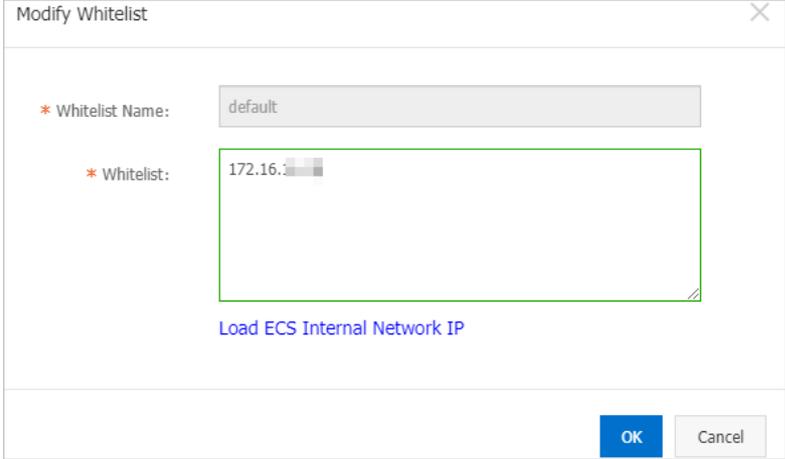
Procedure

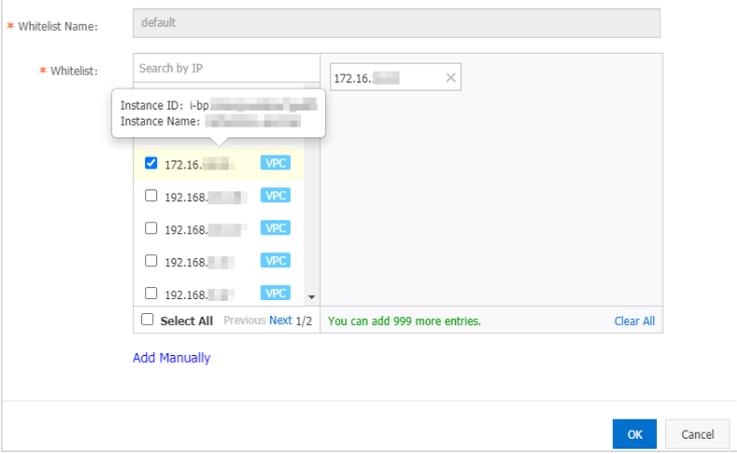
1. [Log on to the KVStore for Redis console](#).
2. On the **Instance List** page, click the ID of the instance.
3. In the left-side navigation pane, click **Whitelist Settings**.

4. Find the IP address whitelist that you want to manage and click **Modify**.

Note You can also click **Add Whitelist** to create an IP address whitelist. The name of the IP address whitelist must be 2 to 32 characters in length and can contain lowercase letters, digits, and underscores (_). The name of the whitelist must start with a lowercase letter and end with a lowercase letter or digit.

5. In the dialog box that appears, perform one of the following operations:

| Action | Procedure |
|---|---|
| <p>Manually modify the IP address whitelist</p> | <p>Enter IP addresses or CIDR blocks.</p> <p>Manually modify the IP address whitelist</p>  <p>Note</p> <ul style="list-style-type: none"> Separate multiple IP addresses with commas (,). You can add up to 1,000 unique IP addresses. Supported formats are specific IP addresses such as 10.23.12.24 and CIDR blocks such as 10.23.12.24/24. /24 indicates the length of the IP address prefix. An IP address prefix can be 1 to 32 bits in length. If you set the prefix length to 0, for example, 0.0.0.0/0 or 127.0.0.1/0, all IP addresses are allowed to access the instance. In this case, the security risk of your instance is high. Proceed with caution. |

| Action | Procedure |
|---|--|
| <p>Add private IP addresses of ECS instances to an IP address whitelist</p> | <p>i. Click Load ECS Internal Network IP.</p> <p>ii. Select IP addresses based on your business requirements.</p> <p>Add private IP addresses of ECS instances</p>  <p>Note To find the ECS instance to which a specific IP address is assigned, you can move the pointer over the IP address. Then, the system displays the ID and name of the ECS instance to which the IP address is assigned.</p> |
| <p>Remove IP addresses from the IP address whitelist</p> | <p>To remove all IP addresses from the IP address whitelist and retain the IP address whitelist, click Delete.</p> |

6. Then, click **OK**.

13.3.3. Change specifications

This topic describes how to change the configuration of a KVStore for Redis instance.

Precautions

After configuration changes, the system migrates data and switches the instance. The instance is disconnected for a few seconds during this process. We recommend that you upgrade or downgrade the instance during off-peak hours.

Procedure

1. [Log on to the KVStore for Redis console](#).
2. On the **Instance List** page, click the ID of the instance.
3. On the page that appears, configure the required parameters.

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|-------------------|---|
| Architecture Type | Select the architecture type of the KVStore for Redis instance. KVStore for Redis supports cluster and standard architectures. The cluster architecture is suitable for scenarios that require large capacity or high performance. Redis-native databases run in a single-threading model. If your database does not require high performance, we recommend that you use a standard instance. If your database requires high performance, select a cluster architecture. |
| Instance Type | Select the specification of the instance. The maximum connections and maximum internal network bandwidth vary based on the instance type. |

4. Click **Submit**.

13.3.4. Specify a maintenance window

You can modify the default maintenance window to perform maintenance on KVStore for Redis during off-peak hours.

Context

To ensure the stability of KVStore for Redis instances on the Alibaba Cloud platform, the backend system performs maintenance on instances and servers occasionally.

To guarantee the stability of the maintenance process, instances will enter the Maintaining Instance status before the preset maintenance window on the day of maintenance. While an instance is in this state, data in the database can still be accessed and query operations such as performance monitoring are still available. However, change operations such as configuration change are temporarily unavailable for this instance in the console.

 **Note** During the maintenance process, instances may be disconnected in the process of maintenance. We recommend that you set the maintenance window to a period during off-peak hours.

Procedure

1. [Log on to the KVStore for Redis console](#).
2. On the **Instance List** page, click the ID of the instance.
3. In the **Basic Information** section, click **Settings** on the right of **Maintenance Window**.
4. Select a time period for maintenance and click **Save**.

 **Note** The time periods are in UTC+8.

13.3.5. Upgrade the minor version

Alibaba Cloud has continuously optimized the kernel of KVStore for Redis to fix security vulnerabilities and provide more stable services. You can upgrade the kernel version (minor version) of a KVStore for Redis instance with one click in the console.

Context

 **Note**

- We recommend that you upgrade instance versions during off-peak hours and ensure that your application supports automatic reconnection.
- The system automatically checks the kernel version of an instance. If the current version is the latest, the **Minor Version Upgrade** button in the upper-right corner of the **Basic Information** section for this instance will appear dimmed.

Procedure

1. [Log on to the KVStore for Redis console](#).
2. On the **Instance List** page, click the ID of the instance.
3. On the **Instance Information** page, click **Minor Version Upgrade** in the upper-right corner of the **Basic Information** section.
4. In the **Minor Version Upgrade** dialog box that appears, click **Upgrade Now**.

On the **Instance Information** page, the instance status will become **Upgrading a minor version**. When the instance status returns to **Available**, the upgrade has been completed.

13.3.6. Configure SSL encryption

The standard and cluster instances of Redis 2.8 and the cluster instances of Redis 4.0 support secure sockets layer (SSL) encryption. You can enable SSL encryption to ensure more secure data transmission.

Context

 **Note** SSL encryption may increase the network response time of instances. We recommend that you enable this feature only when necessary.

Procedure

1. [Log on to the KVStore for Redis console](#).
2. On the **Instance List** page, click the ID of the instance.
3. In the left-side navigation pane, click **SSL Settings**.
4. In the upper-right corner of the **SSL Settings** page, click **Configure SSL**.
5. In the dialog box that appears, turn on **Enable SSL Certificate**.

 **Note** If the version of the instance is not supported, upgrade the minor version. For more information, see [Upgrade the minor version](#).

6. Click **OK**.

 **Note**

- The result of the operation is displayed after a short period of time.
- In the upper-right corner of the **SSL Settings** page, you can also click **Update Validity** and **Download CA Certificate** to perform relevant operations.

13.3.7. Clear data

You can clear the data of a KVStore for Redis instance in the console.

Context

 **Warning** This operation will delete all data contained on an instance. Deleted data cannot be restored. Proceed with caution.

Procedure

1. [Log on to the KVStore for Redis console.](#)
2. On the **Instance List** page, click the ID of the instance.
3. In the upper-right corner of the **Instance Information** page, click **Clear Data**.
4. In the **Clear Data** message that appears, click **OK**.

13.3.8. Release an instance

You can release a KVStore for Redis instance at any time based on your business needs. This topic describes how to release a KVStore for Redis instance.

Procedure

1. [Log on to the KVStore for Redis console.](#)
2. On the **Instances** page, click the instance ID or choose  > **Release** in the **Actions** column.

 **Warning** After an instance is released, the instance cannot be restored. Proceed with caution. We recommend that you back up your data before you release an instance.

3. In the **Release Instance** message that appears, click **OK**.

13.3.9. Manage database accounts

KVStore for Redis allows you to create up to 20 database accounts for an instance. You can grant permissions to these accounts and manage your instance to prevent user errors.

Prerequisites

The engine version of the instance is Redis 4.0 or later.

 **Note** If the engine version of the instance is not Redis 4.0, only the default account that is created when you create the instance is available. For more information about how to change the password of the default account, see [Change the password](#).

Context

You can create accounts, delete accounts, reset the password, and change the permissions. After an account is created, you can use this account to log on to the database and use the command-line tool to perform operations on the database with the account and granted permissions.

Create an account

1. [Log on to the KVStore for Redis console.](#)
2. On the **Instance List** page, click the ID of the instance.
3. On the **Instances** page, find the instance that you want to manage and click the instance ID. In the left-side navigation pane, click **Account Management**.

 **Note** If Account Management is unavailable for an instance of Redis 4.0 or later, you must upgrade the minor version of the instance. For more information, see [Upgrade the minor version](#).

- In the upper-right corner of the **Account Management** page, click **Create**.
- In the dialog box that appears, configure the required parameters and click **OK**. The following table describes the parameters.

| Parameter | Description |
|--------------------------|---|
| Account | The account name must be 1 to 16 characters in length. The name can contain lowercase letters, digits, and underscores (_) and must start with a letter. |
| Privilege | Specify the permissions that are granted to the account. Valid values: Read-only, Read/Write, and Replicate. If you select Replicate, you can run the SYNC and PSYNC commands after you connect to an instance by using your account.  Note You can create accounts that have the replicate permissions only for standard instances of Redis 4.0 or later. |
| Password Settings | Specify a password for the account. The password must be 8 to 32 characters in length. The password must contain at least three of the following types of characters: uppercase letters, lowercase letters, digits, and special characters. The following special characters are supported: !@#\$%^&*()+-=_. |
| Confirm Password | Enter the password again. |
| Description | The description of the account. |

13.3.10. Restart an instance

You can restart an instance from the Instance List page of the console.

Procedure

- Log on to the [KVStore for Redis console](#).
- In the **Actions** column, choose  > **Restart**.

 **Warning** During the restart, the instance may be disconnected for a few seconds. We recommend that you restart instances during off-peak hours. You must also make sure that your application supports automatic reconnection.

- In the dialog box that appears, specify the upgrade time and click **OK**.
 - Restart Immediately: restarts the instance immediately.
 - Restart Within Maintenance Window: restarts the instance within the preset [maintenance window](#).

13.3.11. Export the list of instances

You can export the list of KVStore for Redis instances from the KVStore for Redis console for offline management.

Procedure

- Log on to the [KVStore for Redis console](#).

2. In the upper-right corner of the **Instance List** page, click the **Export Instances** icon.
3. In the **Export Instance List** dialog box that appears, select the columns to export and click **OK**.

Note After you click **OK**, the browser begins to download the CSV file. You can use Excel or a text editor to view this file.

13.3.12. Use a Lua script

KVStore for Redis instances of all editions support Lua commands.

Support for Lua commands

Lua scripts improve the performance of KVStore for Redis. With support for the Lua environment, KVStore for Redis is able to perform check-and-set (CAS) operations, allowing you to combine and run multiple commands in an efficient manner.

Note If the EVAL command cannot be executed, for example, the "ERR command eval not support for normal user" message appears, you can [Upgrade the minor version](#). During the upgrade, the instance may be disconnected and become read-only for a few seconds. We recommend that you upgrade instance versions during off-peak hours.

Limits on Lua scripts

A Lua script, which is supported by the cluster instance of KVStore for Redis, has the following limits to ensure that all operations in the script are performed within the same hash slot:

- The Lua script uses the `redis.call/redis.pcall` function to run Redis commands. For Redis commands, the keys must be passed by using the KEYS array, which cannot be replaced by Lua variables. If the KEYS array is not used, the following error message is returned:

```
-ERR bad lua script for redis cluster, all the keys that the script uses should be passed using the KEYS array\r\n
```

- All keys that are used by the script must be allocated to the same hash slot. If the keys are allocated to different hash slots, the following error message is returned:

```
-ERR eval/evalsha command keys must be in same slot\r\n
```

- Keys must be included in all commands that you want to run. If the keys are not included in all commands, the following error message is returned:

```
-ERR for redis cluster, eval/evalsha number of keys can't be negative or zero\r\n
```

- The following Pub/Sub commands are not supported: `PSUBSCRIBE`, `PUBSUB`, `PUBLISH`, `PUNSUBSCRIBE`, `SUBSCRIBE`, and `UNSUBSCRIBE`.
- The `UNPACK` function is not supported.

Note If you do not want to impose the preceding limits but can make sure that all operations are performed in the same hash slot in the code, you can set the `script_check_enable` parameter to 0 in the console to disable the backend script check.

13.4. Connection management

13.4.1. View endpoints

You can view the internal and public endpoints of instances in the KVStore for Redis console.

Context

Note

- The virtual IP address of a KVStore for Redis instance may change when you maintain or modify the instance. To ensure that the connection is available, we recommend that you use an endpoint to access the KVStore for Redis instance.
- For more information about how to apply for a public endpoint, see [Applies for a public connection string](#).

Procedure

1. [Log on to the KVStore for Redis console](#).
2. On the **Instance List** page, click the ID of the instance.
3. In the **Connection Information** section, you can view the private and public endpoints of the instance.

 **Note** By default, only a private endpoint is provided by a KVStore for Redis instance. If you want to connect to a KVStore for Redis instance over the Internet, you must apply for a public endpoint. For more information, see [Apply for a public endpoint](#).

13.4.2. Apply for a public endpoint

This topic describes how to apply for a public endpoint for a KVStore for Redis instance.

Procedure

1. [Log on to the KVStore for Redis console](#).
2. On the **Instance List** page, click the ID of the instance.
3. In the **Connection Information** section, click **Apply for Public Endpoint**.
4. In the dialog box that appears, enter an endpoint and a port number and click **OK**.

Note

- The custom endpoint must be 8 to 64 characters in length and can contain lowercase letters and digits. The endpoint must start with a lowercase letter.
- The custom port number ranges from 1024 to 65535. The default value is 6379.
- After you apply for a public endpoint, you must add the public endpoint to the IP address whitelist of the instance. This way, you can connect to the instance over the Internet. For more information, see [Configure a whitelist](#).

5. On the **Instance Information** page, view the **Public Endpoint** in the **Connection Information** section.

 **Note** If you no longer use the public endpoint, click **Release Public Endpoint** next to **Public Endpoint** to release the endpoint.

13.4.3. Modify the endpoint of an KVStore for Redis instance

KVStore for Redis allows you to modify internal and public endpoints for instances. When changing the KVStore for Redis instance, you can change the endpoint of the new instance to the endpoint of the original instance without modifying the application.

Prerequisites

The instance is in the Running state.

Procedure

1. [Log on to the KVStore for Redis console.](#)
2. On the **Instance List** page, click the ID of the instance.
3. On the **Instance Information** page, click **Modify Public Endpoint** in the **Connection Information** section.
4. In the **Modify Public Endpoint** dialog box, set the following parameters:

| Parameter | Description |
|-----------------|--|
| Connection type | Select Internal Endpoint or Public Endpoint . |
| Endpoint | Set the prefix of the endpoint. <ul style="list-style-type: none"> ◦ The endpoint can contain lowercase letters and digits. ◦ It must start with a lowercase letter. ◦ The endpoint must be 8 to 64 characters in length. |
| Port | Specify a port number. Valid value: 1024 to 65535. <div style="background-color: #e1f5fe; padding: 5px; margin-top: 10px;"> <p> Note It takes about 10 minutes for the modified port number of the public endpoint to take effect. You can refresh the page to view the latest port number information.</p> </div> |

5. In the **Modify Public Endpoint** dialog box, modify **Connection Type**, **Endpoint**, and **Port**, and then click **OK**.

Note

- The custom endpoint prefix must be 8 to 64 characters in length and can contain lowercase letters and digits. It must start with a lowercase letter.
- The custom port number ranges from 1024 to 65535. The default value is 6379.

13.5. Performance monitoring

13.5.1. Query monitoring data

You can query the monitoring data of a KVStore for Redis instance for a specified period within the last month.

Procedure

1. [Log on to the KVStore for Redis console.](#)
2. On the **Instance List** page, click the ID of the instance.
3. In the left-side navigation pane, click **Performance Monitor**.
4. Select the start and end time and click **OK**.

 **Note** For more information about the metrics, see [Understand metrics](#).

13.5.2. Select metrics

You can select the metrics to be displayed on the Historical Monitoring Data page of the KVStore for Redis console as needed.

Context

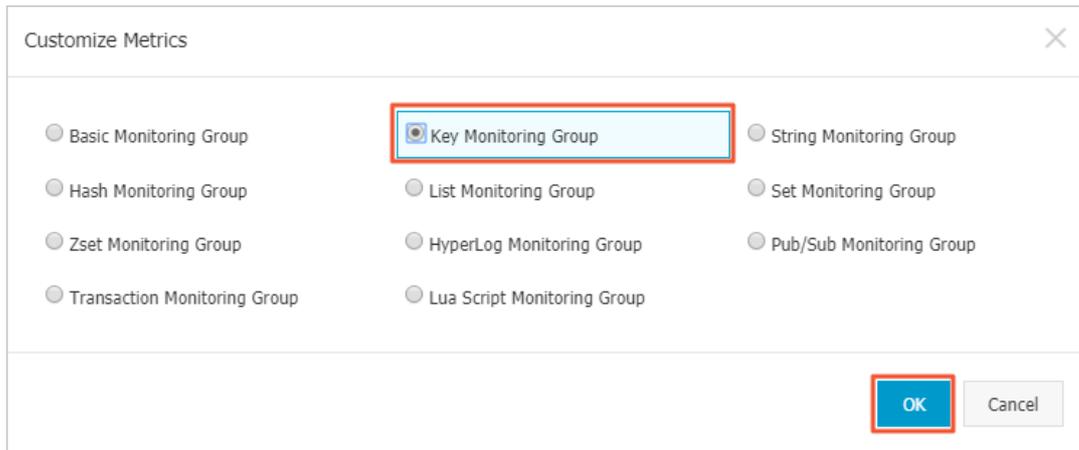
KVStore for Redis supports more than 10 monitoring groups. By default, the Performance Monitor page displays the metrics of the basic monitoring group. You can click **Customize Metrics** to switch to the metrics of other monitoring groups. The following table describes the monitoring groups.

| Monitoring group | Description |
|------------------------------|--|
| Basic monitoring group | The basic monitoring information about an instance, such as the queries per second (QPS), bandwidth, and memory usage of the instance. |
| Key monitoring group | The monitoring information about the use of key-value related commands, such as the number of times that the DEL and EXISTS commands are executed. |
| String monitoring group | The monitoring information about the use of string commands, such as the number of times that the APPEND and MGET commands are executed. |
| Hash monitoring group | The monitoring information about the use of hash commands, such as the number of times that the HGET and HDEL commands are executed. |
| List monitoring group | The monitoring information about the use of list commands, such as the number of times that the BLPOP and BRPOP commands are executed. |
| Set monitoring group | The monitoring information about the use of set commands, such as the number of times that the SADD and SCARD commands are executed. |
| Zset monitoring group | The monitoring information about the use of zset commands, such as the number of times that the ZADD and ZCARD commands are executed. |
| HyperLog monitoring group | The monitoring information about the use of HyperLogLog commands, such as the number of times that the PFADD and PFCOUNT commands are executed. |
| Pub/Sub monitoring group | The monitoring information about the use of publication and subscription commands, such as the number of times that the PUBLISH and SUBSCRIBE commands are executed. |
| Transaction monitoring group | The monitoring information about the use of transaction commands, such as the number of times that the WATCH, MULTI, and EXEC commands are executed. |
| Lua script monitoring group | The monitoring information about the use of Lua script commands, such as the number of times that the EVAL and SCRIPT commands are executed. |

For more information about the definitions of the metrics in each monitoring group, see [Metrics](#).

Procedure

1. [Log on to the KVStore for Redis console](#).
2. On the **Instance List** page, click the ID of the instance.
3. In the left-side navigation pane, click **Performance Monitor**.
4. On the **Historical Monitoring Data** page, click **Customize Metrics** in the **Data Index** section.
5. In the dialog box that appears, specify a monitoring group and click **OK**.



On the Historical Monitoring Data page, the metrics in the selected monitoring group appear.

13.5.3. Modify the data collection interval

KVStore for Redis console allows you to set the frequency at which monitoring data is collected.

Context

You can set the monitoring frequency to either 5 or 60 seconds to specify how often monitoring data to be collected by KVStore for Redis. The default monitoring time of 60 seconds is sufficient to meet common monitoring requirements. If you need to observe certain metrics at a higher frequency and lower latency, you can change the monitoring frequency to 5 seconds as described in the following section. Monitoring data does not occupy instance storage space, and collection of monitoring data does not affect normal running of the instance.

Procedure

1. [Log on to the KVStore for Redis console.](#)
2. On the **Instance List** page, click the ID of the instance.
3. In the left-side navigation pane, click **Performance Monitor**.
4. In the upper-right corner of the **Historical Monitoring Data** page, click **Monitoring Frequency**.
5. In the **Monitoring Frequency** dialog box that appears, select the new monitoring frequency and click **OK**.

13.5.4. Understand metrics

KVStore for Redis updates more than 10 monitoring groups of metrics in real time. This allows you to monitor the status of KVStore for Redis instances. This topic describes the metrics of the monitoring groups.

Metrics of basic monitoring groups

| Metric | Unit | Description | Statistical method |
|------------|----------|--|---|
| CpuUsage | % | The CPU utilization | Check the CPU utilization when the monitoring data is collected. |
| UsedMemory | Bytes | The amount of the used memory. | Check the memory usage when the monitoring data is collected. |
| TotalQps | Counts/s | The number of requests that are received by the instance per second. | Divide the number of requests in a monitoring cycle by the number of seconds in the monitoring cycle. |

| Metric | Unit | Description | Statistical method |
|-------------|----------|--|--|
| ConnCount | Counts | The number of connections. | Check the number of connections when collecting monitoring data. |
| InFlow | KBps | The amount of data received by the instance per second. | Divide the amount of data that is received in a monitoring cycle by the number of seconds in the monitoring cycle. |
| OutFlow | KBps | The amount of data sent by the instance per second. | Divide the amount of data that is sent in a monitoring cycle by the number of seconds in the monitoring cycle. |
| FailedCount | Counts/s | The average number of abnormal requests per second. | Divide the total number of abnormal requests in a monitoring cycle by the number of seconds in the monitoring cycle. |
| AvgRt | us | The average response time of all requests.  Note For more information, see Response time (RT) metrics . | Divide the processing time of all requests in a monitoring cycle by the number of requests in the monitoring cycle. |
| MaxRt | us | The maximum response time of requests.  Note For more information, see Response time (RT) metrics . | The maximum time that is required to process a request in a monitoring cycle. |
| Keys | Counts | The total number of keys. | The number of keys when the monitoring data is collected. |
| Expires | Counts | The total number of keys for which an expiration time is configured. | The total number of keys for which an expiration time is set when the monitoring data is collected. |
| ExpiredKeys | Counts | The total number of expired keys. | The cumulative sum when the monitoring data is collected. After the instance is restarted, the cumulative sum is calculated again. |
| EvictedKeys | Counts | The total number of keys that are evicted because the memory is exhausted. | The cumulative sum when the monitoring data is collected. After the instance is restarted, the cumulative sum is calculated again. |
| request | Bytes | The total amount of request data received by KVStore for Redis nodes in a monitoring cycle. | See the description of this metric. |
| response | Bytes | The total amount of response data sent by KVStore for Redis nodes in a monitoring cycle. | See the description of this metric. |

| Metric | Unit | Description | Statistical method |
|-------------------------------|----------|--|--|
| request_max | Bytes | The maximum amount of data in a request in a monitoring cycle. | See the description of this metric. |
| response_max | Bytes | The maximum amount of data in a response in a monitoring cycle. | See the description of this metric. |
| traffic_control_input | Counts | The number of times that downstream throttling is triggered. | Monitor the cumulative sum in a monitoring cycle. |
| traffic_control_output | Counts | The number of times that uplink throttling is triggered. | Monitor the cumulative sum in a monitoring cycle. |
| traffic_control_input_status | Counts | Indicates whether downstream throttling was triggered in a monitoring cycle. A value of 0 indicates that throttling was not triggered. A value of 1 indicates that throttling was triggered. | See the description of this metric. |
| traffic_control_output_status | Counts | Indicates whether upstream throttling was triggered in a monitoring cycle. A value of 0 indicates that throttling was not triggered. A value of 1 indicates that throttling was triggered. | See the description of this metric. |
| hit_rate | % | The request hit rate, which is the probability that data exists in a KVStore for Redis instance for a data access request. | Calculate the percentage of the hit requests to the total number of requests in a monitoring cycle. |
| hit | Counts | The number of hit requests. | Check the number of hit requests in a monitoring cycle. |
| miss | Counts | The number of missed requests. | Check the number of missed requests in a monitoring cycle. |
| evicted_keys_per_sec | Counts/s | The number of keys that are evicted per second. | Divide the total number of keys that are evicted in a monitoring cycle by the number of seconds in the monitoring cycle. |

Metrics in other monitoring groups

The system also uses other metrics to monitor specific types of data or specific features. The metrics are classified into:

- Metrics that indicate the number of times that commands are used. For example, the del, dump, and exists metrics that are used to monitor keys indicate the number of times that the DEL, DUMP, and EXISTS commands are executed.
- **Response time (RT) metrics** of commands. For example, the metrics that end with avg_rt, such as del_avg_rt, dump_avg_rt, and exists_avg_rt, in the key monitoring group are used to monitor the average response time of the DEL, DUMP, and EXISTS commands in a monitoring cycle.

Response time (RT) metrics

All monitoring groups have RT metrics. RT metrics end with Rt or rt. For example, the AvgRt and MaxRt metrics are used in the basic monitoring group and the del_avg_rt and exists_avg_rt metrics are used to monitor keys.

The AvgRt and MaxRt metrics in the basic monitoring group are the most frequently used RT metrics. These metrics have different meanings for proxy nodes and data nodes.

- For a cluster instance or a read/write splitting instance, the AvgRt metric of a proxy node indicates the average time consumed by the proxy node to process all commands. The following process shows how a proxy node processes a command:
 - i. The proxy node receives a command and forwards the command to a data node.
 - ii. The data node processes the command and responds to the proxy node.
 - iii. The proxy node returns the command processing result.

The AvgRt metric of the proxy node includes the amount of time consumed by the data node to process a command and the time that is required to wait for the command to be processed. This metric also includes the amount of time consumed for network communication between the proxy node and the data node.

- For data nodes of a cluster instance or a read/write splitting instance or for a standard instance, the AvgRt metric indicates the average time consumed by a data node to process all commands. This metric records the period of time from the time when the data node receives the command to the time when the data node returns the result. This metric does not include the time consumed by the proxy node to process a command and the time that is required for network communication.
- The MaxRt metric indicates the maximum response time of requests. The statistical method of this metric is similar to the statistical method of the AvgRt metric for all KVStore for Redis instances.

13.6. Parameter configuration

KVStore for Redis allows you to customize certain instance parameters. This topic describes parameters and the common methods to modify them in the KVStore for Redis console.

Context

KVStore for Redis is completely compatible with the native database services of Redis. The method to set parameters for KVStore for Redis is similar to that of an on-premises Redis database. You can set the parameters described in this topic in the KVStore for Redis console.

Parameters

Parameters

| Parameter | Description |
|----------------------------------|--|
| #no_loose_check-whitelist-always | Specifies whether to check whether the client IP address is in the whitelist of the KVStore for Redis instance after password-free access is enabled in Virtual Private Cloud (VPC). Default value: no. If you set this parameter to yes, the whitelist will still take effect in password-free access mode for VPC. Valid values: <ul style="list-style-type: none"> • yes • no |
| #no_loose_disabled-commands | Specifies the disabled commands. Separate multiple commands with commas (,). You can disable the following commands: FLUSHALL , FLUSHDB , KEYS , HGETALL , EVAL , EVALSHA , and SCRIPT . |
| #no_loose_ssl-enabled | Specifies whether to enable SSL encryption. Default value: no. Valid values: <ul style="list-style-type: none"> • yes • no |

| Parameter | Description |
|-----------------------------------|--|
| #no_loose_sentinel-enabled | <p>Specifies whether to enable Sentinel-compatible mode. Default value: no. Valid values:</p> <ul style="list-style-type: none"> • yes • no |
| client-output-buffer-limit pubsub | <p>Limits the size of output buffers for Pub/Sub clients. This parameter can contain options in the following format: <code><hard limit> <soft limit> <soft seconds></code> .</p> <ul style="list-style-type: none"> • Hard limit: If the output buffer of a Pub/Sub client reaches or exceeds the number of bytes specified by hard limit, the client is immediately disconnected. • soft limit and soft seconds: If the output buffer of a Pub/Sub client reaches or exceeds the size in bytes specified by soft limit for a period of time in seconds specified by soft seconds, the client will be disconnected. |
| dynamic-hz | <p>Specifies whether to enable dynamic frequency control for background tasks. Default value: yes. Valid values:</p> <ul style="list-style-type: none"> • yes • no |
| hash-max-ziplist-entries | <p>Specifies the maximum size of each key-value pair stored within a hash in bytes. A hash is encoded using ziplist when it meets the following conditions:</p> <ol style="list-style-type: none"> 1. The maximum size of each key-value pair stored within the hash in bytes must be less than the value of the hash-max-ziplist-value parameter. 2. The number of key-value pairs stored within the hash must be less than the value of the hash-max-ziplist-entries parameter. |
| hash-max-ziplist-value | <p>Specifies the maximum size of each key-value pair stored within a hash in bytes. A hash is encoded using ziplist when it meets the following conditions:</p> <ol style="list-style-type: none"> 1. The maximum size of each key-value pair stored within the hash in bytes must be less than the value of the hash-max-ziplist-value parameter. 2. The number of key-value pairs stored within the hash must be less than the value of the hash-max-ziplist-entries parameter. |
| hz | <p>Specifies the execution frequency for background tasks, such as tasks to evict expired keys. Valid values: 1 to 500. Default value: 10. The larger the value of the hz parameter, the more frequently background tasks are performed and the more precisely timeout events are handled, but the more CPU KVStore for Redis consumes. We recommend that you do not set the hz parameter to a value greater than 100.</p> |
| lazyfree-lazy-eviction | <p>Specifies whether to enable lazyfree for the eviction feature. Default value: no. Valid values:</p> <ul style="list-style-type: none"> • yes • no |

| Parameter | Description |
|--------------------------|--|
| lazyfree-lazy-expire | <p>Specifies whether to enable lazyfree to delete expired keys. Default value: yes. Valid values:</p> <ul style="list-style-type: none"> • yes • no |
| lazyfree-lazy-server-del | <p>Specifies whether to enable lazyfree to asynchronously delete data with the DEL command. Default value: yes. Valid values:</p> <ul style="list-style-type: none"> • yes • no |
| list-compress-depth | <p>Specifies the number of nodes that are not compressed at each side in a list. Default value: 0. Valid values:</p> <ul style="list-style-type: none"> • 0: does not compress any nodes in the list. • 1: does not compress the first node from each side of the list, but compresses all nodes in between. • 2: does not compress the first two nodes from each side of the list, but compresses all nodes in between. • 3: does not compress the first three nodes from each side of the list, but compresses all nodes in between. • And so on up to 65535. |
| list-max-ziplist-size | <ul style="list-style-type: none"> • Specifies the maximum size of each ziplist in a quicklist. A positive number indicates the maximum number of elements in each ziplist of a quicklist. For example, if you set this parameter to 5, each ziplist of a quicklist can contain a maximum of five elements. • A negative number indicates the maximum number of bytes in each ziplist of a quicklist. Default value: -2. Valid values: <ul style="list-style-type: none"> ◦ -5: indicates that each ziplist of a quicklist cannot exceed 64 KB (1 KB = 1,024 bytes). ◦ -4: indicates that each ziplist of a quicklist cannot exceed 32 KB. ◦ -3: indicates that each ziplist of a quicklist cannot exceed 16 KB. ◦ -2: indicates that each ziplist of a quicklist cannot exceed 8 KB. ◦ -1: indicates that each ziplist of a quicklist cannot exceed 4 KB. |
| maxmemory-policy | <p>Specifies the policy used to evict keys if the memory is fully occupied. Valid values: LRU means least recently used. LFU means least frequently used. LRU, LFU, and TTL are implemented by using approximated randomized algorithms.</p> <ul style="list-style-type: none"> • volatile-lru: evicts the approximated least recently used (LRU) keys among keys with a preset expiration time. • allkeys-lru: evicts the approximated LRU keys. • volatile-lfu: evicts the approximated least frequently used (LFU) keys among keys with a preset expiration time. • allkeys-lfu: evicts the approximated LFU keys. • volatile-random: evicts random keys among keys with a preset expiration time. • allkeys-random: evicts random keys. • volatile-ttl: evicts keys with the nearest time to live (TTL) among keys with a preset expiration time. • noeviction: does not evict any keys, but returns an error on write operations. |

| Parameter | Description |
|-------------------------|---|
| notify-keyspace-events | <p>Specifies the events that the Redis server can notify clients of. The value of this parameter is any combination of the following characters, each of which specifies a type of event to be notified:</p> <ul style="list-style-type: none"> • K: keyspace events, published with the <code>__keyspace@<db>__</code> prefix. • E: keyevent events, published with the <code>__keyevent@<db>__</code> prefix. • g: generic commands that are non-type specific, such as DEL, EXPIRE, and RENAME. • l: list commands. • s: set commands. • h: hash commands. • z: sorted set commands. • x: expired key events. An expired key event is generated when a key expires. • e: evicted key events. An evicted key event is generated when a key is evicted due to the policy specified by the <code>maxmemory-policy</code> parameter. • A: the alias for <code>gshzxe</code>. |
| set-max-intset-entries | <p>Specifies the maximum number of data entries in a set. A set is encoded by using <code>intset</code> when it meets the following conditions:</p> <ol style="list-style-type: none"> 1. The set is composed of just strings. The number of strings is less than the value of this parameter. 2. All strings are integers in radix 10 in the range of 64-bit signed integers. |
| slowlog-log-slower-than | <p>Specifies whether to log slow queries.</p> <ul style="list-style-type: none"> • Negative number: does not log slow queries. • 0: logs all queries. • Positive number: logs queries that exceed an execution time specified by this positive number, in microseconds. <p>Valid values: 0 to 10,000,000. Default value: 10,000.</p> |
| slowlog-max-len | <p>Specifies the maximum number of slow query log entries that can be stored.</p> <p>Valid values: 100 to 10,000. Default value: 1,024.</p> |
| stream-node-max-bytes | <p>Specifies the maximum memory that can be used by each macro node in streams. Valid values: 0 to 999,999,999,999. If you set the parameter to 0, each macro node can use an unlimited amount of memory.</p> |
| stream-node-max-entries | <p>Specifies the maximum number of stream entries that can be stored within each macro node. Valid values: 0 to 999,999,999,999. If you set the parameter to 0, each macro node can store unlimited stream entries.</p> |
| timeout | <p>Specifies a timeout period for client connections. Unit: seconds. Valid values: 0 to 100,000. 0 indicates that client connections never time out.</p> |

| Parameter | Description |
|--------------------------|---|
| zset-max-ziplist-entries | <p>Specifies the maximum size of each key-value pair stored within a sorted set in bytes. A sorted set is encoded using ziplist when it meets the following conditions:</p> <ol style="list-style-type: none"> 1. The maximum size of each key-value pair stored within the sorted set in bytes must be less than the value of the zset-max-ziplist-value parameter. 2. The number of key-value pairs stored within the sorted set must be less than the value of the zset-max-ziplist-entries parameter. |
| zset-max-ziplist-value | <p>Specifies the maximum size of each key-value pair stored within a sorted set in bytes. A sorted set is encoded using ziplist when it meets the following conditions:</p> <ol style="list-style-type: none"> 1. The maximum size of each key-value pair stored within the sorted set in bytes must be less than the value of the zset-max-ziplist-value parameter. 2. The number of key-value pairs stored within the sorted set must be less than the value of the zset-max-ziplist-entries parameter. |
| list-max-ziplist-entries | <p>Specifies the maximum size of each key-value pair stored within a list in bytes. A list is encoded using ziplist when it meets the following conditions:</p> <ol style="list-style-type: none"> 1. The maximum size of each key-value pair stored within the list in bytes must be less than the value of the list-max-ziplist-value parameter. 2. The number of elements stored within the list is less than the value of the list-max-ziplist-entries parameter. |
| list-max-ziplist-value | <p>Specifies the maximum size of each key-value pair stored within a list in bytes. A list is encoded using ziplist when it meets the following conditions:</p> <ol style="list-style-type: none"> 1. The maximum size of each key-value pair stored within the list in bytes must be less than the value of the list-max-ziplist-value parameter. 2. The number of elements stored within the list is less than the value of the list-max-ziplist-entries parameter. |
| cluster_compat_enable | <p>Specifies whether to enable compatibility with the syntax of Redis Cluster. Default value: 1. Valid values:</p> <ul style="list-style-type: none"> • 0: no • 1: yes |
| script_check_enable | <p>Specifies whether to confirm that all the keys used in a Lua script are in the same hash slot. Default value: 1. Valid values:</p> <ul style="list-style-type: none"> • 0: no • 1: yes |

 **Note** The maxclients parameter, which is used to specify the maximum number of connections to Redis data nodes, is fixed to 10,000. You cannot modify the value of this parameter.

Configure parameters in the KVStore for Redis console

1. [Log on to the KVStore for Redis console.](#)
2. On the **Instance List** page, click the ID of the instance.
3. In the left-side navigation pane of the **Instance Information** page, click **System Parameters**.
4. Find the target parameter and click **Modify** in the **Action** column.
5. In the dialog box that appears, modify the parameter value and click **OK**.

13.7. Backup and recovery

13.7.1. Automatically back up data

An increasing number of applications use Redis for persistent storage. In this case, an automatic backup mechanism is required to back up data on a regular basis so that you can restore data if user errors occur. KVStore for Redis uses Redis database backup (RDB) snapshots to back up data on replica nodes. The backup process does not have negative impacts on the performance of your instance. You can configure a custom backup policy in the console.

Procedure

1. [Log on to the KVStore for Redis console.](#)
2. On the **Instance List** page, click the ID of the instance.
3. In the left-side navigation pane, click **Backup and Recovery**.
4. Click the **Backup Settings** tab.
5. Click **Edit** and specify Backup Cycle and Backup Time.
 - **Retention Days:** The number of days for which backups are retained. This parameter is set to seven days and cannot be changed.
 - **Backup Cycle:** You can select one or more days in a week. By default, one backup is created per day.
 - **Backup Time:** You can specify a period of time in hours within a day. We recommend that you back up data during off-peak hours.
6. Click **OK**.

13.7.2. Back up an instance

You can initiate a manual backup task in the console at any time.

Procedure

1. [Log on to the KVStore for Redis console.](#)
2. On the **Instance List** page, click the ID of the instance.
3. In the left-side navigation pane, click **Backup and Recovery**.
4. In the upper-right corner of the page, click **Create Backup**.
5. Click **OK**.

 **Note** On the **Data Backup** tab, you can select a time range to query existing backups. Backups are retained for seven days.

13.7.3. Download backup files

To archive backup files for a long period, you can copy the URLs in the console and download the database backup files to an on-premises machine.

Procedure

1. [Log on to the KVStore for Redis console.](#)
2. On the **Instance List** page, click the ID of the instance.
3. In the left-side navigation pane, click **Backup and Recovery**.
4. Find the backup that you want to download and click in the **Actions** column.
5. In the dialog box that appears, perform one of the actions that are described in the following table.

| Action | Procedure |
|--|--|
| Download the backup file over the Internet | <ol style="list-style-type: none"> i. Click Get URL for Internet. ii. Paste the URL into the address bar of a browser and press Enter. The browser automatically downloads the backup file. |
| Download the backup file over an internal network. For example, you can download the backup file on an Elastic Compute Service (ECS) instance. | <ol style="list-style-type: none"> i. Click Get URL for Intranet. ii. Select a downloading method based on the operating system of the client. <ul style="list-style-type: none"> ▪ Windows: Paste the URL into the address bar of a browser and press Enter. The browser automatically downloads the backup file. ▪ Linux: Run the following command in the specified format: <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <pre>wget -c '<URL that is used to download the backup file over the internal network>' -O <Name of the downloaded backup file>.<The extension that is added to the name of the downloaded backup file></pre> </div> <p>Example:</p> <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> <pre>wget -c 'http://rds****.oss-cn-hangzhou-internal.aliyuncs.com/custins416****/hins1****.rdb?...!' -O backupfile.rdb</pre> </div> <div style="background-color: #e0f2f7; padding: 5px; margin-top: 10px;"> <p> Note If you download the backup file on an ECS instance, the ECS instance and the KVStore for Redis instance can be deployed in different types of networks.</p> </div> |

13.7.4. Restore data

KVStore for Redis allows you to restore data from a specified backup set to the current KVStore for Redis instance.

Procedure

1. [Log on to the KVStore for Redis console.](#)
2. On the **Instance List** page, click the ID of the instance.
3. In the left-side navigation pane, click **Backup and Recovery**.
4. Perform one of the following operations based on the architecture of your KVStore for Redis instance:
 - Master-replica instance: Find the backup set that you want to restore and click **Restore Data** in the **Actions** column.
 - Cluster instance: Select the backup sets of all data shards that were generated at the same point in time and click **Restore Data** in the upper-right corner.

 **Warning** Risks may occur when you restore data. Proceed with caution. Verify the data that you want to restore before you restore the data.

5. In the message that appears, read the content and click **Continue**.

You can also restore backup data by cloning an instance. For more information, see [Clone an instance](#).

13.7.5. Clone an instance

KVStore for Redis allows you to create an instance from a specified backup set. The data in the new instance is the same as the data in the backup set. This feature can be applied in scenarios such as data recovery, quick workload deployment, and data verification.

Prerequisites

The KVStore for Redis instance is of the master-replica type.

Procedure

1. [Log on to the KVStore for Redis console](#).
2. On the **Instance List** page, click the ID of the instance.
3. In the left-side navigation pane, click **Backup and Recovery**.
4. Find the backup set that you want to restore and click **Clone Instance** in the **Actions** column.
5. In the message that appears, click **OK**.
6. On the Restore Instance page, configure the parameters and click **Submit**.

 **Note** For more information about the configurations of the new instance, see [Create an instance](#).

13.8. CloudDBA

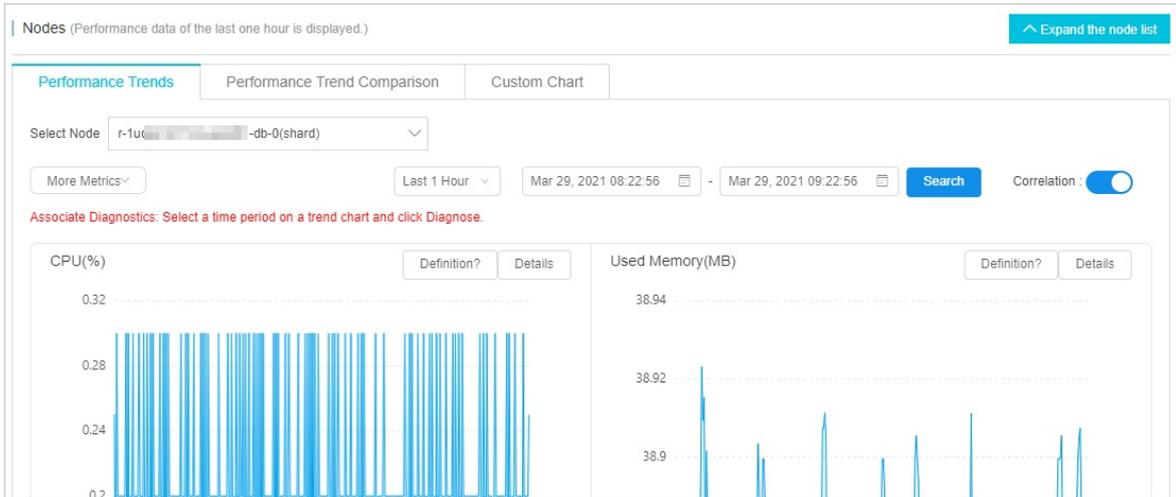
13.8.1. Performance trends

CloudDBA provides the performance trends feature that allows you to monitor the basic performance of a KVStore for Redis instance and the operational trends within a specified period of time. The performance trends include the CPU utilization, memory usage, queries per second (QPS), total connections, response time, data transfer, and key hit ratio.

Procedure

1. [Log on to the KVStore for Redis console](#).
2. On the **Instance List** page, click the ID of the instance.
3. In the left-side navigation pane, choose **CloudDBA > Performance Trends**.
4. You can use the following methods to view performance trends:

 **Note** If the KVStore for Redis instance uses a cluster architecture, the Performance Trends page displays the information about the nodes. The performance data during the last 1 hour is displayed. If you click the node ID, you can view the details of a specified node.



o Performance trends

On the Performance Trends tab, specify a time range and click Search.

Note

- By default, Correlation is enabled. If you move the pointer over the CPU chart to view the CPU metric of the KVStore for Redis instance at 09:00, other charts also display other metrics of the instance at 09:00.
- To view the definition of the performance metric and the performance trend, click Definition? and Details in the upper-right corner of the chart.

o Performance trend comparison

To compare the performance trends within two periods of time, click the Performance Trend Comparison tab, specify two periods of time, select more metrics, and then click Search.

o Custom chart

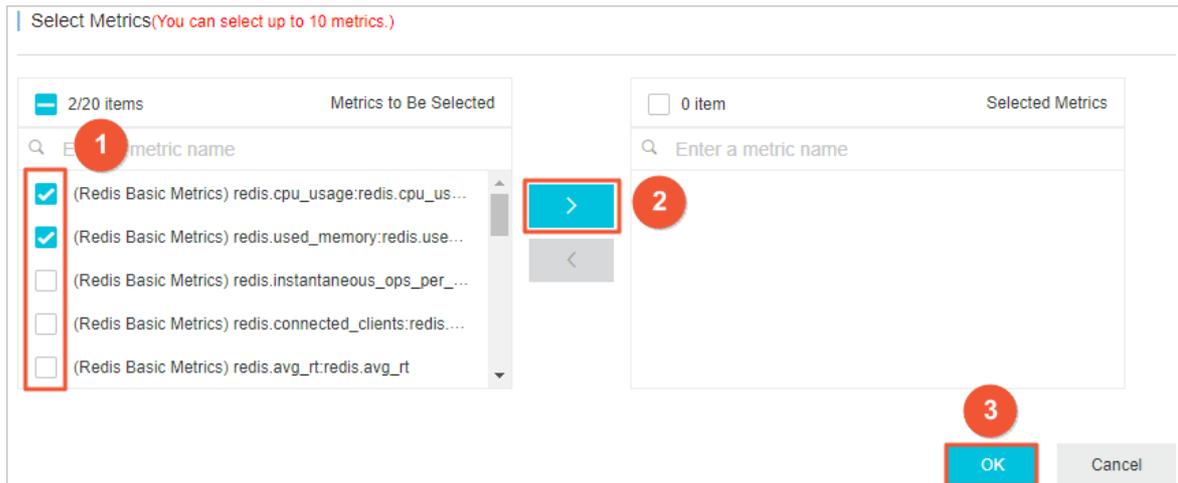
The preceding two methods display the basic metrics of a KVStore for Redis instance. If you want to display only basic metrics, you can configure custom performance trend charts. For more information, see [Add a performance trend chart](#).

13.8.2. Add a performance trend chart

The default performance trends tab displays the basic performance metrics of a KVStore for Redis instance. You can add a chart that contains only specified performance metrics to analyze the performance trends of your instances. This topic describes how to add a performance trend chart to a dashboard for KVStore for Redis instances.

Procedure

1. [Log on to the KVStore for Redis console](#).
2. On the Instance List page, click the ID of the instance.
3. In the left-side navigation pane, choose **CloudDBA > Performance Trends**.
4. Click the **Custom Chart** tab.
5. Click **Add Monitoring Dashboard**. In the Create Monitoring Dashboard dialog box, enter a dashboard name and click **OK**.
6. Click **+ Add Chart** or **Add Monitoring Chart**. Select the metrics that you want to add and click **OK**.



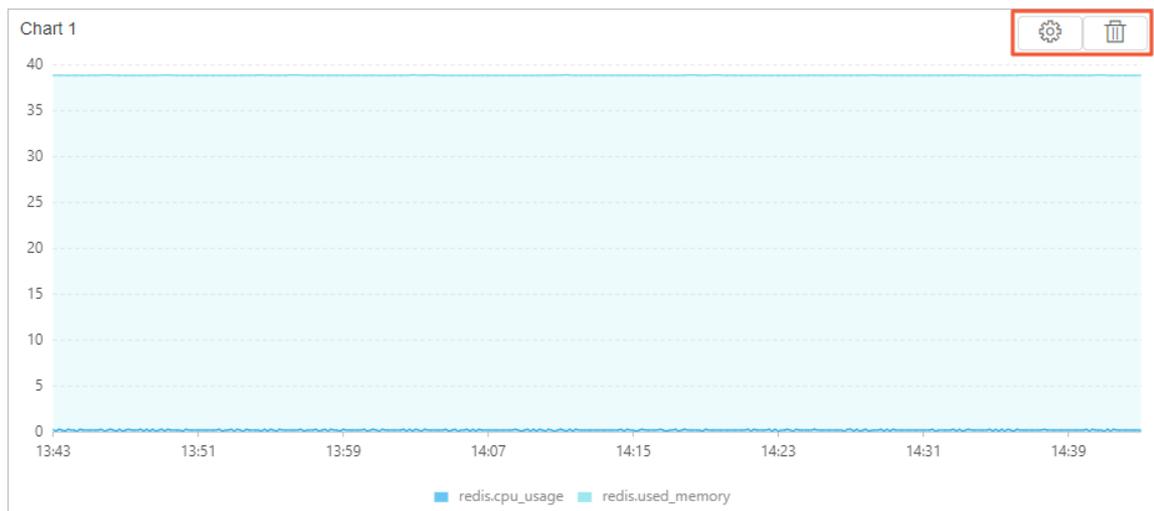
7. (Optional) You can view, modify, and delete monitoring dashboards.

- View a monitoring dashboard

Select the monitoring dashboard, specify a time range, and then click **Search**.

- Modify a monitoring dashboard

Click the following icons to modify or delete a chart in the monitoring dashboard.



- Delete a monitoring dashboard

Choose **Operate Dashboard > Delete Monitoring Dashboard**.

13.8.3. View performance metrics in real time

CloudDBA allows you to view the performance metrics of KVStore for Redis instances in real time. The performance metrics include information about CPU utilization, memory usage, queries per second (QPS), network traffic, servers, keys, clients, and connections.

Procedure

1. [Log on to the KVStore for Redis console](#).
2. On the **Instance List** page, click the ID of the instance.
3. In the left-side navigation pane, choose **CloudDBA > Real-time Performance**.
4. Select a view based on your business requirements.

In the upper part of the page, performance metrics are displayed in real time. The metrics include information about the server, keys, memory, clients, and connections. The details about the performance metrics are displayed in **Real-time Charts** and **Real-time Tables**.

The screenshot shows a 'Real-time Monitoring' dashboard with the following data:

- Server Information:** Version/Port/Uptime: 5.0.5 / 6379 / 17 Days 23 Hours 21 Minutes
- Key Information:** Total/Expiration Configured/Expired/Evicted: 1 / 0 / 0 / 0
- Memory Information:** Max /Used/System Memory/Fragmentation Rate: 2.00 GB / 77.70 MB / -- / 0.25
- Connection Information:** Established/Rejected: 14100961 / 0

Note The metrics are automatically updated every 5 minutes. Therefore, you can view real-time changes in performance. The remaining updating times are displayed in the upper-right corner. To stop updating the performance metrics, you can click **Pause**.

| View | Description |
|------------------|---|
| Real-time Charts | <p>Displays the real-time performance metrics of an instance in curve charts, such as key information, key hit information, key hit ratio, CPU utilization, memory information, QPS, and network traffic.</p> <p>The Real-time Charts section displays two line graphs. The left graph, 'Key Information (Total/Expiration Configured/Expired/Evicted)', shows four data series: Total Keys (blue), Expiration-configured Keys (green), Expired Keys (yellow), and Evicted Keys (purple). The right graph, 'Key Hit Information (hit/miss)', shows hit (blue) and miss (green) rates over time.</p> |
| Real-time Tables | <p>Displays the information about keys, QPS, memory usage, CPU utilization, network traffic, clients, and connections. The table displays up to 999 records. A new record is added to the table every 5 seconds.</p> <p>The Real-time Tables section shows a table with the following columns: time, Key (hit, miss, hitRate), QPS, Memory (used, fragment), CPU (sys, user), Network (In, Out), Clients (connected, blocked), and Connections (received, rejected). The table contains several rows of performance data.</p> |

13.8.4. Instance sessions

Instance sessions allow you to view the information about sessions between a KVStore for Redis instance and a client in real time, which includes the client information, commands that are run, and connection duration. You can also terminate abnormal sessions based on your business requirements.

Procedure

1. Log on to the KVStore for Redis console.
2. On the Instance List page, click the ID of the instance.
3. In the left-side navigation pane, choose **CloudDBA > Instance Sessions**.
4. Perform the operations that are described in the following table based on your business requirements.

- View sessions: By default, the details of all sessions are displayed. You can move the pointer over a specific parameter name to view information.

 **Note**

- You can enter keywords in the **search box** to filter session information.
- To refresh instance session information, click **Refresh** in the upper-left corner or enable **Auto Refresh** to automatically refresh the page every 30 seconds.

- Terminate sessions: Press the **Shift** key and select the specified session. To terminate the selected session, click **Kill Selected** in the upper-right corner. To terminate all sessions, click **Kill All**.

 **Warning** To avoid unexpected results, we recommend that you do not terminate system-level sessions.

- View session statistics: Session statistics record the total number of clients, active clients, and source IP addresses of instance sessions.

 **Note** In the **Statistics by Source** table, click the icon on the right of the source IP address to modify the source alias. In the **Total Sessions** column, click a value to view the details about a source IP address.

13.8.5. Slow queries

Slow queries reduce the stability of KVStore for Redis instances. To monitor and analyze slow queries, you can view the details about slow query logs in CloudDBA.

Procedure

1. [Log on to the KVStore for Redis console](#).
2. On the **Instance List** page, click the ID of the instance.
3. In the left-side navigation pane, choose **CloudDBA > Slow Queries**.
4. Query the details about the slow query logs.

Current Log

Slow Log Entries

Select Number of Log Entries: **100** **500** **1024** **Export**

| ID ↓↑ | Time ↓↑ | Query Duration (ms) ↓↑ | Query Statement | Client Address ↓↑ | Client Name |
|-------|------------------------|------------------------|------------------|-------------------|-------------|
| 71 | Mar 25, 2021, 22:22:26 | 89.63 | info all | ██████████:40728 | |
| 70 | Mar 25, 2021, 22:22:14 | 115.08 | info all | ██████████:57748 | |
| 69 | Mar 25, 2021, 22:22:14 | 93.89 | INFO replication | ██████████:65453 | |

 **Note** You can select the number of log entries to be displayed or enter keywords in the search box to filter log entries.

14. ApsaraDB for MongoDB

14.1. Usage notes

You must get familiar with the precautions and limits of ApsaraDB for MongoDB before you start.

To ensure the stability and security of ApsaraDB for MongoDB instances, take note of the limits, see Instance types in *ApsaraDB for MongoDB ApsaraDB for MongoDB limits*.

ApsaraDB for MongoDB limits

| Item | Limit |
|---------------------|--|
| Scale out nodes | You cannot scale out secondary nodes. |
| | <ul style="list-style-type: none">• When a replica set instance is created, three nodes are added.• ApsaraDB for MongoDB provides a primary node, a secondary node, and a hidden node for each replica set instance. The hidden node is invisible to you.• You cannot scale out secondary nodes. |
| Restart an instance | You must restart an ApsaraDB for MongoDB instance in the ApsaraDB for MongoDB console or by calling the API operation. |

14.2. Log on to the ApsaraDB for MongoDB console

This topic describes how to log on to the ApsaraDB for MongoDB console.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

Note When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.

4. In the top navigation bar, choose **Products > Database Services > ApsaraDB for MongoDB**.

14.3. Quick start

14.3.1. Use ApsaraDB for MongoDB

This topic is a quick start guide to basic usage operations for ApsaraDB for MongoDB, such as creating an instance, configuring a whitelist, and connecting to an instance. Flowcharts are used to describe the basic procedures in ApsaraDB for MongoDB, and guide you to create an ApsaraDB for MongoDB instance.



- **Create an ApsaraDB for MongoDB instance**

An instance is a virtual database server on which you can create and manage multiple databases.

- **Configure a whitelist for an ApsaraDB for MongoDB instance**

After you create an ApsaraDB for MongoDB instance, you need to configure a whitelist for the instance to allow external devices to access the instance.

A whitelist can enhance access security for ApsaraDB for MongoDB instances. We recommend that you update the whitelist on a regular basis. The normal services of the instance are not affected if you configure a whitelist.

- **Connect to a replica set instance by using the mongo shell**

After you create an instance and configure a whitelist, you can use the mongo shell to connect to the instance.

14.3.2. Create an ApsaraDB for MongoDB instance

This topic describes how to create an instance in the ApsaraDB for MongoDB console.

Prerequisites

An account is obtained to log on to the ApsaraDB for MongoDB console.

Create a replica set instance

1. **Log on to the ApsaraDB for MongoDB console.**
2. On the **Replica Set Instances** page, click **Create Instance** in the upper-left corner. On the **Create ApsaraDB for MongoDB Instance** page, configure the parameters.

The following table describes the required parameters.

Parameters for creating a replica set instance

| Section | Parameter | Description |
|----------------|-----------------------|---|
| Basic Settings | Organization | Select an organization for the new instance. |
| | Resource Set | Select a resource set for the new instance. |
| Region | Region | Select a region for the new instance. |
| | Zone | Select a zone for the new instance |
| Specifications | Database Engine | Select a database engine for the new instance. In this case, you can select only MongoDB . |
| | Engine Version | Select a database engine version for the new instance. Valid values: <ul style="list-style-type: none"> ◦ 3.0 ◦ 3.4 ◦ 4.0 ◦ 4.2 |
| | Node Type | ApsaraDB for MongoDB supports the following options: <ul style="list-style-type: none"> ◦ Three-node Replica Set: uses dedicated memory and I/O resources but shares CPU and storage resources with other general-purpose instances on the same server. ◦ Dedicated Instance: uses dedicated CPU, memory, storage, and I/O resources to ensure long-term performance. In this case, an instance is not affected by other instances on the same server. ◦ Dedicated Host: exclusively uses all resources of a server. This is the top configuration of exclusive specifications. |
| | Node Specifications | Select a node specification for the new instance. For more information, see descriptions in the ApsaraDB for MongoDB console. |
| | Storage Capacity (GB) | The storage space of the instance, which contains the space for data, system files, log files, and transaction files. For more information, see <i>Instance types</i> in <i>ApsaraDB for MongoDB Product Introduction</i> . |
| Network | Network Type | ApsaraDB for MongoDB supports the following options: <ul style="list-style-type: none"> ◦ Classic Network: Cloud services on the classic network are not isolated. Unauthorized access can be blocked only by security groups or whitelists. ◦ VPC: Virtual Private Cloud (VPC) is an isolated network environment built in Apsara Stack. You can customize route tables, CIDR blocks, and gateways in a VPC. We recommend that you select VPC for improved security. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p> Note If you select the VPC network type, you must configure the VPC and vSwitch parameters.</p> </div> |
| | VPC | Select a VPC. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p> Note When Network Type is set to VPC, you must specify this parameter.</p> </div> |

| Section | Parameter | Description |
|-------------------|------------------|---|
| | vSwitch | Select a vSwitch.  Note When Network Type is set to VPC , you must specify this parameter. |
| Password Settings | Instance Name | Set the name of the new instance. The name must be 2 to 256 characters in length, The name must start with a letter, and can contain digits, letters, underscores (_), and hyphens (-). |
| | Password Setting | Determine when to set the password for logging on to databases in the new instance. You can select Set Now to set the logon password immediately, or select Set after Purchase to set the logon password after you create the instance. For more information, see Reset the password for an ApsaraDB for MongoDB instance . |
| | Logon Password | Set a password. The password must meet the following requirements: <ul style="list-style-type: none"> It must be 8 to 32 characters in length. It must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. The name can contain special characters. Special characters include ! # \$ % ^ & * () _ + = |
| | Confirm Password | Enter the password again. The password you enter here must be the same as that in New Password. |

- Click **Submit** to create the instance.

Create a sharded cluster instance

- [Log on to the ApsaraDB for MongoDB console](#).
- In the left-side navigation pane, click **Sharded Cluster Instances**.
- Click **Create Instance** in the upper-left corner to go to the **Create MongoDB Instance** page. Configure the parameters.

The following table describes the required parameters.

Parameters for creating a sharded cluster instance

| Section | Parameter | MNS logs |
|----------------|-----------------|---|
| Basic Settings | Organization | Select an organization for the new instance. |
| | Resource Set | Select a resource set for the new instance. |
| Region | Region | Select a region for the new instance. |
| | Zone | Select a zone for the new instance |
| | Database Engine | Select a database engine for the new instance. In this case, you can select only MongoDB . |

| Specification Section | Parameter | MNS logs |
|------------------------------|------------------------------|--|
| | Engine Version | Select a database engine version for the new instance. Valid values: <ul style="list-style-type: none"> ◦ 3.4 ◦ 4.0 ◦ 4.2 |
| Network | Network Type | ApsaraDB for MongoDB supports the following options: <ul style="list-style-type: none"> ◦ Classic Network: Cloud services on the classic network are not isolated. Unauthorized access can be blocked only by security groups or whitelists. ◦ VPC: Virtual Private Cloud (VPC) is an isolated network environment built in Apsara Stack. You can customize route tables, CIDR blocks, and gateways in a VPC. We recommend that you select VPC for improved security. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note If you select the VPC network type, you must configure the VPC and vSwitch parameters. </div> |
| | VPC | Select a VPC. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note When Network Type is set to VPC, you must specify this parameter. </div> |
| | vSwitch | Select a vSwitch. If no vSwitch exists, you can click Create vSwitch to create a vSwitch. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note When Network Type is set to VPC, you must specify this parameter. </div> |
| Mongos Specifications | Mongos Specifications | The specifications of the mongos node. For more information, see descriptions in the ApsaraDB for MongoDB console. |
| | Quantity | The number of the mongos nodes. You can select 2 to 32 mongos nodes. |
| Shard Specifications | Shard Specifications | The specifications of the shard node. For more information, see descriptions in the ApsaraDB for MongoDB console. |
| | Storage Capacity (GB) | The storage space of the instance, which contains the space for data, system files, log files, and transaction files. For more information, see <i>Instance types in ApsaraDB for MongoDB Product Introduction</i> . |
| | Quantity | The number of the shard nodes. You can select 2 to 32 shard nodes. |
| Config Server Specifications | Config Server Specifications | The specifications of Configserver nodes. It is fixed at 1 vCPU, 2 GiB memory and cannot be customized. |
| | Storage Capacity (GB) | The storage space of the Configserver node. It is fixed at 20 GB and cannot be customized. |

| Section | Parameter | MNS logs |
|-------------------|------------------|---|
| Password Settings | Instance Name | Set the name of the new instance. The name must be 2 to 256 characters in length, and can contain digits, letters, underscores (_), and hyphens (-). It must start with a letter. |
| | Password Setting | Determine when to set the password for logging on to databases in the new instance. You can select Set Now to set the logon password immediately, or select Set after Purchase to set the logon password after you create the instance. For more information, see Reset the password for an ApsaraDB for MongoDB instance . |
| | Logon Password | Set a password. The password must meet the following requirements: <ul style="list-style-type: none"> ◦ It must be 8 to 32 characters in length. ◦ It must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. ◦ The name can contain special characters. Special characters include ! # \$ % ^ & * () _ + = |
| | Confirm Password | Enter the password again. The password you enter here must be the same as that in New Password. |

4. Click **Submit** to create the instance.

14.3.3. Reset the password for an ApsaraDB for MongoDB instance

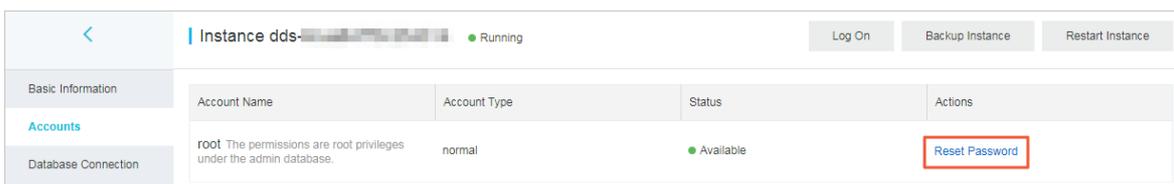
This topic describes how to reset your password in the ApsaraDB for MongoDB console.

Context

 **Notice** We recommend that you change your password on a regular basis to ensure data security.

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, click **Accounts**.
5. Click **Reset Password** in the **Actions** column and configure the parameters in the **Reset Password** panel.



[Parameters for resetting a password](#) describes the parameters.

Parameters for resetting a password

| Parameter | Description |
|----------------------|--|
| New Password | <ul style="list-style-type: none"> ◦ The password must be 8 to 32 characters in length and must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. Special characters include ! # \$ % ^ & * () _ + - = ◦ ! # \$ % ^ & * () _ + - = ◦ |
| Confirm New Password | Enter the password again. The password you enter here must be the same as that in New Password. |

6. Click OK.

14.3.4. Configure a whitelist for an ApsaraDB for MongoDB instance

This topic describes how to configure a whitelist for an ApsaraDB for MongoDB instance. Before you use an ApsaraDB for MongoDB instance, you must add the IP addresses or Classless Inter-Domain Routing (CIDR) blocks that you use for database access to a whitelist of this instance. This improves database security and stability. Proper configuration of whitelists can enhance access security of ApsaraDB for MongoDB. We recommend that you maintain the whitelists on a regular basis.

Context

The system creates a default whitelist for each instance. This whitelist can be modified or cleared, but it cannot be deleted. After an ApsaraDB for MongoDB instance is created, the system automatically adds the IP address 127.0.0.1 to the default whitelist of this instance. The IP address 0.0.0.0/0 indicates that all IP addresses are allowed to access this instance. You must add the IP addresses or CIDR blocks that you allow to access this ApsaraDB for MongoDB instance.

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the

Basic Information page appears.

4. In the left-side navigation pane, choose **Data Security** > **Whitelist Settings**.
5. You can manually configure a whitelist or import ECS Internal IP addresses to the whitelist.

Manually modify a whitelist

- i. Find the whitelist you want to modify and choose  > **Manually Modify** in the **Actions** column.

- ii. Enter IP addresses or CIDR blocks.

 **Note**

- Separate multiple IP addresses with commas (,). You can add a maximum of 1,000 different IP addresses to a whitelist. Supported formats are IP addresses such as 0.0.0.0/0 and 10.23.12.24, or CIDR blocks such as 10.23.12.24/24. /24 indicates the length of the IP address prefix in the CIDR block. An IP address prefix can contain 1 to 32 bits.
- If the IP whitelist is empty or only contains 0.0.0.0/0, all devices are granted access. This is risky for your ApsaraDB for MongoDB instance. We recommend that you add only the IP addresses or CIDR blocks of your own web servers to the whitelist.

- iii. Click **OK**.

Load IP addresses of ECS instances

- i. Find the whitelist, and choose  > **Import ECS Intranet IP** in the **Actions** column.
- ii. From the displayed internal IP addresses of ECS instances that belong to the current account, select the IP addresses and click  to add them to the whitelist.
- iii. Click **OK**.

14.3.5. Connect to an instance

14.3.5.1. Use DMS to log on to an ApsaraDB for MongoDB instance

You can use DMS to connect to an ApsaraDB for MongoDB instance.

Prerequisites

The IP address whitelist is configured. For more information about how to configure the IP address whitelist, see [Configure a whitelist for an ApsaraDB for MongoDB instance](#).

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. Click **Log On** in the upper-right corner of the page.

 **Note** For a sharded cluster instance, you must also select the Mongos node.

5. In the **Login instance** dialog box of the **DMS** console, check values of **Database type**, **Instance Area**, and **Connection string address**. If the information is correct, enter **Database account** and **Database password**, as shown in the following figure.

| Parameter | Description |
|---------------------------|---|
| Database type | The engine of the database. By default, the engine of the database to be connected is displayed. |
| Instance Area | The region where the instance is deployed. By default, the region of the current instance is displayed. |
| Connection string address | The endpoint of the instance. By default, the endpoint of the current instance is displayed. |
| Database account | The account of the database to be connected. |
| Database password | The password of the account used to connect to the database. |

6. Click **Login**.

Note If you want your web browser to remember the password, select **Remember password** before you click **Login**.

14.3.5.2. Use the mongo shell to connect to an ApsaraDB for MongoDB instance

This topic describes how to use the mongo shell to connect to an ApsaraDB for MongoDB instance. The mongo shell is a database management tool provided by ApsaraDB for MongoDB. You can install it on your client or in an Elastic Compute Service (ECS) instance.

Prerequisites

- The version of the mongo shell is the same as your ApsaraDB for MongoDB instance. This ensures successful authentication. For more information about the installation procedure, see [Install MongoDB](#). Choose the version

in the upper-left corner of the page based on your client version.

- The IP address of your client is added to a whitelist of the ApsaraDB for MongoDB instance. For more information, see [Configure a whitelist for an ApsaraDB for MongoDB instance](#).

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, click **Database Connections** to view connection strings.

Note

- Replica set instances: View the connection string or connection string URI of a node.
- Sharded cluster instances: View the connection string or connection string URI of a mongos node.

For more information about connection strings, see [Overview of replica set instance connections](#) or [Overview of sharded cluster instance connections](#).

5. Connect to the ApsaraDB for MongoDB instance from your client or ECS instance where the mongo shell is installed.

Replica set instances:

- Connection string of a node

During routine tests, you can directly connect to a primary or secondary node. Note that if the primary node fails, the system automatically switches to the secondary node, and the roles of connected nodes change. This affects the read and write operations of your application.

Command syntax:

```
mongo --host <host> -u <username> -p --authenticationDatabase <database>
```

Note

- <host>: the connection string of the primary or secondary node.
- <username>: the username used to log on to a database of the instance. The initial username is root.
- <database>: the name of the database corresponding to the username if authentication is enabled. If the database username is root, enter admin as the database name.

Example:

```
mongo --host dds-bp*****.mongodb.rds.aliyuncs.com:3717 -u root -p --authenticationDatabase admin
```

When **Enter password:** is displayed, enter the password of the database user and press the Enter key. If you forget the password of the root user, you can reset it. For more information, see [Reset the password for an ApsaraDB for MongoDB instance](#).

 **Note** The password characters are not displayed when you enter the password.

- HA connection (recommended): You can use a connection string URI to connect to both the primary and secondary nodes of a replica set instance. This ensures that your application is always connected to the primary node and the read and write operations of your application are not affected even if the roles of the primary and secondary nodes are switched.

Command syntax:

```
mongo "<ConnectionStringURI>"
```

 **Note**

- The connection string URI must be enclosed in a pair of double quotation marks ("").
- <ConnectionStringURI>: the connection string URI of the instance.
You must replace **** in the connection string URI with the database password. For more information about how to set a database password, see [Reset the password for an ApsaraDB for MongoDB instance](#).

Example:

```
mongo "mongodb://root:****@dds-*****.mongodb.rds.intra.env17e.shuguang.com:3717,dds-*****.mongodb.rds.intra.env17e.shuguang.com:3717/admin?replicaSet=mgset-*****"
```

Sharded cluster instances:

- Connection string of a mongos node

Command syntax:

```
mongo --host <mongos_host> -u <username> -p --authenticationDatabase <database>
```

 **Note**

- <mongos_host>: the connection string of a mongos node in the sharded cluster instance.
- <username>: the username used to log on to a database of the instance. The initial username is root.
- <database>: the name of the database corresponding to the username if authentication is enabled. If the database username is root, enter admin as the database name.

Example:

```
mongo --host s-bp*****.mongodb.rds.aliyuncs.com:3717 -u root -p --authenticationDatabase admin
```

- HA connection (recommended): You can use a connection string URI to connect to a database. If one mongos node fails, another mongos node takes over business to ensure the high availability of the connection.

Command syntax:

```
mongo "<ConnectionStringURI>"
```

 **Note**

- The connection string URI must be enclosed in a pair of double quotation marks ("").
- <ConnectionStringURI>: the connection string URI of the instance.
You must replace **** in the connection string URI with the database password. For more information about how to set a database password, see [Reset the password for an ApsaraDB for MongoDB instance](#).

Example:

```
mongo "mongodb://root:****@s-*****.mongodb.rds.intra.env17e.shuguang.com:3717,s-*****.mongodb.rds.intra.env17e.shuguang.com:3717/admin"
```

14.3.5.3. Introduction to connection strings and URIs

14.3.5.3.1. Overview of replica set instance connections

ApsaraDB for MongoDB supports both connection strings and connection string URIs. You can use a connection string to connect to the primary or secondary node, and use a connection string URI to connect to both of them. For high availability, we recommend that you use connection string URIs to connect your application to both primary and secondary nodes. This topic provides an overview of replica set instance connections.

Prerequisites

A whitelist is configured for the instance. For more information, see [Configure a whitelist for an ApsaraDB for MongoDB instance](#).

View connection strings

1. [Log on to the ApsaraDB for MongoDB console](#).
2. On the **Replica Set Instances** page, find the instance and then click the instance ID or choose  **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
3. In the left-side navigation pane, click **Database Connections** to view connection strings.



Description of connection strings

| Item | Description |
|-------------------|--|
| Connection type | <ul style="list-style-type: none"> • Internal Connections - Classic Network: Cloud services on the classic network are not isolated. Unauthorized access can be blocked only by security groups or whitelists. • Internal Connections - VPC: A virtual private cloud (VPC) is an isolated network with higher security and performance than the classic network. By default, ApsaraDB for MongoDB provides VPC endpoints. |
| Role | <ul style="list-style-type: none"> • Primary: the primary node in the replica set instance. If you connect to this node, you can perform read and write operations on the databases of the replica set instance. • Secondary: the secondary node in the replica set instance. If you connect to this node, you can perform only read operations on the databases of the replica set instance. • Connection String URI: ApsaraDB for MongoDB allows you to use a connection string URI to connect to a replica set instance to achieve load balancing and high availability. |
| Connection string | <p>The connection string of a primary or secondary node is in the following format:</p> <pre><host>:<port></pre> <ul style="list-style-type: none"> • <host>: the endpoint used to connect to the instance. • <port>: the port used to connect to the instance. |

| Item | Description |
|-----------------------|---|
| Connection string URI | <p>A connection string URI is in the following format:</p> <pre style="background-color: #f0f0f0; padding: 5px;">mongodb://[username:password@[host1[:port1][,host2[:port2],...[,hostN[:portN]]]]/[database][?options]]</pre> <ul style="list-style-type: none"> mongodb://: the prefix of a connection string URI. username:password@: the username and password used to log on to a database of the replica set instance. You must separate them with a colon (:). hostX:portX: the endpoint and port of a node in the replica set instance. /database: the name of the database corresponding to the username if authentication is enabled. ?options: additional connection options. <div style="background-color: #e0f0ff; padding: 5px; border: 1px solid #add8e6;"> <p> Note If your application is in a production environment, we recommend that you use a connection string URI to connect to the instance. This way, when a node fails, the read and write operations of your application are not affected as a result of the failover.</p> </div> |

Related information

- [Connect to a replica set instance by using the mongo shell](#)

14.3.5.3.2. Overview of sharded cluster instance connections

ApsaraDB for MongoDB supports both connection strings and connection string URIs. You can use a connection string to connect to a single mongos node, and use a connection string URI to connect to multiple mongos nodes. For high availability, we recommend that you use connection string URIs to connect your application to multiple mongos nodes. This topic provides an overview of sharded cluster instance connections.

View connection strings

1. [Log on to the ApsaraDB for MongoDB console.](#)
2. In the left-side navigation pane, click **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, click **Database Connections** to view connection strings.

| Intranet Connection - Classic Network | | | | | Update Connection String | |
|---------------------------------------|-----------|------|---|---------|------------------------------------|--------------------------|
| ID | Node Type | Node | Address | Actions | | |
| s-xxxxxxxxxxxx | Mongos | - | s-xxxxxxxxxxxx.mongodb.rds.thirteenth-inc.com:3717 | Release | | |
| s-xxxxxxxxxxxx | Mongos | - | s-xxxxxxxxxxxx.mongodb.rds.thirteenth-inc.com:3717 | Release | | |
| ConnectionStringURI | Mongos | - | mongodb://root:****@s-xxxxxxxxxxxx.mongodb.rds.thirteenth-inc.com:3717,s-xxxxxxxxxxxx.mongodb.rds.thirteenth-inc.com:3717/admin | Release | | |
| Public IP Connection | | | | | Apply for Public Connection String | Update Connection String |
| ID | Node Type | Node | Address | Actions | | |
| s-xxxxxxxxxxxx | Mongos | - | s-xxxxxxxxxxxx.pub.mongodb.rds.thirteenth-inc.com:3717 | Release | | |
| s-xxxxxxxxxxxx | Mongos | - | s-xxxxxxxxxxxx.pub.mongodb.rds.thirteenth-inc.com:3717 | Release | | |
| ConnectionStringURI | Mongos | - | mongodb://root:****@s-xxxxxxxxxxxx.pub.mongodb.rds.thirteenth-inc.com:3717,s-xxxxxxxxxxxx.pub.mongodb.rds.thirteenth-inc.com:3717/admin | | | |

Description of connection strings

| Item | Description |
|-----------------------|---|
| Connection type | <ul style="list-style-type: none"> Internal Connections - Classic Network: Cloud services on the classic network are not isolated. Unauthorized access can be blocked only by security groups or whitelists. Internal Connections - VPC: A virtual private cloud (VPC) is an isolated network with higher security and performance than the classic network. By default, ApsaraDB for MongoDB provides VPC endpoints. Public Connections: Security risks may arise if you connect to a sharded cluster instance over the Internet. Therefore, ApsaraDB for MongoDB does not provide public endpoints. If you want to connect to an ApsaraDB for MongoDB instance from a device outside of Alibaba Cloud (such as an on-premises device), you must apply for a public endpoint. For more information, see Apply for a public endpoint for an ApsaraDB for MongoDB instance. |
| Mongos ID | <p>The connection string of a mongos node is in the following format:</p> <pre><host>:<port></pre> <ul style="list-style-type: none"> <host>: the endpoint used to connect to the instance. <port>: the port used to connect to the instance. <p>Note During routine tests, you can use a connection string to directly connect to a mongos node.</p> |
| Connection string URI | <p>A connection string URI is in the following format:</p> <pre>mongodb://[username:password@]host1[:port1][,host2[:port2],...[,hostN[:portN]]][/[database][?options]]</pre> <ul style="list-style-type: none"> mongodb://: the prefix of a connection string URI. username:password@: the username and password used to log on to a database of the sharded cluster instance. You must separate them with a colon (:). hostX:portX: the endpoint and port of a mongos node in the sharded cluster instance. /database: the name of the database corresponding to the username if authentication is enabled. ?options: additional connection options. <p>Note If your application is in a production environment, we recommend that you use a connection string URI to connect to the instance. Then, your client can automatically distribute your requests to multiple mongos nodes to balance loads. If a mongos node fails, your client automatically redirects requests to other mongos nodes in the normal state.</p> |

Log on to a database of an ApsaraDB for MongoDB instance

- Obtain the [connection string](#) and the following information:
 - The username used to log on to the database. The initial username is root.
 - The password of the database user. If you forget the password of the root user, you can reset it. For more information, see [Reset the password for an ApsaraDB for MongoDB instance](#).
 - The name of the database corresponding to the username if authentication is enabled. If the database username is root, enter admin as the database name.
- Log on to the database.
 - [Use DMS to log on to a replica set instance of ApsaraDB for MongoDB](#)
 - [Connect to a replica set instance by using the mongo shell](#)

14.4. Instances

14.4.1. Create an ApsaraDB for MongoDB instance

This topic describes how to create an instance in the ApsaraDB for MongoDB console.

Prerequisites

An account is obtained to log on to the ApsaraDB for MongoDB console.

Create a replica set instance

1. [Log on to the ApsaraDB for MongoDB console.](#)
2. On the **Replica Set Instances** page, click **Create Instance** in the upper-left corner. On the **Create ApsaraDB for MongoDB Instance** page, configure the parameters.

The following table describes the required parameters.

Parameters for creating a replica set instance

| Section | Parameter | Description |
|----------------|---------------------|--|
| Basic Settings | Organization | Select an organization for the new instance. |
| | Resource Set | Select a resource set for the new instance. |
| Region | Region | Select a region for the new instance. |
| | Zone | Select a zone for the new instance |
| Specifications | Database Engine | Select a database engine for the new instance. In this case, you can select only MongoDB . |
| | Engine Version | Select a database engine version for the new instance. Valid values: <ul style="list-style-type: none"> ◦ 3.0 ◦ 3.4 ◦ 4.0 ◦ 4.2 |
| | Node Type | ApsaraDB for MongoDB supports the following options: <ul style="list-style-type: none"> ◦ Three-node Replica Set: uses dedicated memory and I/O resources but shares CPU and storage resources with other general-purpose instances on the same server. ◦ Dedicated Instance: uses dedicated CPU, memory, storage, and I/O resources to ensure long-term performance. In this case, an instance is not affected by other instances on the same server. ◦ Dedicated Host: exclusively uses all resources of a server. This is the top configuration of exclusive specifications. |
| | Node Specifications | Select a node specification for the new instance. For more information, see descriptions in the ApsaraDB for MongoDB console. |

| Section | Parameter | Description |
|--------------------------|------------------------------|---|
| | Storage Capacity (GB) | The storage space of the instance, which contains the space for data, system files, log files, and transaction files. For more information, see <i>Instance types</i> in <i>ApsaraDB for MongoDB Product Introduction</i> . |
| Network | Network Type | <p>ApsaraDB for MongoDB supports the following options:</p> <ul style="list-style-type: none"> ◦ Classic Network: Cloud services on the classic network are not isolated. Unauthorized access can be blocked only by security groups or whitelists. ◦ VPC: Virtual Private Cloud (VPC) is an isolated network environment built in Apsara Stack. You can customize route tables, CIDR blocks, and gateways in a VPC. We recommend that you select VPC for improved security. <p>Note If you select the VPC network type, you must configure the VPC and vSwitch parameters.</p> |
| | VPC | <p>Select a VPC.</p> <p>Note When Network Type is set to VPC, you must specify this parameter.</p> |
| | vSwitch | <p>Select a vSwitch.</p> <p>Note When Network Type is set to VPC, you must specify this parameter.</p> |
| Password Settings | Instance Name | Set the name of the new instance. The name must be 2 to 256 characters in length, The name must start with a letter, and can contain digits, letters, underscores (_), and hyphens (-). |
| | Password Setting | Determine when to set the password for logging on to databases in the new instance. You can select Set Now to set the logon password immediately, or select Set after Purchase to set the logon password after you create the instance. For more information, see Reset the password for an ApsaraDB for MongoDB instance . |
| | Logon Password | <p>Set a password. The password must meet the following requirements:</p> <ul style="list-style-type: none"> ◦ It must be 8 to 32 characters in length. ◦ It must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. ◦ The name can contain special characters. Special characters include ! # \$ % ^ & * () _ + = |
| | Confirm Password | Enter the password again. The password you enter here must be the same as that in New Password. |

3. Click **Submit** to create the instance.

Create a sharded cluster instance

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Sharded Cluster Instances**.
3. Click **Create Instance** in the upper-left corner to go to the **Create MongoDB Instance** page. Configure the

parameters.

The following table describes the required parameters.

Parameters for creating a sharded cluster instance

| Section | Parameter | MNS logs |
|-----------------------|-----------------------|---|
| Basic Settings | Organization | Select an organization for the new instance. |
| | Resource Set | Select a resource set for the new instance. |
| Region | Region | Select a region for the new instance. |
| | Zone | Select a zone for the new instance |
| Specifications | Database Engine | Select a database engine for the new instance. In this case, you can select only MongoDB . |
| | Engine Version | Select a database engine version for the new instance. Valid values: <ul style="list-style-type: none"> ◦ 3.4 ◦ 4.0 ◦ 4.2 |
| Network | Network Type | ApsaraDB for MongoDB supports the following options: <ul style="list-style-type: none"> ◦ Classic Network: Cloud services on the classic network are not isolated. Unauthorized access can be blocked only by security groups or whitelists. ◦ VPC: Virtual Private Cloud (VPC) is an isolated network environment built in Apsara Stack. You can customize route tables, CIDR blocks, and gateways in a VPC. We recommend that you select VPC for improved security. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note If you select the VPC network type, you must configure the VPC and vSwitch parameters. </div> |
| | VPC | Select a VPC. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note When Network Type is set to VPC, you must specify this parameter. </div> |
| | vSwitch | Select a vSwitch. If no vSwitch exists, you can click Create vSwitch to create a vSwitch. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note When Network Type is set to VPC, you must specify this parameter. </div> |
| Mongos Specifications | Mongos Specifications | The specifications of the mongos node. For more information, see descriptions in the ApsaraDB for MongoDB console. |
| | Quantity | The number of the mongos nodes. You can select 2 to 32 mongos nodes. |

| Section | Parameter | MNS logs |
|------------------------------|------------------------------|---|
| Shard Specifications | Shard Specifications | The specifications of the shard node. For more information, see descriptions in the ApsaraDB for MongoDB console. |
| | Storage Capacity (GB) | The storage space of the instance, which contains the space for data, system files, log files, and transaction files. For more information, see <i>Instance types in ApsaraDB for MongoDB Product Introduction</i> . |
| | Quantity | The number of the shard nodes. You can select 2 to 32 shard nodes. |
| Config Server Specifications | Config Server Specifications | The specifications of Configserver nodes. It is fixed at 1 vCPU, 2 GiB memory and cannot be customized. |
| | Storage Capacity (GB) | The storage space of the Configserver node. It is fixed at 20 GB and cannot be customized. |
| Password Settings | Instance Name | Set the name of the new instance. The name must be 2 to 256 characters in length, and can contain digits, letters, underscores (_), and hyphens (-). It must start with a letter. |
| | Password Setting | Determine when to set the password for logging on to databases in the new instance. You can select Set Now to set the logon password immediately, or select Set after Purchase to set the logon password after you create the instance. For more information, see Reset the password for an ApsaraDB for MongoDB instance . |
| | Logon Password | Set a password. The password must meet the following requirements: <ul style="list-style-type: none"> ◦ It must be 8 to 32 characters in length. ◦ It must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. ◦ The name can contain special characters. Special characters include ! # \$ % ^ & * () _ + = |
| | Confirm Password | Enter the password again. The password you enter here must be the same as that in New Password. |

4. Click **Submit** to create the instance.

14.4.2. View the details of an ApsaraDB for MongoDB instance

This topic describes how to view the details of an ApsaraDB for MongoDB instance, such as the basic information, internal network connection information, status, and configurations. This topic describes how to view the details of an ApsaraDB for MongoDB instance.

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Go to the instance details page. Either of the following methods can be used:
 - Find the instance and click its ID to go to the **Basic Information** page, where you can view the details of

the instance.

- In the Operations column corresponding to the instance, choose  > **Manage** to go to the **Basic Information** page, where you can view the details of the instance.

14.4.3. Restart an ApsaraDB for MongoDB instance

You can manually restart an instance when the number of connections exceeds the threshold or any performance issue occurs on the instance. This topic describes how to restart an ApsaraDB for MongoDB instance.

Prerequisites

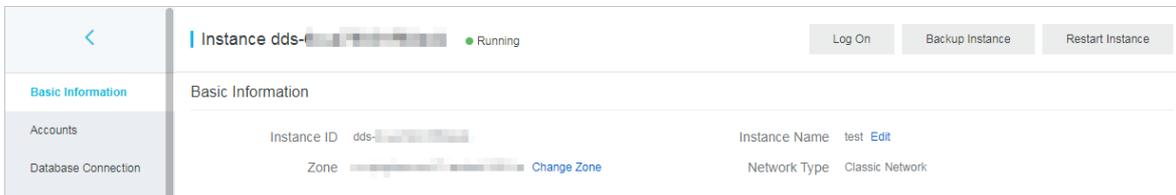
The instance is in the **Running** state.

Context

 **Note** When an ApsaraDB for MongoDB instance is restarted, all its connections are terminated. Plan your operations in advance before you restart an ApsaraDB for MongoDB instance. Proceed with caution.

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the upper-right corner of the page, click **Restart Instance**.



 **Note** You can also choose  > **Restart** in the **Actions** column corresponding to the instance.

5. In the **Restart Instance** message, click **OK**.

14.4.4. Change the specifications of an ApsaraDB for MongoDB instance

This topic describes how to change the specifications of an ApsaraDB for MongoDB instance. You can upgrade or downgrade an ApsaraDB for MongoDB instance to meet your business needs.

Prerequisites

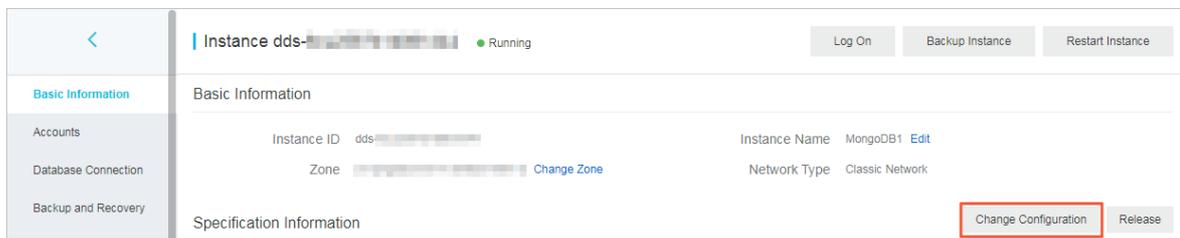
The instance must be an ApsaraDB for MongoDB replica set instance.

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. On the **Replica Set Instances** page, find the instance and then click the instance ID or choose  **Manage** in

the **Actions** column. Then, the **Basic Information** page appears.

3. Click **Change Configuration** in the upper right corner of the Specification Information section to go to the **Modify Instance** page.



Note To go to the **Modify Instance** page, you can also choose **Change Configuration** in the **Actions** column corresponding to the instance on the **Replica Set Instances** page.

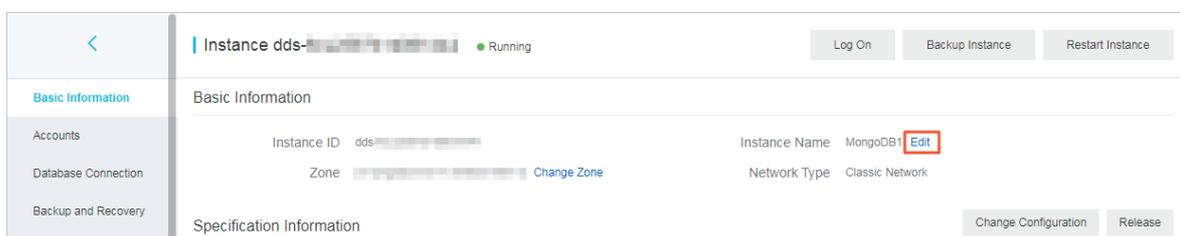
4. On the **Modify Instance** page, change the instance specifications.
You can change values of the following parameters:
 - o **Node Type**
 - o **Node Specifications**
 - o **Storage Capacity**
5. Click **Submit**.

14.4.5. Change the name of an ApsaraDB for MongoDB instance

This topic describes how to change the name of an ApsaraDB for MongoDB instance for better management.

Procedure

1. **Log on to the ApsaraDB for MongoDB console.**
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. Click **Edit** next to **Instance Name**.



Note

- o The instance name must start with an English letter. It cannot start with `http://` or `https://`.
- o The instance name can contain letters, underscores (`_`), hyphens (`-`), and digits.
- o The instance name must be 2 to 128 characters in length.

5. Click **OK**.

14.4.6. Reset the password for an ApsaraDB for MongoDB instance

This topic describes how to reset your password in the ApsaraDB for MongoDB console.

Context

 **Notice** We recommend that you change your password on a regular basis to ensure data security.

Procedure

1. Log on to the ApsaraDB for MongoDB console.
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, click **Accounts**.
5. Click **Reset Password** in the **Actions** column and configure the parameters in the **Reset Password** panel.



[Parameters for resetting a password](#) describes the parameters.

Parameters for resetting a password

| Parameter | Description |
|-----------------------------|---|
| New Password | <ul style="list-style-type: none"> o The password must be 8 to 32 characters in length and must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters. Special characters include ! # \$ % ^ & * () _ + - = o ! # \$ % ^ & * () _ + - = |
| Confirm New Password | Enter the password again. The password you enter here must be the same as that in New Password. |

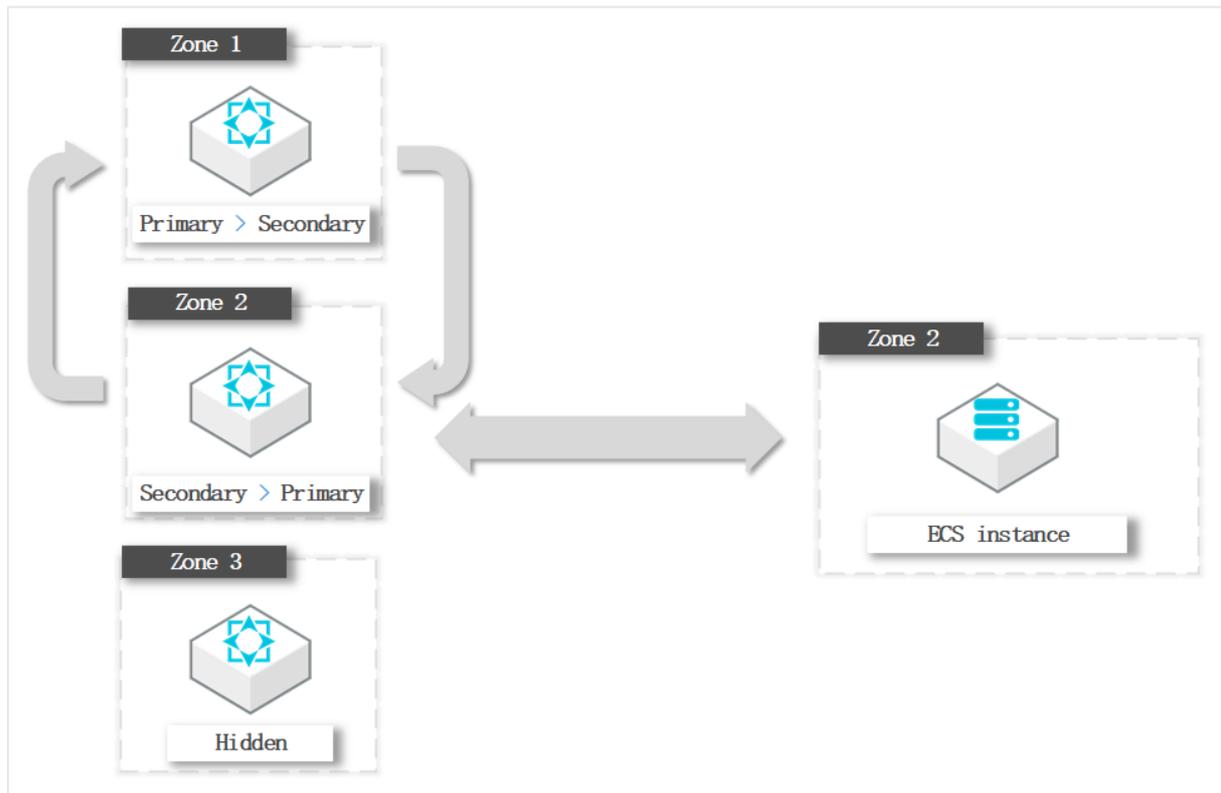
6. Click **OK**.

14.4.7. Switch node roles

You can switch the node roles of an ApsaraDB for MongoDB instance in the ApsaraDB for MongoDB console based on your business deployment.

Typical scenario

When an ECS instance and an ApsaraDB for MongoDB instance are in the same zone and connected over the internal network, the latency is minimal. If they are connected across different zones, the latency increases and the performance of ApsaraDB for MongoDB instances and your business will be affected.



In this example, the ECS instance to which the application belongs is in Zone 2. If the primary node of the ApsaraDB for MongoDB instance is in Zone 1, the ECS instance needs to connect to the primary node across zones.

To optimize the business deployment architecture, you can switch the roles of the primary and secondary nodes. In this example, you can change the role of the node in Zone 2 to primary and the role of the node in Zone 1 to secondary. Note that only the node roles are changed. ECS and ApsaraDB for MongoDB instances can be connected in the same zone without changing the actual zones and role IDs.

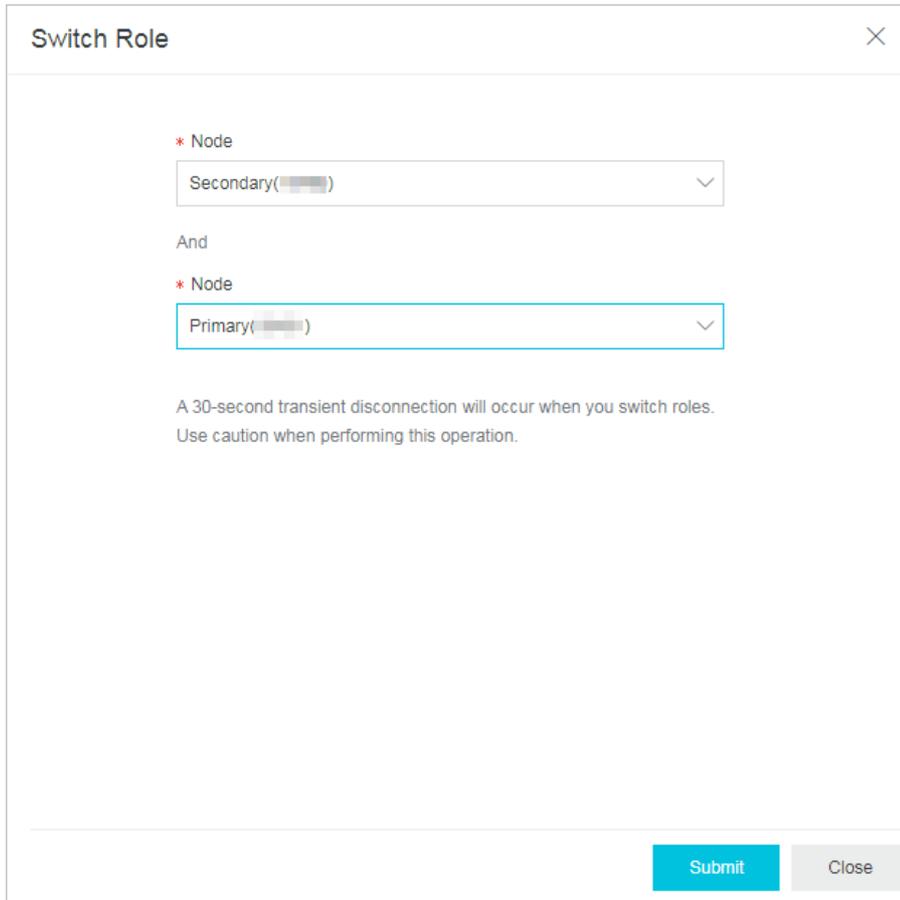
Precautions

- Switching node roles will cause a transient disconnection of up to 30 seconds. Perform this operation during off-peak hours or ensure that your application has a reconnection mechanism.
- Switching node roles only changes the roles of nodes, but not the zones and role IDs of nodes.

Procedure

1. [Log on to the ApsaraDB for MongoDB console.](#)
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, click **Service Availability**.
5. Subsequent steps on the **Service Availability** page vary depending on instance types.
 - Replica set instance
 - a. Click **Switch Role** in the upper-right corner of the page.

- b. In the **Switch Role** dialog box that appears, select the nodes.



Switch Role [X]

* Node
Secondary()

And

* Node
Primary()

A 30-second transient disconnection will occur when you switch roles.
Use caution when performing this operation.

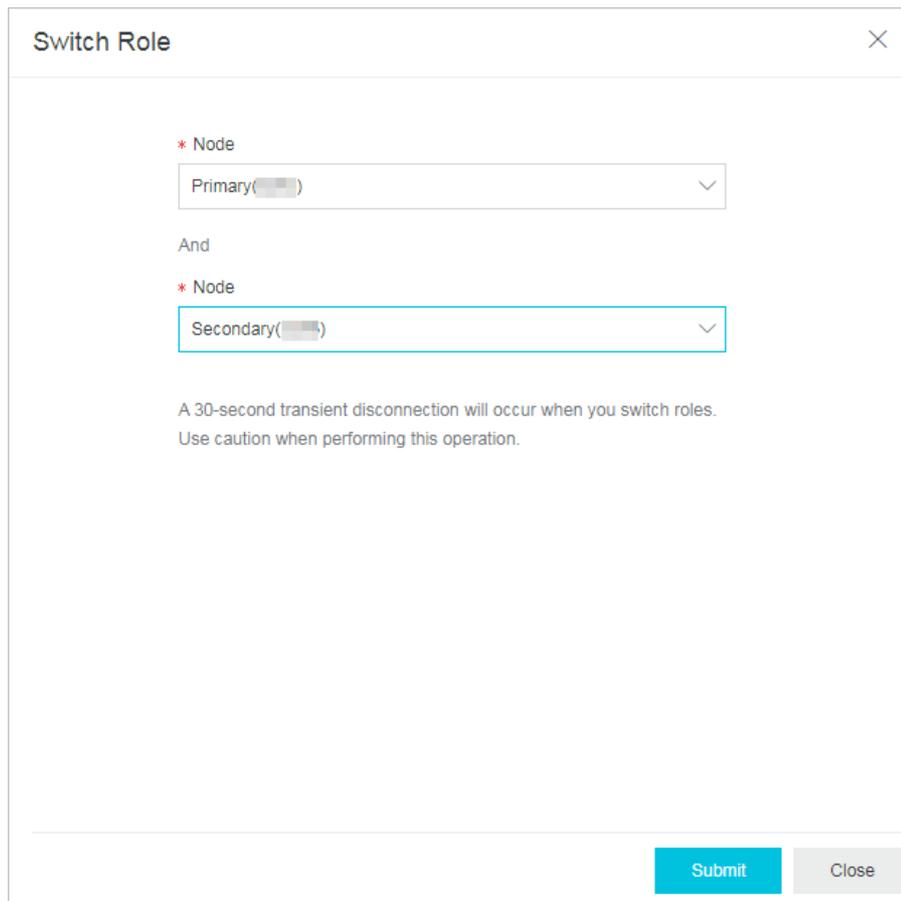
Submit Close

- c. Click **OK**.
- o Sharded cluster instance

Note For sharded cluster instances, you can only manage the zone distribution of shard and Configserver nodes.

- a. In the upper-right corner of the **Zone Distribution for Shards** or **Zone Distribution for Configservers** section, click **Switch Role**.

- b. In the **Switch Role** dialog box that appears, select the nodes.



Switch Role

* Node
Primary()

And

* Node
Secondary()

A 30-second transient disconnection will occur when you switch roles.
Use caution when performing this operation.

Submit Close

- Click **OK**.

14.4.8. Migrate an ApsaraDB for MongoDB instance across zones in the same region

This topic describes how to migrate an ApsaraDB for MongoDB instance across zones in the same region. After instances are migrated to other zones, the attributes, specifications, and connection strings of the instances remain unchanged.

Prerequisites

- The ApsaraDB for MongoDB instance is a replica set instance or sharded cluster instance that runs MongoDB 4.2 or later.
- Transparent data encryption (TDE) is not enabled for the ApsaraDB for MongoDB instance.
- The destination zone must be in the same region as the current zone of the instance.
- If the instance is in a VPC, make sure that a vSwitch is created in the destination zone before you start migration. For more information, see *Create a VPC* and *Create a vSwitch* in *VPC User Guide*.
- The instance does not have a public endpoint. If you have applied for a public endpoint, you must release it before migration. For more information, see [Release a public connection string](#).

Precautions

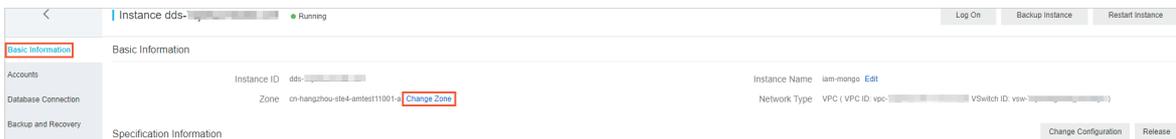
- If the instance is in a VPC, you cannot change the VPC when the instance is in the migration process.
- The time required varies based on factors such as the network conditions, task queue status, and data volume. We recommend that you migrate the instance across zones during off-peak hours.

- During the migration, a transient connection of 30 seconds occurs. Make sure that your application is configured to reconnect to the instance after it is disconnected.
- The virtual IP addresses (VIPs) of the instance, such as 172.16.88.60, are changed when the instance is migrated across zones. If your application uses the original VIP, the application cannot connect to the instance after the migration.

Note We recommend that you use a connection string URI to connect to the instance, which ensures high availability. For more information, see [Overview of replica set instance connections](#) or [Overview of sharded cluster instance connections](#).

Procedure

1. Log on to the ApsaraDB for MongoDB console.
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the **Basic Information** section, click **Change Zone**.



5. In the **Migrate Instance to Other Zone** panel, configure parameters based on the network type of the instance.
 - o VPC

Migrate Instance to Other Zone ✕

Instance: dds-██████████

Current Zone: cn-hangzhou-ste4-amtest11001-a

Migrate To:

VPC: vpc-██████████

Select a VSwitch:

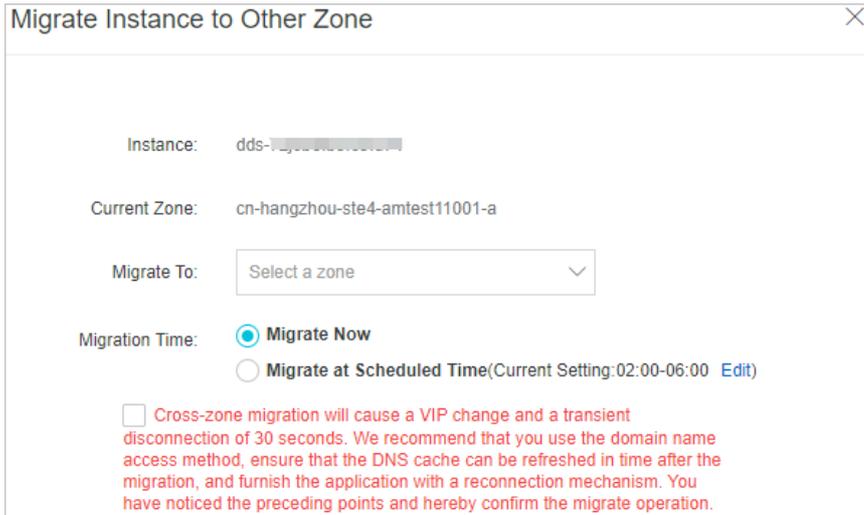
Migration Time: **Migrate Now**
 Migrate at Scheduled Time(Current Setting: 02:00-06:00 [Edit](#))

Cross-zone migration will cause a VIP change and a transient disconnection of 30 seconds. We recommend that you use the domain name access method, ensure that the DNS cache can be refreshed in time after the migration, and furnish the application with a reconnection mechanism. You have noticed the preceding points and hereby confirm the migrate operation.

| Parameter | Description |
|-------------------------|---------------------------------|
| Migrate To | Select the destination zone. |
| Select a VSwitch | Select the destination vSwitch. |

| Parameter | Description |
|-----------------------|--|
| Migration Time | Select the time when you want to start the migration. <ul style="list-style-type: none"> ▪ Switch Immediately after Migration: The migration immediately starts. When the instance status changes to Running, the migration is complete. ▪ Switch within Maintenance Window: The migration starts during the specified period. You can click Edit next to Switch within Maintenance Window to change the period. |

- o Classic network



| Parameter | Description |
|-----------------------|--|
| Migrate To | Select the destination zone. |
| Migration Time | Select the time when you want to start the migration. <ul style="list-style-type: none"> ▪ Switch Immediately after Migration: The migration immediately starts. When the instance status changes to Running, the migration is complete. ▪ Switch within Maintenance Window: The migration starts during the specified period. You can click Edit next to Switch within Maintenance Window to change the period. |

6. Read the message that is displayed and select the check box next to the message.
7. Click OK.

14.4.9. Release an ApsaraDB for MongoDB instance

This topic describes how to manually release an ApsaraDB for MongoDB instance to meet your business needs. This topic describes how to manually release an ApsaraDB for MongoDB instance.

Procedure

1. [Log on to the ApsaraDB for MongoDB console.](#)
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.

4. In the lower-right corner of the **Basic Information** section, click **Release**.

 **Note** You can also go to the Replica Set Instances page, find the instance, click the  icon in the **Actions** column, and then choose **Release**.

5. In the **Release Instance** message, click **OK**.

 **Warning** After you release an ApsaraDB for MongoDB instance, data in the instance can no longer be recovered. Proceed with caution.

14.4.10. Primary/secondary failover

14.4.10.1. Trigger a primary/secondary failover for a replica set instance

An ApsaraDB for MongoDB replica set instance consists of three nodes by default. ApsaraDB for MongoDB provides connection strings for you to connect to the primary node and a secondary node. The other secondary node is hidden as a backup to ensure high availability. If a node is faulty, the high availability system of ApsaraDB for MongoDB automatically triggers a primary/secondary failover to ensure the availability of the instance. You also can manually trigger a primary/secondary failover for an ApsaraDB for MongoDB instance in scenarios such as routine disaster recovery drills.

Context

After you log on to the ApsaraDB for MongoDB console or call the `SwitchDBInstanceHA` operation to trigger a primary/secondary failover for a replica set instance, ApsaraDB for MongoDB interchanges the roles of the primary and secondary nodes.

Note

- You can trigger a primary/secondary failover only for replica set and sharded cluster instances, but not for standalone instances due to their single-node architecture.
- Each time you trigger a primary/secondary failover for an instance, the instance may have a transient connection error of about 30 seconds. Ensure that your applications can automatically re-establish a connection.
- You can trigger a primary/secondary failover only for instances in the running state.

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. On the **Replica Set Instances** page, find the instance and then click the instance ID or choose  **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
3. In the **Node List** section, click **Failover**, as shown in the following figure.

| Node List | | | | Failover |
|-----------|---------|--------------------|------|-----------|
| Role | Role ID | Domain Information | Port | Operation |
| Primary | | | 3717 | ⋮ |
| Secondary | | | 3717 | ⋮ |

- In the dialog box that appears, click **OK**.
- The instance status changes to **HA Switching**. The failover is successful when the instance status changes back to **Running**.

The failover operation is complete in about one minute. Then the instance returns to normal.

Note If you have connected to the connection string of the primary node for an instance, you are connecting to a secondary node after a failover and you have no write permissions on the instance. In this case, you must connect to the connection string of the new primary node and obtain read and write permissions. For more information, see [Overview of replica set instance connections](#).

14.4.10.2. Trigger a primary/secondary failover for a sharded cluster instance

Each shard or Configserver of a sharded cluster instance consists of three nodes by default. If a node is faulty, the high availability system of ApsaraDB for MongoDB automatically triggers a primary/secondary failover to ensure the availability of the shard. You can also manually trigger a primary/secondary failover for an ApsaraDB for MongoDB instance in scenarios such as routine disaster recovery drills.

Precautions

ApsaraDB for MongoDB provides connection strings for you to connect to the primary and secondary nodes. The hidden node is invisible to you and only used to ensure high availability. After you log on to the ApsaraDB for MongoDB console or call the `SwitchDBInstanceHA` operation to trigger a primary/secondary failover for a shard of a sharded cluster instance, ApsaraDB for MongoDB switches the roles of the primary and secondary nodes.

Note

- You can trigger a primary/secondary failover only for the shard or Configserve node in the running state.
- Each time you trigger a primary/secondary failover for an instance, the instance may have a transient connection error of about 30 seconds. We recommend that you perform this operation during off-peak hours and ensure that your applications can automatically re-establish a connection.

Procedure

- [Log on to the ApsaraDB for MongoDB console](#).
- In the left-side navigation pane, click **Sharded Cluster Instances**.
- Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
- In the **Shard List** or **ConfigServer List** section, find the node and choose  > **Failover** in the **Actions** column.

 **Note** You can trigger a primary/secondary failover separately for each shard node. The failover operation takes effect only for the current shard node and does not affect other shard nodes of the same sharded cluster instance.

5. In the **Failover** message, click **OK**.

 **Note** The failover operation is complete in about one minute.

14.4.11. Monitoring

This topic describes the performance metrics provided by ApsaraDB for MongoDB to check the status of ApsaraDB for MongoDB instances. You can view these performance metrics in the ApsaraDB for MongoDB console.

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the

Basic Information page appears.

4. In the left-side navigation pane, click **Monitoring Info**.

You can select a time range to query historical performance metrics. The following table describes metric details.

Performance metrics

| Performance metric | Description | Monitoring frequency |
|----------------------------|--|----------------------|
| CPU Utilization Percentage | cpu_usage: the CPU utilization of the instance | Every 300 seconds |
| Memory Usage Percentage | mem_usage: the memory usage of the instance | Every 300 seconds |
| IOPS Usage | The input/output operations per second (IOPS) of the instance. The following items are included: <ul style="list-style-type: none"> ◦ data_iops: the IOPS of the data disk. ◦ log_iops: the IOPS of the log disk. | Every 300 seconds |
| IOPS Usage Percentage | iops_usage: the ratio of the IOPS used by the instance to the maximum IOPS allowed | Every 300 seconds |
| Disk Usage | The total disk space used by the instance. The following items are included: <ul style="list-style-type: none"> ◦ ins_size: the total space used. ◦ data_size: the space used on the data disk. ◦ log_size: the space used on the log disk. | Every 300 seconds |

| Performance metric | Description | Monitoring frequency |
|-----------------------|---|----------------------|
| Disk Usage Percentage | disk_usage: the ratio of the total disk space used by the instance to the maximum disk space that can be used | Every 300 seconds |
| Operation QPS | The queries per second (QPS) of the instance. The following items are included: <ul style="list-style-type: none"> ◦ The number of insert operations. ◦ The number of query operations. ◦ The number of delete operations. ◦ The number of update operations. ◦ The number of getmore operations. ◦ The number of command operations. | Every 300 seconds |
| Connections | current_conn: the number of current connections to the instance | Every 300 seconds |
| Cursors | The number of cursors used by the instance. The following items are included: <ul style="list-style-type: none"> ◦ total_open: the number of cursors that are opened. ◦ timed_out: the number of cursors that timed out. | Every 300 seconds |
| Network Traffic | The network traffic of the instance. The following items are included: <ul style="list-style-type: none"> ◦ bytes_in: the inbound network traffic. ◦ bytes_out: the outbound network traffic. ◦ num_requests: the number of requests that are processed. | Every 300 seconds |

| Performance metric | Description | Monitoring frequency |
|---------------------------------------|---|----------------------|
| Global Lock Waiting Queues | <p>The length of the queues that are waiting for global locks for the instance. The following items are included:</p> <ul style="list-style-type: none"> gl_cq_total: the length of the queue that is waiting for both global read and write locks. gl_cq_readers: the length of the queue that is waiting for global read locks. gl_cq_writers: the length of the queue that is waiting for global write locks. | Every 300 seconds |
| WiredTiger | <p>The cache metrics of the WiredTiger engine used by the instance. The following items are included:</p> <ul style="list-style-type: none"> bytes_read_into_cache: the amount of data that is read into the cache. bytes_written_from_cache: the amount of data that is written from the cache to the disk. maximum_bytes_configured: the size of the maximum available disk space that is configured. | Every 300 seconds |
| Primary/Secondary Replication Latency | <p>repl_lag: the latency in data synchronization between the primary and secondary nodes of the instance</p> | Every 300 seconds |

14.5. Backup and restoration

14.5.1. Configure automatic backup for an ApsaraDB for MongoDB instance

This topic describes how to configure automatic backup for an ApsaraDB for MongoDB instance. ApsaraDB for MongoDB automatically backs up data based on the backup policy you specify.

Procedure

1. [Log on to the ApsaraDB for MongoDB console.](#)
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, click **Backup and Recovery**.
5. In the upper-left corner of the page, click **Backup Settings**. Configure the parameters in the Backup Settings panel.

The following table describes the parameters.

Backup policy parameters

| Parameter | Description |
|-----------------------|--|
| Retention Days | The number of days for which you want to retain backup files. It can only be seven days. |
| Backup Time | The hour at which you want to perform the backup task. |
| Day of Week | The backup cycle. You can select one or more days in a week. |

6. Click **OK**.

14.5.2. Manually back up an ApsaraDB for MongoDB instance

This topic describes how to manually back up an ApsaraDB for MongoDB instance.

Backup methods

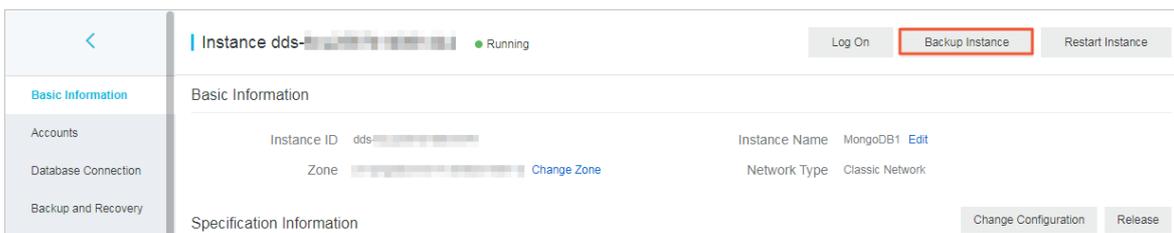
- **Physical backup:** This method backs up physical database files of an ApsaraDB for MongoDB instance. Compared with logical backup, physical backup provides faster data backup and recovery.
- **Logical backup:** The mongodump tool is used to store operation records of databases to a logical backup file. Then, data can be restore data by using the mongorestore tool.

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the

Basic Information page appears.

4. In the upper-right corner of the page, click **Backup Instance**.



5. In the dialog box that appears, set **Backup Method** and click **OK**.

14.5.3. Restore data to the current ApsaraDB for MongoDB instance

This topic describes how to restore data to the current ApsaraDB for MongoDB instance. This helps minimize the data loss caused by incorrect operations.

Prerequisites

The instance is a replica set instance with three nodes.

Background information

- The time required to restore data to your current instance varies depending on factors such as the data volume, task queue status, and network conditions. When the status of the instance changes to **Running**, the restoration is complete.
- If you restore data to your current instance, all existing data is overwritten and cannot be restored.

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. On the **Replica Set Instances** page, find the instance and then click the instance ID or choose  **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
3. In the left-side navigation pane, click **Backup and Recovery**.
4. On the **Backup and Recovery** page that appears, find the backup set and choose  **Data Recovery** in the **Actions** column.

 **Note** If you have upgraded the database version, you cannot use the backup files of the earlier database version to restore data.

5. In the **Roll Back Instance** message, click **OK**.

 **Note** The instance status becomes **Restoring from Backup** after you click **OK**. You can click **Refresh** in the upper-right corner of the **Backup and Recovery** page to update the instance status. The restoration is complete when the instance status changes to **Running**.

14.6. Database connections

14.6.1. Modify a public or internal endpoint of an ApsaraDB for MongoDB instance

This topic describes how to modify a public or internal endpoint of an ApsaraDB for MongoDB instance in the ApsaraDB for MongoDB console.

Limits

| Architecture | Limit |
|--------------------------|---|
| Replica set instance | You can modify the public and internal endpoints of both primary and secondary nodes. |
| Sharded cluster instance | You can modify only the public and internal endpoints of mongos nodes. |

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, click **Database Connections**.

- In the upper-right corner of the **Internal Connections** or **Public Connections** section, click **Update Connection String**.
- In the panel that appears, enter a new endpoint.

Replica set instance

Update Connection String
✕

*** Node**

Primary(871)
▼

Current Connection String

dds-XXXXXXXXXX.mongodb.rds.thirteenth-inc.com

*** New Connection String**

newconnection123
.mongodb.rds.thirteenth-inc.com

Submit
Close

Sharded cluster instance

Update Connection String
✕

*** Node**

s-q8ia34fec7bfc604
▼

Note: You can only modify the endpoints of Mongos nodes.

Current Connection String

s-q8ia34fec7bfc604.mongodb.rds.thirteenth-inc.com

*** New Connection String**

newconnection123
.mongodb.rds.thirteenth-inc.com

Submit
Close

? **Note**

- You can modify only the prefix of the endpoint.
- The endpoint must start with a lowercase letter and end with a lowercase letter or a digit. It must be 8 to 64 characters in length and can contain lowercase letters, digits, and hyphens (-).

- Click **Submit**.

What's next

After you modify the public or internal endpoint, you must connect a client or an application to your ApsaraDB for MongoDB instance by using the new endpoint.

14.6.2. Use DMS to log on to an ApsaraDB for MongoDB instance

You can use DMS to connect to an ApsaraDB for MongoDB instance.

Prerequisites

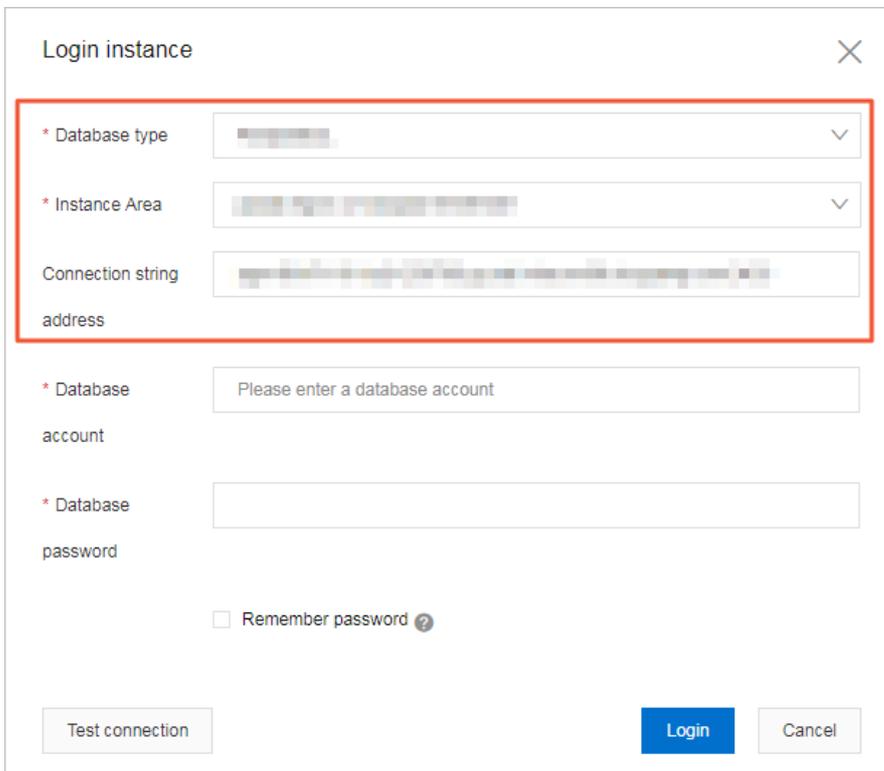
The IP address whitelist is configured. For more information about how to configure the IP address whitelist, see [Configure a whitelist for an ApsaraDB for MongoDB instance](#).

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. Click **Log On** in the upper-right corner of the page.

 **Note** For a sharded cluster instance, you must also select the Mongos node.

5. In the **Login instance** dialog box of the **DMS** console, check values of **Database type**, **Instance Area**, and **Connection string address**. If the information is correct, enter **Database account** and **Database password**, as shown in the following figure.



The screenshot shows a 'Login instance' dialog box with the following fields and controls:

- Database type**: A dropdown menu with a red box around it.
- Instance Area**: A dropdown menu with a red box around it.
- Connection string address**: A text input field with a red box around it.
- Database account**: A text input field with the placeholder text 'Please enter a database account'.
- Database password**: A text input field.
- Remember password**: A checkbox with a help icon.
- Buttons**: 'Test connection', 'Login', and 'Cancel'.

| Parameter | Description |
|----------------------------------|---|
| Database type | The engine of the database. By default, the engine of the database to be connected is displayed. |
| Instance Area | The region where the instance is deployed. By default, the region of the current instance is displayed. |
| Connection string address | The endpoint of the instance. By default, the endpoint of the current instance is displayed. |
| Database account | The account of the database to be connected. |
| Database password | The password of the account used to connect to the database. |

6. Click **Login**.

 **Note** If you want your web browser to remember the password, select **Remember password** before you click **Login**.

14.6.3. Use the mongo shell to connect to an ApsaraDB for MongoDB instance

This topic describes how to use the mongo shell to connect to an ApsaraDB for MongoDB instance. The mongo shell is a database management tool provided by ApsaraDB for MongoDB. You can install it on your client or in an Elastic Compute Service (ECS) instance.

Prerequisites

- The version of the mongo shell is the same as your ApsaraDB for MongoDB instance. This ensures successful authentication. For more information about the installation procedure, see [Install MongoDB](#). Choose the version in the upper-left corner of the page based on your client version.
- The IP address of your client is added to a whitelist of the ApsaraDB for MongoDB instance. For more information, see [Configure a whitelist for an ApsaraDB for MongoDB instance](#).

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, click **Database Connections** to view connection strings.

 **Note**

- Replica set instances: View the connection string or connection string URI of a node.
- Sharded cluster instances: View the connection string or connection string URI of a mongos node.

For more information about connection strings, see [Overview of replica set instance connections](#) or [Overview of sharded cluster instance connections](#).

5. Connect to the ApsaraDB for MongoDB instance from your client or ECS instance where the mongo shell is installed.

Replica set instances:

- Connection string of a node

During routine tests, you can directly connect to a primary or secondary node. Note that if the primary node fails, the system automatically switches to the secondary node, and the roles of connected nodes change. This affects the read and write operations of your application.

Command syntax:

```
mongo --host <host> -u <username> -p --authenticationDatabase <database>
```

 **Note**

- <host>: the connection string of the primary or secondary node.
- <username>: the username used to log on to a database of the instance. The initial username is **root**.
- <database>: the name of the database corresponding to the username if authentication is enabled. If the database username is root, enter admin as the database name.

Example:

```
mongo --host dds-bp*****.mongodb.rds.aliyuncs.com:3717 -u root -p --authenticationDatabase admin
```

When **Enter password:** is displayed, enter the password of the database user and press the Enter key. If you forget the password of the root user, you can reset it. For more information, see [Reset the password for an ApsaraDB for MongoDB instance](#).

 **Note** The password characters are not displayed when you enter the password.

- HA connection (recommended): You can use a connection string URI to connect to both the primary and secondary nodes of a replica set instance. This ensures that your application is always connected to the primary node and the read and write operations of your application are not affected even if the roles of the primary and secondary nodes are switched.

Command syntax:

```
mongo "<ConnectionStringURI>"
```

 **Note**

- The connection string URI must be enclosed in a pair of double quotation marks ("").
- <ConnectionStringURI>: the connection string URI of the instance.
You must replace ******** in the connection string URI with the database password. For more information about how to set a database password, see [Reset the password for an ApsaraDB for MongoDB instance](#).

Example:

```
mongo "mongodb://root:****@dds-*****.mongodb.rds.intra.env17e.shuguang.com:3717,dds-*****.mongodb.rds.intra.env17e.shuguang.com:3717/admin?replicaSet=mgset-*****"
```

Sharded cluster instances:

- Connection string of a mongos node

Command syntax:

```
mongo --host <mongos_host> -u <username> -p --authenticationDatabase <database>
```

 **Note**

- <mongos_host>: the connection string of a mongos node in the sharded cluster instance.
- <username>: the username used to log on to a database of the instance. The initial username is root.
- <database>: the name of the database corresponding to the username if authentication is enabled. If the database username is root, enter admin as the database name.

Example:

```
mongo --host s-bp*****.mongodb.rds.aliyuncs.com:3717 -u root -p --authenticationDatabase admin
```

- HA connection (recommended): You can use a connection string URI to connect to a database. If one mongos node fails, another mongos node takes over business to ensure the high availability of the connection.

Command syntax:

```
mongo "<ConnectionStringURI>"
```

 **Note**

- The connection string URI must be enclosed in a pair of double quotation marks ("").
 - <ConnectionStringURI>: the connection string URI of the instance.
- You must replace **** in the connection string URI with the database password. For more information about how to set a database password, see [Reset the password for an ApsaraDB for MongoDB instance](#).

Example:

```
mongo "mongodb://root:****@s-*****.mongodb.rds.intra.env17e.shuguang.com:3717,s-*****.mongodb.rds.intra.env17e.shuguang.com:3717/admin"
```

14.6.4. Apply for a public endpoint for an ApsaraDB for MongoDB instance

This topic describes how to apply for a public endpoint for an ApsaraDB for MongoDB instance when you want to connect to this instance over the Internet.

Context

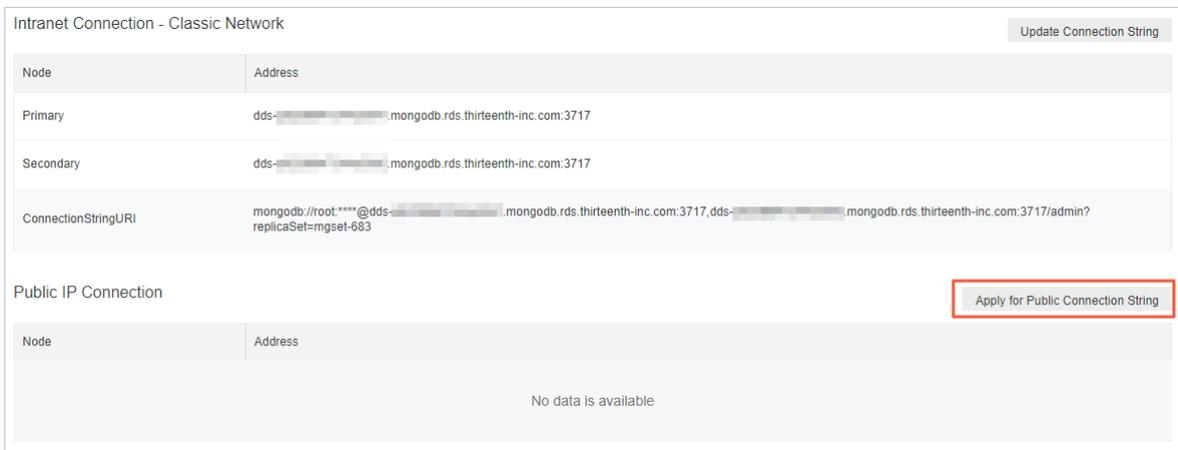
The following table describes the Virtual Private Cloud (VPC) and classic network endpoints supported by ApsaraDB for MongoDB.

| Type | Description |
|--------------------------|--|
| VPC endpoint | <ul style="list-style-type: none"> • A VPC is an isolated network with higher security and performance than the classic network. • By default, ApsaraDB for MongoDB provides endpoints on a VPC. |
| Classic network endpoint | Cloud services on the classic network are not isolated. Unauthorized access can be blocked only by using security groups or whitelists. |

| Type | Description |
|-----------------|--|
| Public endpoint | <ul style="list-style-type: none">Security risks may arise if you connect to an ApsaraDB for MongoDB instance over the Internet. Therefore, ApsaraDB for MongoDB does not provide public endpoints.If you want to connect to an ApsaraDB for MongoDB instance from a device outside of Alibaba Cloud (such as an on-premises device), you must apply for a public endpoint. |

Apply for a public endpoint for a replica set instance

- Log on to the [ApsaraDB for MongoDB console](#).
- On the **Replica Set Instances** page, find the instance and then click the instance ID or choose  **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
- In the left-side navigation pane, click **Database Connections**.
- In the upper-right corner of the **Public Connections** section, click **Apply for Public Connection String**.



- In the **Apply for Public Connection String** message, click **OK**.

 **Note** If you want to connect to an ApsaraDB for MongoDB instance by using a public endpoint, you must add the public IP address of your client to a whitelist of this instance. For more information, see [Configure a whitelist for an ApsaraDB for MongoDB instance](#).

After the application is complete, the replica set instance generates new endpoints for both the primary and secondary nodes and the corresponding connection string URI. For more information, see [Overview of replica set instance connections](#).

Apply for a public endpoint for a sharded cluster instance

- Log on to the [ApsaraDB for MongoDB console](#).
- In the left-side navigation pane, click **Sharded Cluster Instances**.
- Find the instance and then click the instance ID or choose  **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
- In the left-side navigation pane, click **Database Connections**.
- In the upper-right corner of the **Public Connections** section, click **Apply for Public Connection String**.

| Public IP Connection | | | | |
|----------------------|-----------|------|---------|---------|
| ID | Node Type | Node | Address | Actions |
| No data is available | | | | |

6. In the panel that appears, specify **Node Type** and **Node ID**, and click **OK**.

Apply for Public Connection String ✕

• **Node Type**

Mongos ▼

• **Node ID**

Select ▼

OK
Cancel

| Parameter | Value | Description |
|------------------|--|---|
| Node Type | Mongos | The mongos node. You can apply for public endpoints only for mongos nodes. Your applications are connected to mongos nodes in most cases. |
| Node ID | The ID of the component for which you want to apply for a public endpoint. | None |

Note You can repeat this step to apply for public endpoints for other mongos nodes as your needs change. You can apply for a new public endpoint only after the current one is created.

References

- To ensure data security, we recommend that you release a public endpoint if you no longer need it. For more information, see [Release a public connection string](#).
- Before you connect to a replica set instance over the Internet, we recommend that you enable SSL encryption. For more information, see [Use the mongo shell to connect to an ApsaraDB for MongoDB database in SSL encryption mode](#).

14.6.5. Release a public endpoint

To ensure data security, you can release a public endpoint that is no longer needed in the console.

Precautions

- You can release one or more public endpoints of the mongos nodes for a sharded cluster instance.
- After the public endpoint is released for an instance or node, you cannot connect to the instance or node by

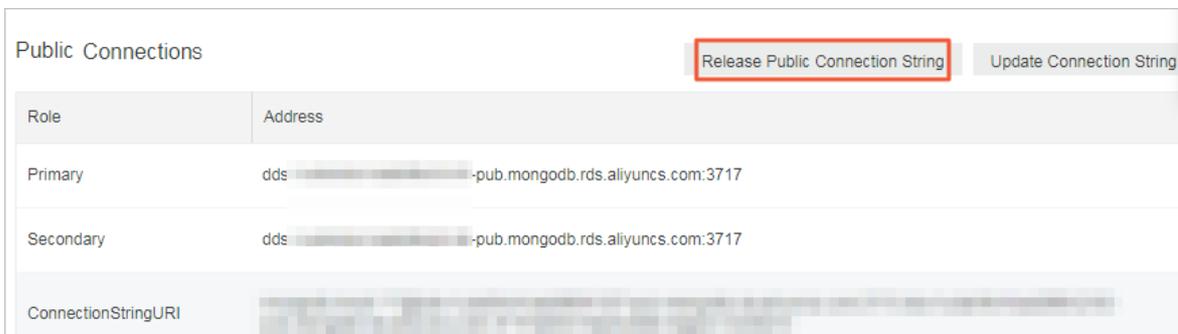
using the original public endpoint.

- After the public endpoint is released, we recommend that you delete the corresponding public IP address from the whitelist to ensure data security. For more information, see [Configure a whitelist](#).

Release a public endpoint for a replica set instance

Note After the public endpoint of a replica set instance is released, the public endpoints of the primary and secondary nodes are released.

1. [Log on to the ApsaraDB for MongoDB console](#).
2. On the **Replica Set Instances** page, find the instance and then click the instance ID or choose  **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
3. In the left-side navigation pane, click **Database Connections**.
4. In the upper-right corner of the **Public Connections** section, click **Release Public Connection String**.



5. In the message that appears, click **OK**.

Release a public endpoint for a sharded cluster instance

You can release one or more public endpoints of the mongos, shard, and Configserver nodes for a sharded cluster instance.

Note After the public endpoint of a shard or Configserver node is released, the public endpoints of the primary and secondary nodes are released.

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, click **Database Connections**.
5. In the **Public Connections** section, find the mongos, shard, or Configserver node for which you want to release the public endpoint.
6. In the **Actions** column corresponding to the node, click **Release**.

| Public IP Connection | | | | Apply for Public Connection String | Update Connection String |
|----------------------|-----------|------|------------|------------------------------------|--------------------------|
| ID | Node Type | Node | Address | Actions | |
| s- [redacted] | Mongos | - | [redacted] | Release | |
| ConnectionStringURI | Mongos | - | [redacted] | | |

Note You can repeat this step to release the public endpoints of other nodes as your needs change. To release the public endpoint of the next node, you must wait until the public endpoint of the current node is released or the status of the current node becomes **Running**.

7. In the message that appears, click **OK**.

14.6.6. Overview of replica set instance connections

ApsaraDB for MongoDB supports both connection strings and connection string URIs. You can use a connection string to connect to the primary or secondary node, and use a connection string URI to connect to both of them. For high availability, we recommend that you use connection string URIs to connect your application to both primary and secondary nodes. This topic provides an overview of replica set instance connections.

Prerequisites

A whitelist is configured for the instance. For more information, see [Configure a whitelist for an ApsaraDB for MongoDB instance](#).

View connection strings

1. [Log on to the ApsaraDB for MongoDB console](#).
2. On the **Replica Set Instances** page, find the instance and then click the instance ID or choose **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
3. In the left-side navigation pane, click **Database Connections** to view connection strings.

| Intranet Connection - Classic Network | | Update Connection String |
|---------------------------------------|---|--------------------------|
| Node | Address | |
| Primary | dds-[redacted].mongodb.rds.thirteenth-inc.com:3717 | |
| Secondary | dds-[redacted].mongodb.rds.thirteenth-inc.com:3717 | |
| ConnectionStringURI | mongodb://root:***@dds-[redacted].mongodb.rds.thirteenth-inc.com:3717,dds-[redacted].mongodb.rds.thirteenth-inc.com:3717/admin?replicaSet=mgset-683 | |

Description of connection strings

| Item | Description |
|-----------------|---|
| Connection type | <ul style="list-style-type: none"> • Internal Connections - Classic Network: Cloud services on the classic network are not isolated. Unauthorized access can be blocked only by security groups or whitelists. • Internal Connections - VPC: A virtual private cloud (VPC) is an isolated network with higher security and performance than the classic network. By default, ApsaraDB for MongoDB provides VPC endpoints. |

| Item | Description |
|-----------------------|---|
| Role | <ul style="list-style-type: none"> • Primary: the primary node in the replica set instance. If you connect to this node, you can perform read and write operations on the databases of the replica set instance. • Secondary: the secondary node in the replica set instance. If you connect to this node, you can perform only read operations on the databases of the replica set instance. • Connection String URI: ApsaraDB for MongoDB allows you to use a connection string URI to connect to a replica set instance to achieve load balancing and high availability. |
| Connection string | <p>The connection string of a primary or secondary node is in the following format:</p> <pre><host>:<port></pre> <ul style="list-style-type: none"> • <host>: the endpoint used to connect to the instance. • <port>: the port used to connect to the instance. |
| Connection string URI | <p>A connection string URI is in the following format:</p> <pre>mongodb://[username:password@]host1[:port1][,host2[:port2],...[,hostN[:portN]]][/[database][?options]]</pre> <ul style="list-style-type: none"> • mongodb://: the prefix of a connection string URI. • username:password@: the username and password used to log on to a database of the replica set instance. You must separate them with a colon (:). • hostX:portX: the endpoint and port of a node in the replica set instance. • /database: the name of the database corresponding to the username if authentication is enabled. • ?options: additional connection options. <p>Note If your application is in a production environment, we recommend that you use a connection string URI to connect to the instance. This way, when a node fails, the read and write operations of your application are not affected as a result of the failover.</p> |

Related information

- [Connect to a replica set instance by using the mongo shell](#)

14.6.7. Overview of sharded cluster instance connections

ApsaraDB for MongoDB supports both connection strings and connection string URIs. You can use a connection string to connect to a single mongos node, and use a connection string URI to connect to multiple mongos nodes. For high availability, we recommend that you use connection string URIs to connect your application to multiple mongos nodes. This topic provides an overview of sharded cluster instance connections.

View connection strings

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, click **Database Connections** to view connection strings.

| Intranet Connection - Classic Network | | | | | Update Connection String |
|---------------------------------------|-----------|------|---|---------|--------------------------|
| ID | Node Type | Node | Address | Actions | |
| s- [redacted] | Mongos | - | s- [redacted].mongodb.rds.thirteenth-inc.com:3717 | Release | |
| s- [redacted] | Mongos | - | s- [redacted].mongodb.rds.thirteenth-inc.com:3717 | Release | |
| ConnectionStringURI | Mongos | - | mongodb://root:****@s- [redacted].inc.com:3717/admin | Release | |

| Public IP Connection | | | | | Apply for Public Connection String | Update Connection String |
|----------------------|-----------|------|---|---------|------------------------------------|--------------------------|
| ID | Node Type | Node | Address | Actions | | |
| s- [redacted] | Mongos | - | s- [redacted]-pub.mongodb.rds.thirteenth-inc.com:3717 | Release | | |
| s- [redacted] | Mongos | - | s- [redacted]-pub.mongodb.rds.thirteenth-inc.com:3717 | Release | | |
| ConnectionStringURI | Mongos | - | mongodb://root:****@s- [redacted]-pub.mongodb.rds.thirteenth-inc.com:3717/s- [redacted]-pub.mongodb.rds.thirteenth-inc.com:3717/admin | | | |

Description of connection strings

| Item | Description |
|-----------------|---|
| Connection type | <ul style="list-style-type: none"> Internal Connections - Classic Network: Cloud services on the classic network are not isolated. Unauthorized access can be blocked only by security groups or whitelists. Internal Connections - VPC: A virtual private cloud (VPC) is an isolated network with higher security and performance than the classic network. By default, ApsaraDB for MongoDB provides VPC endpoints. Public Connections: Security risks may arise if you connect to a sharded cluster instance over the Internet. Therefore, ApsaraDB for MongoDB does not provide public endpoints. If you want to connect to an ApsaraDB for MongoDB instance from a device outside of Alibaba Cloud (such as an on-premises device), you must apply for a public endpoint. For more information, see Apply for a public endpoint for an ApsaraDB for MongoDB instance. |
| Mongos ID | <p>The connection string of a mongos node is in the following format:</p> <pre><host>:<port></pre> <ul style="list-style-type: none"> <host>: the endpoint used to connect to the instance. <port>: the port used to connect to the instance. <div style="background-color: #e0f2f1; padding: 5px;"> <p> Note During routine tests, you can use a connection string to directly connect to a mongos node.</p> </div> |

| Item | Description |
|-----------------------|--|
| Connection string URI | <p>A connection string URI is in the following format:</p> <pre>mongodb://[username:password@]host1[:port1][,host2[:port2],...[,hostN[:portN]]][/[database][?options]]</pre> <ul style="list-style-type: none"> • <code>mongodb://</code>: the prefix of a connection string URI. • <code>username:password@</code>: the username and password used to log on to a database of the sharded cluster instance. You must separate them with a colon (:). • <code>hostX:portX</code>: the endpoint and port of a mongos node in the sharded cluster instance. • <code>/database</code>: the name of the database corresponding to the username if authentication is enabled. • <code>?options</code>: additional connection options. <p>Note If your application is in a production environment, we recommend that you use a connection string URI to connect to the instance. Then, your client can automatically distribute your requests to multiple mongos nodes to balance loads. If a mongos node fails, your client automatically redirects requests to other mongos nodes in the normal state.</p> |

Log on to a database of an ApsaraDB for MongoDB instance

1. Obtain the **connection string** and the following information:
 - The username used to log on to the database. The initial username is root.
 - The password of the database user. If you forget the password of the root user, you can reset it. For more information, see [Reset the password for an ApsaraDB for MongoDB instance](#).
 - The name of the database corresponding to the username if authentication is enabled. If the database username is root, enter admin as the database name.
2. Log on to the database.
 - [Use DMS to log on to a replica set instance of ApsaraDB for MongoDB](#)
 - [Connect to a replica set instance by using the mongo shell](#)

14.7. Data security

14.7.1. Configure a whitelist for an ApsaraDB for MongoDB instance

This topic describes how to configure a whitelist for an ApsaraDB for MongoDB instance. Before you use an ApsaraDB for MongoDB instance, you must add the IP addresses or Classless Inter-Domain Routing (CIDR) blocks that you use for database access to a whitelist of this instance. This improves database security and stability. Proper configuration of whitelists can enhance access security of ApsaraDB for MongoDB. We recommend that you maintain the whitelists on a regular basis.

Context

The system creates a default whitelist for each instance. This whitelist can be modified or cleared, but it cannot be deleted. After an ApsaraDB for MongoDB instance is created, the system automatically adds the IP address 127.0.0.1 to the **default** whitelist of this instance. The IP address 0.0.0.0/0 indicates that all IP addresses are allowed to access this instance. You must add the IP addresses or CIDR blocks that you allow to access this ApsaraDB for MongoDB instance.

Procedure

1. [Log on to the ApsaraDB for MongoDB console.](#)
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, choose **Data Security > Whitelist Settings**.
5. You can manually configure a whitelist or import ECS Internal IP addresses to the whitelist.

Manually modify a whitelist

- i. Find the whitelist you want to modify and choose  > **Manually Modify** in the **Actions** column.
- ii. Enter IP addresses or CIDR blocks.

Note

- Separate multiple IP addresses with commas (,). You can add a maximum of 1,000 different IP addresses to a whitelist. Supported formats are IP addresses such as 0.0.0.0/0 and 10.23.12.24, or CIDR blocks such as 10.23.12.24/24. /24 indicates the length of the IP address prefix in the CIDR block. An IP address prefix can contain 1 to 32 bits.
- If the IP whitelist is empty or only contains 0.0.0.0/0, all devices are granted access. This is risky for your ApsaraDB for MongoDB instance. We recommend that you add only the IP addresses or CIDR blocks of your own web servers to the whitelist.

- iii. Click **OK**.

Load IP addresses of ECS instances

- i. Find the whitelist, and choose  > **Import ECS Intranet IP** in the **Actions** column.
- ii. From the displayed internal IP addresses of ECS instances that belong to the current account, select the IP addresses and click  to add them to the whitelist.
- iii. Click **OK**.

14.7.2. Create or delete a whitelist

This topic describes how to create or delete whitelists that consist of the IP addresses allowed to access specific databases.

Context

If your business involves multiple applications and you need to add a whitelist for each of them, you can sort the IP addresses into different whitelists.

Create a whitelist

1. [Log on to the ApsaraDB for MongoDB console.](#)
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, choose **Data Security > Whitelist Settings**.

5. Click **Create Whitelist** in the upper-left corner of the page.
6. In the panel that appears, set **Group Name** and **IP White List** and click **OK**.

 **Note**

- **Group Name:** The name must be 2 to 32 characters in length and can contain lowercase letters, digits, and underscores (_). It must start with a lowercase letter and end with a lowercase letter or digit.
- **IP White List :**
 - Separate multiple IP addresses with commas (,). You can add a maximum of 1,000 different IP addresses to a whitelist. A whitelist can include IP addresses such as 0.0.0.0/0 and 10.23.12.24, or CIDR blocks such as 10.23.12.24/24. /24 indicates that the prefix of the CIDR block is 24-bit long. You can replace 24 with a value within the range of 1 to 32.
 - If the whitelist is empty or contains only 0.0.0.0/0, all devices are granted access. This poses risks to your ApsaraDB for MongoDB instance. We recommend that you add only the IP addresses or CIDR blocks of your own web servers to the whitelist.

Delete a whitelist

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, choose **Data Security** > **Whitelist Settings**.
5. Find the whitelist that you want to delete and choose  > **Delete Whitelist Group** in the **Actions** column.

 **Note** You cannot delete the **default** whitelist.

6. In the Delete Whitelist Group message, click **OK**.

14.7.3. Audit logs

This topic describes audit logs provided in the ApsaraDB for MongoDB console. You can query the statement execution logs, operations logs, and error logs of an ApsaraDB for MongoDB instance to locate and analyze faults.

Context

The audit log feature records all operations that a client performs on a connected database. This feature provides references for you to perform fault analysis, behavior analysis, and security auditing because you can obtain the operation execution details from the audit logs. Audit logs are essential in the regulatory operations of Finance Cloud and other core business scenarios.

 **Note** Audit logs are stored for seven days, after which they are deleted.

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the

Basic Information page appears.

4. In the left-side navigation pane, choose **Data Security > Audit Logs**.
5. Click **Enable Audit Log** in the upper-left corner. In the Enable Audit message, click **OK**.

Result

On the **Audit Log** page, specify the time range, database name, database user, and keyword to query audit logs. You can also perform the following operations:

- **Export File**: exports an audit log file.
- **File List**: displays a list of audit logs.
- **Disable Audit Log**: stops the collection of information on database operations and deletes the saved audit logs.

14.7.4. Configure SSL encryption for an ApsaraDB for MongoDB instance

This topic describes how to enhance link security by enabling Secure Sockets Layer (SSL) encryption and installing SSL CA certificates on your application services. The SSL encryption feature encrypts network connections at the transport layer to improve data security and ensure data integrity during communication.

Prerequisites

- The instance is a replica set instance.
- The MongoDB version of the instance is 3.4, 4.0, or 4.2.

Notes

When you enable or disable SSL encryption or update SSL CA certificates for an instance, the instance is restarted. Plan your operations in advance and make sure that your application is configured to reconnect to the instance after it is disconnected.

 **Note** When an instance is restarted, all its nodes are restarted in turn and each node has a transient connection error of about 30 seconds. If the instance contains more than 10,000 collections, the transient connection error last longer.

Precautions

- You can download SSL CA certificate files only from the ApsaraDB for MongoDB console.
- After you enable SSL encryption for an instance, the CPU utilization of the instance is significantly increased. We recommend that you enable SSL encryption only when necessary. For example, you can enable SSL encryption when you connect to an ApsaraDB for MongoDB instance over the Internet.

 **Note** Internal network connections are more secure than Internet connections and do not need SSL encryption.

- After you enable SSL encryption for an instance, both SSL and non-SSL connections are supported.

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  **> Manage** in the **Actions** column. Then, the

Basic Information page appears.

4. In the left-side navigation pane, choose **Data Security > SSL**.
5. Perform one of the following operations.

Note When you enable or disable SSL encryption or update SSL CA certificates for an instance, the instance is restarted. Plan your operations in advance and make sure that your applications can automatically re-establish a connection.

| Operation | Prerequisite | Procedure |
|-------------------------------------|--|--|
| Enable SSL encryption | The SSL encryption status is Disabled . | Turn on SSL Status . In the message that appears, click OK . |
| Update an SSL CA certificate | The SSL encryption status is Enabled . | Click Update Certificate . In the message that appears, click OK . |
| Download an SSL CA certificate file | The SSL encryption status is Enabled . | Click Download Certificate to download an SSL CA certificate file to your computer. |
| Disable SSL encryption | The SSL encryption status is Enabled . | Turn off SSL Status . In the message that appears, click OK . |

14.7.5. Configure TDE for an ApsaraDB for MongoDB instance

This topic describes how to configure Transparent Data Encryption (TDE) for an ApsaraDB for MongoDB instance. Before data files are written to disks, TDE encrypts the data files. When data files are loaded from disks to the memory, TDE decrypts the data files. TDE does not increase the sizes of data files. When you use TDE, you do not need to modify your application that uses the ApsaraDB for MongoDB instance. To enhance data security, you can enable the TDE feature for an instance in the ApsaraDB for MongoDB console.

Prerequisites

The MongoDB version of the instance is 4.0 or 4.2.

Note Before you enable TDE, you can create a MongoDB 4.0 or 4.2 instance to test the compatibility between your application and the database version. You can release the instance after the test is complete.

Notes

- When you enable TDE, your instance is restarted, and your application is disconnected from the instance. We recommend that you enable TDE during off-peak hours and make sure that your application can reconnect to the instance after it is disconnected.
- TDE increases the CPU utilization of your instance.

Precautions

- You cannot disable TDE after it is enabled.
- You can enable TDE for an instance and disable encryption for a collection.

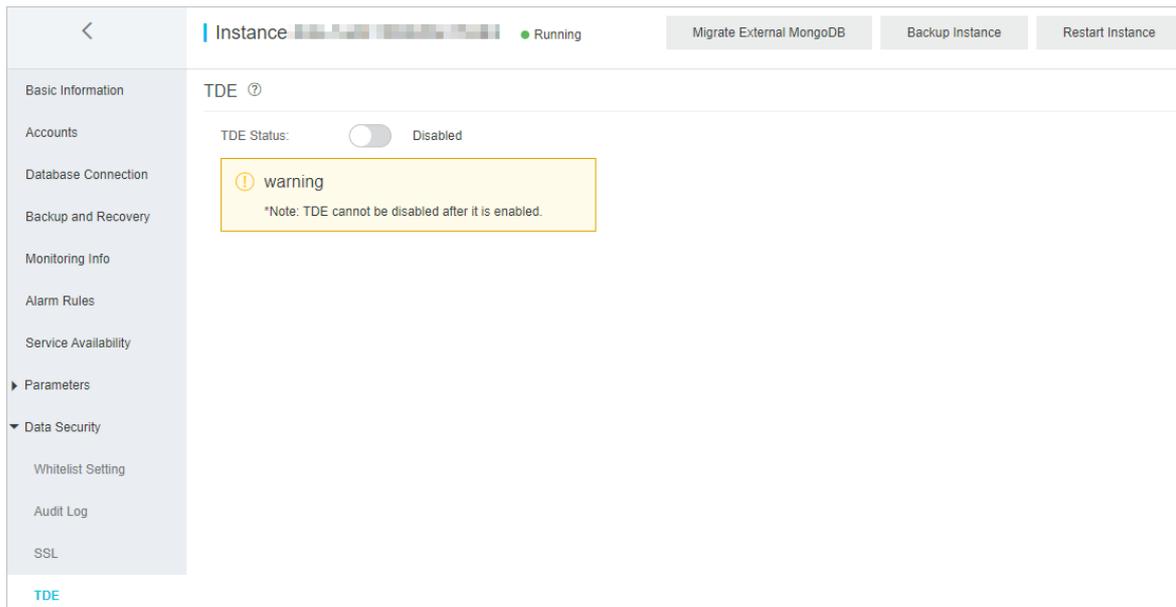
Note In special business scenarios, you can choose not to encrypt a collection when you create it. For more information, see [Disable encryption for a specified collection](#).

- After you enable TDE, only new collections are encrypted. Existing collections are not encrypted.
- Key Management Service (KMS) generates and manages the keys used by TDE. ApsaraDB for MongoDB does not

provide keys or certificates required for encryption.

Procedure

1. [Log on to the ApsaraDB for MongoDB console.](#)
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, choose **Data Security > TDE**.
5. Turn on **TDE Status** to enable TDE.



6. In the **Restart Instance** dialog box, click **OK**.
The instance status changes to **Modifying TDE**. After the status changes to **Running**, TDE is enabled.

Disable encryption for a specified collection

After you enable TDE, all new collections are encrypted. When you create a collection, you can perform the following steps to disable encryption for the collection:

1. Connect to a replica set instance by using the mongo shell. For more information, see [Connect to a replica set instance by using the mongo shell](#).
2. Run the following command to create a collection and disable the encryption feature:

```
db.createCollection("<collection_name>",{ storageEngine: { wiredTiger: { configString: "encryption=(name=none)" } } })
```

 **Note** <collection_name>: the name of the collection.

Example:

```
db.createCollection("customer",{ storageEngine: { wiredTiger: { configString: "encryption=(name=none)" } } })
```

14.7.6. Use the mongo shell to connect to an ApsaraDB for MongoDB database in SSL encryption mode

This topic describes how to use the mongo shell to connect to an ApsaraDB for MongoDB database in Secure Sockets Layer (SSL) encryption mode. SSL encryption can encrypt network connections at the transport layer to improve data security and ensure data integrity.

Prerequisites

- The instance is a replica set instance, and the database version of the instance is 3.4, 4.0, or 4.2.
- The mongo shell of the required version is installed on the local server or ECS instance from which you want to connect to the database. For more information about the installation procedure, visit [MongoDB official documentation](#).
- SSL encryption is enabled for the instance. For more information, see [Configure SSL encryption](#).
- The IP address of the local server or the ECS instance is added to a whitelist of the ApsaraDB for MongoDB instance. For more information, see [Configure a whitelist for an ApsaraDB for MongoDB instance](#).

Precautions

After you enable SSL encryption for an instance, the CPU utilization of the instance is significantly increased. We recommend that you enable SSL encryption only when necessary.

Procedure

A local server with a Linux operating system is used in the following example.

1. Download an SSL CA certificate package. For more information, see [Configure SSL encryption](#).
2. Decompress the package and upload the certificate files to the local server or the ECS instance where the mongo shell is installed.

 **Note** In this example, the `.pem` file is uploaded to the `/root/sslcafile/` directory of the local server.

3. On the local server or the ECS instance, run the following command to connect to a database of the ApsaraDB for MongoDB instance:

```
mongo --host <host> -u <username> -p --authenticationDatabase <database> --ssl --sslCAFile <sslCAFile_path> --ssl AllowInvalidHostnames
```

Note

- `<host>`: the connection string of the primary or secondary node for a replica set instance or of mongos node for a sharded cluster instance. For more information, see [Overview of replica set instance connections](#) or [Overview of sharded cluster instance connections](#). If you want to connect to a database of the ApsaraDB for MongoDB instance over an internal network, make sure that the ApsaraDB for MongoDB instance has the same network type as the ECS instance. If the network type is VPC, make sure that the two instances are in the same VPC.
- `<username>`: the username you use to log on to a database of the ApsaraDB for MongoDB instance. The default username is root.
- `<database>`: the name of database corresponding to the username if authentication is enabled. If the database username is root, enter admin.
- `<sslCAFile_path>`: the path of the SSL CA certificate files.

Example:

```
mongo --host dds-bpxxxxxxx-pub.mongodb.rds.aliyuncs.com:3717 -u root -p --authenticationDatabase admin --ssl --sslCAFile /root/sslcafile/ApsaraDB-CA-Chain.pem --sslAllowInvalidHostnames
```

4. When **Enter password:** is displayed, enter the password of the database user and press Enter.

 **Note**

- The password characters are not displayed when you enter the password.
- If you forget the password of the root user, you can reset it. For more information, see [Reset the password for an ApsaraDB for MongoDB instance](#).

14.8. Zone-disaster recovery

14.8.1. Create a dual-zone replica set instance

This topic describes how to create a dual-zone replica set instance. ApsaraDB for MongoDB provides a zone-disaster recovery solution to ensure the reliability and availability of your replica set instance. This solution deploys the three nodes of a replica set instance across two different zones within one region. The components in these zones exchange data over an internal network. When one of the two zones becomes unavailable due to unexpected events such as a power or network failure, the high-availability (HA) system switches services over to another zone.

Deployment policies

The primary, secondary, and hidden nodes of a replica set instance are deployed in two different zones within one region.

Procedure

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances**.
3. On the **Replica Set Instances** page, click **Create Instance**.
4. On the **Create ApsaraDB for MongoDB Instance** page, configure parameters.

 **Note**

- For the **Zone** parameter, you must select dual zones, such as the amtest17001-a and amtest17001-b zones of the cn-qingdao-env11e-MAZ1 region.
- For more information, see [Create a replica set instance](#).

5. Click **Submit**.

14.8.2. Create a dual-zone sharded cluster instance

This topic describes how to create a dual-zone sharded cluster instance. ApsaraDB for MongoDB provides a zone-disaster recovery solution to ensure the reliability and availability of your sharded cluster instance. This solution deploys the components of a sharded cluster instance across two different zones within one region. The components in these zones exchange data over an internal network. When one zone becomes unavailable due to unexpected events such as a power or network failure, the high availability (HA) system automatically switches business over to the other zone.

Deployment policies

The components of a sharded cluster instance are deployed across two different zones within one region.

- Mongos nodes are evenly deployed across all data centers. At least two mongos nodes are deployed at a time, with each in one zone. Each new mongos node added later is deployed to one of the zones in turn.
- The primary, secondary, and hidden nodes in each shard node are not deployed to the two zones in sequence. The deployment of these nodes may change when manual switchover or HA failover between primary and secondary nodes is triggered.

Procedure

1. [Log on to the ApsaraDB for MongoDB console.](#)
2. In the left-side navigation pane, click **Sharded Cluster Instances**.
3. On the **Sharded Cluster Instances** page, click **Create Instance**.
4. On the **Create ApsaraDB for MongoDB Sharded Cluster Instance** page, configure parameters.

Note

- For the **Zone** parameter, you must select dual zones, such as the `amttest17001-a` and `amttest17001-b` zones of the `cn-qingdao-env11e-MAZ1` region.
- For more information, see [Create a sharded cluster instance](#).

5. Click **Submit**.

14.9. CloudDBA

14.9.1. Performance trends

This topic describes how to view performance trends in specific ranges, compare performance trends, and customize charts to view performance trends on your ApsaraDB for MongoDB instances.

Go to the Performance page

1. [Log on to the ApsaraDB for MongoDB console.](#)
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, choose **CloudDBA > Performance**.

 **Note** For more information about performance trends, see relevant topics in Database Autonomy Service (DAS) User Guide.

14.9.2. Real-time performance

This topic describes how to view real-time monitoring statistics of your ApsaraDB for MongoDB instances, such as read/write latency, queries per second (QPS), operations, connections, and network traffic.

Prerequisites

Database Autonomy Service (DAS) is authorized to manage ApsaraDB for MongoDB instances.

Procedure

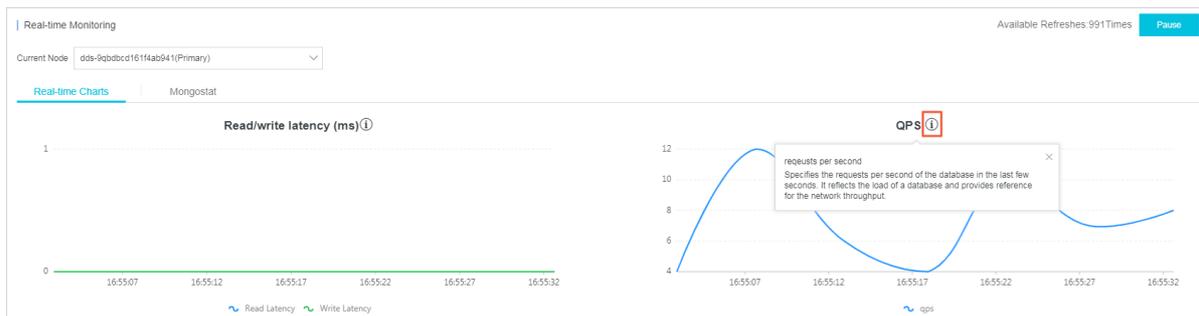
1. [Log on to the ApsaraDB for MongoDB console.](#)

- In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
- Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
- In the left-side navigation pane, choose **CloudDBA > Real-time Performance**.

Overview of the Real-time Monitoring page

On the Real-time Monitoring page, you can click the Real-time Charts or Mongostat tab to view monitoring statistics. When you refresh the **Real-time Monitoring** page, the information on the Real-time Charts and Mongostat tabs is refreshed, and **Available Refreshes** is reset in the upper-right corner.

Real-time charts



By default, content on the **Real-time Charts** tab is displayed when you go to the Real-time Monitoring page. Line charts on the tab are refreshed every 5 seconds.

 **Note** For more information about performance metrics, click the  icon above each chart.

mongostat

| time | query | insert | update | delete | getmore | cmd | dirty | used | qriqw | arfaw | vsize | mapped | in(Byte/s) | out(Byte/s) |
|----------|-------|--------|--------|--------|---------|-----|-------|------|-------|-------|-------|--------|------------|-------------|
| 16:55:33 | 0 | 0 | 0 | 0 | 4 | 35 | 0% | 2% | 0/0 | 0/0 | 1.6G | 0 | 2.29 k | 24.82 k |
| 16:55:28 | 0 | 0 | 0 | 0 | 2 | 66 | 0.3% | 2% | 0/0 | 0/0 | 1.6G | 0 | 2.04 k | 14.32 k |
| 16:55:22 | 0 | 0 | 0 | 0 | 4 | 43 | 0.3% | 2% | 0/0 | 0/0 | 1.6G | 0 | 2.13 k | 10.73 k |
| 16:55:17 | 0 | 0 | 0 | 0 | 2 | 19 | 0.3% | 2% | 0/0 | 0/0 | 1.6G | 0 | 1.30 k | 9.76 k |
| 16:55:13 | 0 | 0 | 0 | 0 | 2 | 59 | 0.3% | 2% | 0/0 | 0/0 | 1.6G | 0 | 1.78 k | 13.34 k |
| 16:55:07 | 0 | 0 | 2 | 5 | 8 | 98 | 0.3% | 2% | 0/0 | 0/0 | 1.6G | 0 | 3.98 k | 13.08 k |
| 16:55:02 | 0 | 0 | 0 | 0 | 4 | 19 | 0.29% | 2% | 0/0 | 0/0 | 1.6G | 0 | 1.75 k | 9.01 k |
| 16:55:57 | 0 | 0 | 0 | 0 | 2 | 56 | 0.29% | 2% | 0/0 | 0/0 | 1.6G | 0 | 1.91 k | 13.92 k |
| 16:55:52 | 0 | 0 | 0 | 0 | 4 | 22 | 0.29% | 2% | 0/0 | 0/0 | 1.6G | 0 | 1.96 k | 9.57 k |
| 16:55:47 | 0 | 0 | 0 | 0 | 2 | 19 | 0.29% | 2% | 0/0 | 0/0 | 1.6G | 0 | 1.19 k | 9.01 k |
| 16:55:42 | 0 | 0 | 0 | 0 | 4 | 36 | 0.29% | 2% | 0/0 | 0/0 | 1.6G | 0 | 2.12 k | 11.42 k |
| 16:55:37 | 0 | 0 | 0 | 0 | 2 | 43 | 0% | 2% | 0/0 | 0/0 | 1.6G | 0 | 1.69 k | 11.79 k |
| 16:55:32 | 0 | 0 | 0 | 0 | 4 | 36 | 0% | 2% | 0/0 | 0/0 | 1.6G | 0 | 2.43 k | 24.82 k |
| 16:55:27 | 0 | 0 | 0 | 0 | 2 | 38 | 0.29% | 2% | 0/0 | 0/0 | 1.6G | 0 | 1.40 k | 11.06 k |

Click the **Mongostat** tab. On the tab, you can view Mongostat command outputs. A new line of monitoring data is added every 5 seconds. The tab can contain up to 999 lines of information.

 **Note** For more information about Mongostat command outputs, see [MongoDB official documentation](#).

14.9.3. Instance sessions

This topic describes how to view real-time monitoring statistics of your ApsaraDB for MongoDB instances, such as read/write latency, queries per second (QPS), operations, connections, and network traffic.

Prerequisites

Database Autonomy Service (DAS) is authorized to manage your ApsaraDB for MongoDB instances.

View instance sessions

1. [Log on to the ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, choose **CloudDBA > Sessions**.

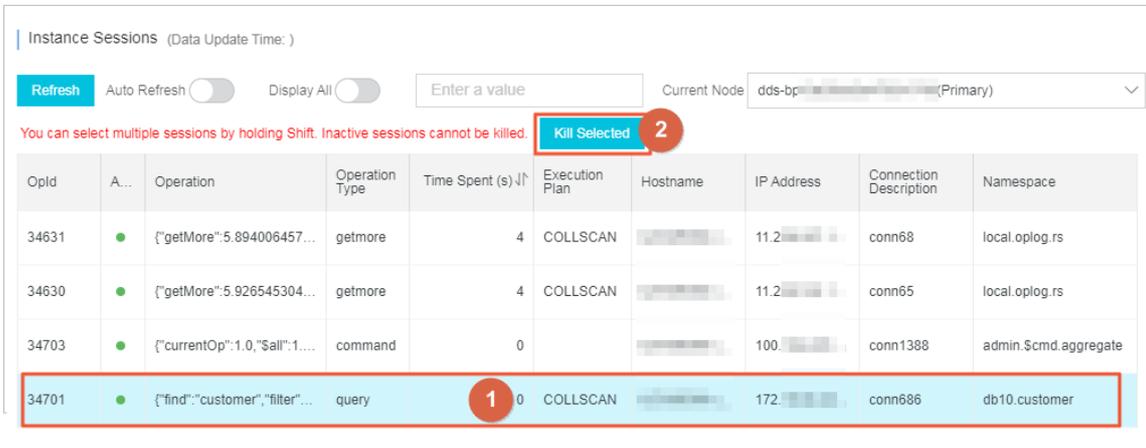
Note

- If you turn on **Auto Refresh**, the system updates session data on the page every 5 seconds.
- By default, the system displays only active sessions. You can turn on **Display All** to view both active and inactive sessions.
- In the **Session Statistics** section, you can view information about sessions in the **Overview**, **Statistics by Client**, and **Statistics by Namespace** charts.

Terminate instance sessions

 **Warning** To avoid unexpected results, we recommend that you do not terminate system-level sessions.

1. In the **Instance Sessions** section, select the sessions that you want to terminate and click **Kill Selected**.



Instance Sessions (Data Update Time:)

Refresh Auto Refresh Display All Enter a value Current Node dds-bp-... (Primary)

You can select multiple sessions by holding Shift. Inactive sessions cannot be killed. Kill Selected 2

| OpId | A... | Operation | Operation Type | Time Spent (s) ↓↑ | Execution Plan | Hostname | IP Address | Connection Description | Namespace |
|-------|------|---------------------------------|----------------|-------------------|----------------|----------|------------|------------------------|-----------------------|
| 34631 | ● | {"getMore":5.894006457... | getmore | 4 | COLLSCAN | ... | 11.2... | conn68 | local.oplog.rs |
| 34630 | ● | {"getMore":5.926545304... | getmore | 4 | COLLSCAN | ... | 11.2... | conn65 | local.oplog.rs |
| 34703 | ● | {"currentOp":1.0,"\$all":1.... | command | 0 | | ... | 100... | conn1388 | admin.\$cmd.aggregate |
| 34701 | ● | {"find":{"customer","filter"... | query | 0 | COLLSCAN | ... | 172... | conn686 | db10.customer |

2. In the message that appears, click **OK**.

14.9.4. Storage analysis

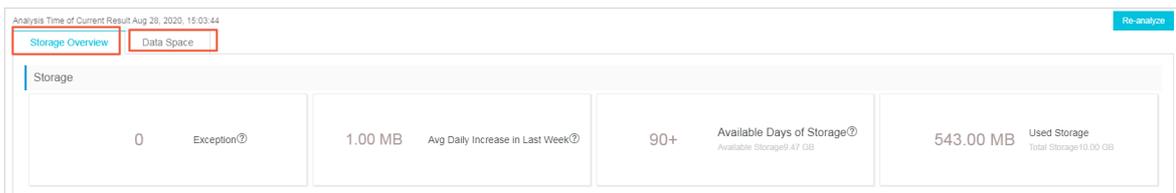
This topic describes how to view information about the storage analysis feature, including **Storage**, **Exceptions**, **Storage trend**, **Tablespaces**, and **Data space**. The information helps you identify and resolve exceptions in the database storage to ensure database stability.

Prerequisites

Database Autonomy Service (DAS) is authorized to manage ApsaraDB for MongoDB instances.

Procedure

1. Log on to the [ApsaraDB for MongoDB console](#).
2. In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
3. Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
4. In the left-side navigation pane, choose **CloudDBA** > **Storage analysis**.
5. In the upper-right corner, click **Re-analyze**. Then, wait until the analysis is complete.
6. On the **Storage Overview** or **Data Space** tab, view the analysis results.



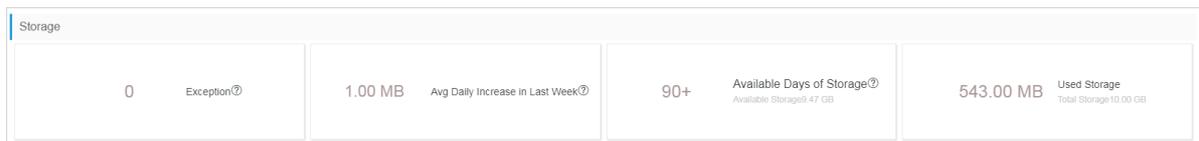
For more information about Storage Overview, see [Storage Overview](#).

For more information about Data Space, see [Data Space](#).

Storage Overview

On the Storage Overview tab, you can view information in the **Storage**, **Exceptions**, **Storage Trend**, and **Tablespaces** sections.

- **Storage**



| Item | Description |
|-----------|--|
| Exception | The number of detected storage exceptions. ApsaraDB for MongoDB can detect the following types of exceptions: <ul style="list-style-type: none"> Over 90% of the storage capacity is used. The total physical storage can become unavailable in seven days. The number of indexes in a collection exceeds 10. |

| Item | Description |
|--|--|
| Avg Daily Increase in Last Week | <p>The average daily increase of storage usage over the last seven days. Formula: (Storage usage at the time of collection - Storage usage seven days ago)/7.</p> <div style="border: 1px solid #ADD8E6; padding: 10px; margin-top: 10px;"> <p>Note</p> <ul style="list-style-type: none"> The increase speed is the average value during the seven days before the collection time. This parameter is used only as a reference for scenarios in which the traffic remains stable. Abrupt storage changes caused by batch imports, deletion of historical data, instance migration, or instance re-creation affect the accuracy of the data. </div> |
| Available Days of Storage | <p>The estimated number of days during which storage space is available. Formula: Size of available storage space/Average daily increase over the last week.</p> <div style="border: 1px solid #ADD8E6; padding: 10px; margin-top: 10px;"> <p>Note</p> <ul style="list-style-type: none"> 90+ indicates that the storage space is sufficient for more than 90 days of usage. This parameter is used only as a reference for scenarios in which the traffic remains stable. Abrupt storage changes caused by batch imports, deletion of historical data, instance migration, or instance re-creation affect the accuracy of the data. </div> |
| Used Storage | The size of used storage space in contrast to the total size of storage space. |

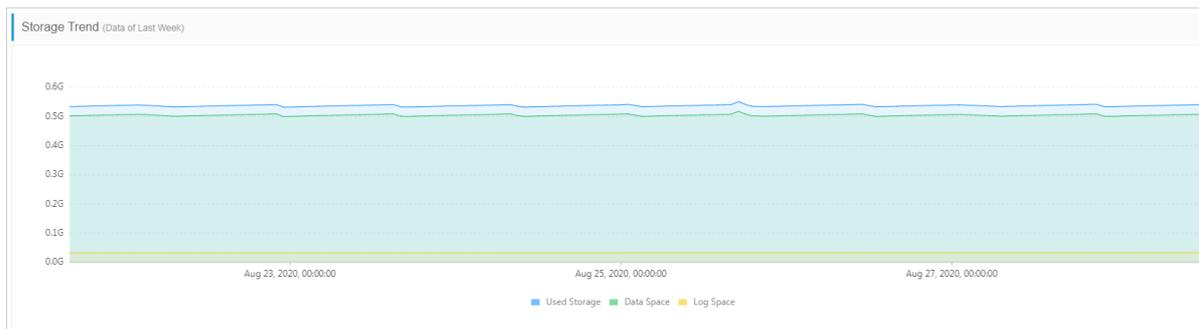
• **Exceptions**

Information about detected storage exceptions. You can resolve the exceptions based on the information in this section.

| Exceptions | | | |
|---------------------------------------|----|-----------|------------|
| Table/Collection Name (Click to View) | DB | Exception | Start Time |
| No storage exceptions found | | | |

• **Storage Trend**

Changes in storage usage over the last week, such as changes in the used storage, data space, and log space.



• **Tablespaces**

Information about all tables, such as the database name, storage engine, and collection storage.

Note You can click the name of a collection to view its indexes.

Data Space

The Data Space tab shows the total storage capacity and tablespace information of each database.

- Note**
- You can click the name of a data space to view its tablespace information.
 - You can click the name of a collection to view its indexes.

14.9.5. Slow query logs

This topic describes how to view slow query logs of ApsaraDB for MongoDB instances. You can identify, analyze, diagnose, and track slow query logs to create indexes. This improves the utilization of the instance resources.

Procedure

- Log on to the **ApsaraDB for MongoDB console**.
- In the left-side navigation pane, click **Replica Set Instances** or **Sharded Cluster Instances**.
- Find the instance and then click the instance ID or choose  > **Manage** in the **Actions** column. Then, the **Basic Information** page appears.
- In the left-side navigation pane, choose **CloudDBA > Slow Query Logs**.

Note By default, slow query logs generated in the past 15 minutes are displayed in the trend chart. You can specify the time range and click **Search** to view slow query logs in specific periods of time. The maximum time range is one day.

- View details of slow query logs by using one of the following methods:

Method 1:

- Click the **Slow Log Details** tab in the lower part of the page.
- On the **Slow Log Details** tab, select the database that you want to query.

Note If the request content of the database is hidden, you can move the pointer over the corresponding content in the **Request Content** column and view the complete content.

Method 2:

- i. In the Slow Log Trend chart, click a point in time. Then, you can view the statistics of the slow query logs generated at the point in time on the **Slow Log Statistics** tab.



- ii. On the **Slow Log Statistics** tab, click **Sample** in the **Actions** column corresponding to a slow query log. In the **Slow Log Sample** dialog box, you can view details of the slow query log.

Slow Log Sample Note: Binary data in the sample is replaced with the \$binData string

| Execution Finish Time | Actions | Namespace | Request Content | User | Client | Avg Execution Duration (ms) | docsExamined | Avg keysExamined | Avg Returned Rows |
|------------------------|----------|-------------|--|------|--------------|-----------------------------|--------------|------------------|-------------------|
| Aug 28, 2020, 14:04:25 | ismaster | admin.\$cmd | ["op":"command","ns":"admin.\$cmd","command":{"ismaster":1},"client":{"driver":"..."}] | | 11.200.150.7 | 295.00 | - | - | - |

| Operation Type | Namespace | Request Template | Total Executions (↑) | Avg Execution Duration (ms) (↓) | Max Execution Duration (ms) (↓) | Avg DocsExamined (↓) | Max DocsExamined (↓) | Avg KeysExamined (↓) | Max KeysExamined (↓) | Avg Returned Rows (↓) | Max Returned Rows (↓) | Actions |
|----------------|-------------|------------------|----------------------|---------------------------------|---------------------------------|----------------------|----------------------|----------------------|----------------------|-----------------------|-----------------------|-----------------|
| ismaster | admin.\$cmd | 0 | 2 | 247.000 | 295 | - | - | - | - | - | - | Sample Optimize |
| isMaster | admin.\$cmd | 0 | 1 | 117.000 | 117 | - | - | - | - | - | - | Sample Optimize |

Note If the request content of the database is hidden, you can move the pointer over the corresponding content in the **Request Content** column and view the complete content.

Export slow query logs

You can click **Export Slow Log** on the **Slow Log Statistics** tab to save the slow query log information to your computer.

15. Data Management (DMS)

15.1. What is DMS?

Data Management (DMS) is a fully managed service that is provided by Alibaba Cloud. This service allows you to manage data, schema, development procedures, development specifications, users, and permissions. DMS also provides security control to ensure secure access to databases.

Supported databases

- Relational databases: MySQL, SQL Server, PostgreSQL, PolarDB-X (previously called DRDS), Oracle, and ApsaraDB for OceanBase.
- NoSQL databases: ApsaraDB for Redis and ApsaraDB for MongoDB.
- Analytical databases: AnalyticDB for MySQL and AnalyticDB for PostgreSQL.

Features

- DMS provides support for the database development process. The process includes the following stages: 1. Design table structures in an on-premises environment based on the predefined design specification. 2. Publish and produce SQL reviews that are included in code and schemas to a specified environment on demand. The SQL reviews are included in code in schemas. The SQL reviews in code are used to add, remove, modify, or query rows.
- DMS provides fine-grained access control at the database, table, or field level. You can perform the required operations on databases in the DMS console. The operations can be traced and audited.
- DMS allows you to configure the required operation specifications and approval processes for multiple modules. These modules include the schema design, data change, data export, and permission request.
- DMS integrates database development with database interaction. You can manage databases without the need to switch between database endpoints at a high frequency by using database accounts and passwords.
- DMS provides the task orchestration feature. This feature allows you to orchestrate and schedule SQL tasks for databases on a regular basis. You can use this feature to perform the required operations with ease. For example, you can dump historical data, analyze periodical reports, and generate analytical results.

15.2. Quick start

15.2.1. Log on to the DMS console

This topic uses Google Chrome to describe how to log on to the Data Management (DMS) console.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

 **Note** When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top navigation bar, move the pointer over **Products** and click **Data Management** under **Database Services**.
5. Set the **Organization** and **Region** parameters and click **DMS**.

 **Note** If you log on to the DMS console as a DMS administrator and your account is added to multiple tenants, you can move the pointer over the  icon in the upper-right corner and select **Switch tenant** to switch to another tenant.

15.2.2. Customize the top navigation bar

Data Management (DMS) allows you to customize the top navigation bar so that you can access the features that are commonly used in DMS with ease.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, move the pointer over the **More** icon.
3. Turn on the  switch in the upper-right corner so that you can customize the top navigation bar.
 - on: You can customize the top navigation bar.
 - off: By default, all features are displayed. You cannot customize the top navigation bar.
4. Find a feature that you want to add to favorites and click the  icon. The feature is displayed in the top navigation bar.

 **Note** If a great many features are added to favorites, move the pointer over **More** in the top navigation bar to view the favorite features that are hidden.

15.2.3. Add an instance

Before you can manage database instances in Data Management (DMS), you must add database instances to DMS. DMS allows you to add ApsaraDB instances and self-managed database instances for which public IP addresses are specified. This topic describes how to add an ApsaraDB RDS for MySQL instance to DMS.

Procedure

1. [Log on to the DMS console](#).
2. In the left-side navigation pane, click **Add Instance**.

 **Note** You can also click the + icon in the left-side navigation pane to add an instance.

3. In the Add Instance dialog box, click the **Cloud** tab.
4. On the Cloud tab, select the type of the instance.
5. In the Basic Information section of the dialog box, set the required parameters. This example shows how to add an ApsaraDB RDS for MySQL instance.

Add instance
✕

✔ Database Source
 2 Basic Information/Advanced information

▾ Basic Information

* Database Source

* Database type
 ▾

* Instance Area
 ▾

* Entry mode
 Connection string address

Connection string address

Database account

Database password

* Control Mode

[Click here to learn](#)

>
Advanced information (View environment type, name, DBA, and more advanced features)

| Section | Parameter | Description |
|-------------------|----------------------------------|--|
| Basic Information | Data Source | The source of the database. In this example, select Cloud . |
| | Database Type | The type of the database. In this example, select MySQL . |
| | Instance Area | The area where the database instance resides. Select a region from the drop-down list. |
| | Entry mode | Valid value: Connection string address . |
| | Connection string address | The endpoint of the database instance. The endpoint contains a port number. |
| | Database account | The username of the account that is used to log on to the database of the instance. |
| | Database password | The password of the account. |

| Section | Parameter | Description |
|-----------------------------|--------------------------------|--|
| | Control Mode | <p>The control mode that is used to manage the database. For more information, see Control modes.</p> <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p> Note If you set this parameter to Security Collaboration, you must set the Security Rules parameter.</p> </div> |
| Advanced Information | Environment type | The type of the environment to which the database instance belongs. |
| | Instance Name | The name of the database instance. |
| | Enable DSQL | Specifies whether to enable the cross-database query feature. To enable the cross-database query feature, you must specify the name of a custom database link. For more information, see Cross-database query . |
| | Lock-free Schema Change | Specifies whether to enable the lock-free schema change feature. |
| | DBA | The database administrator (DBA) of the database instance. The DBA is used to handle approval processes, such as permission requests. |
| | query timeout(s) | The security policy. When the specified timeout period is reached, the execution of the SQL statement is terminated. This way, the security of the database is ensured. |
| | export timeout(s) | The security policy. When the specified timeout period is reached, the execution of the SQL export statement is terminated. This way, the security of the database is ensured. |

6. After the preceding parameters are set, click **Basic Information** and click **Test connection** to verify the settings.

 **Note** If the connection test fails, check the specified parameters based on the error message that appears.

7. Click **Submit**.

Result

After the preceding steps are performed, the ApsaraDB instance is added to DMS. You can view and manage your database instance in the left-side navigation pane of the DMS console.

15.2.4. Add a user

Data Management (DMS) allows you to manage users. You can add users and assign the required roles to each user based on your business requirements.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **System > User**.

 **Note** On the User tab, you can perform the required operations on the existing users. For example, you can edit, disable, enable, or delete a user.

3. Click **New**.
4. In the Add User dialog box, set the required parameters. The following table describes the parameters.

| Parameter | Description |
|-----------------------|--|
| Alibaba Cloud Account | <p>The ID of an Apsara Stack tenant account. You can enter one of the following IDs:</p> <ul style="list-style-type: none"> ◦ The ID of an Apsara Stack tenant account. You can obtain the ID from the account owner. ◦ The ID of a Resource Access Management (RAM) user. You can obtain the required ID from the Service-linked Roles page of the Apsara Uni-manager Management Console. |
| Role | <p>The role that you want to assign to the user based on your business requirements. Valid values:</p> <ul style="list-style-type: none"> ◦ Regular User ◦ DBA ◦ Administrator ◦ Security Administrator ◦ Technical Support <p> Note For more information about the features that are supported by each role, see Features that are supported by each role.</p> |

5. Click **OK**.

15.2.5. Use the sharing feature

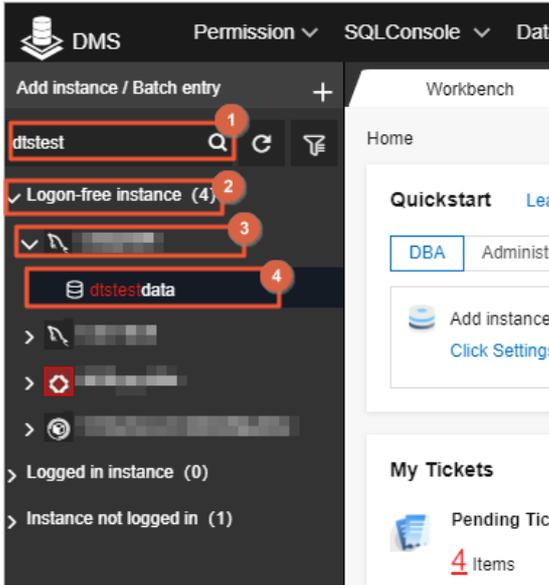
Data Management (DMS) provides the sharing feature for you to share an SQLConsole tab or a ticket. You can use a share link to access an SQLConsole tab or a ticket with ease. This improves efficiency. This topic describes how to use the sharing feature in DMS.

Scenarios

- You can share historical tickets. This improves the audit efficiency of security administrators.
- You can obtain the information about historical changes to databases based on the share links of historical data change tickets.
- You can query and update databases based on the share links of SQLConsole tabs.

Share an SQLConsole tab

1. [Log on to the DMS console](#).
2. In the left-side navigation pane, enter the name of a database in the search box and press the Enter key. From the matched results, click the instance in which the database resides and double-click the database. The SQLConsole tab appears.



- On the **SQLConsole** tab, click the  icon in the upper-right corner.

 **Note** You can share the SQLConsole tab of a database only after you obtain the permissions on the database. Otherwise, the  icon is not displayed.

- In the dialog box that appears, click **Open sharing**.
- Set the parameters as required.

| Parameter | Description |
|----------------|---|
| Sharing form | The content to be shared. Valid values: Console only and Console + SQL . <ul style="list-style-type: none"> ◦ Console only: the SQLConsole tab without SQL statements. ◦ Console + SQL: the SQLConsole tab including the SQL statements that you entered. |
| Sharing period | The validity period of the share link. Valid values: 7 Days, 30 Days, 180 Days, and 360 Days . |

- Click **Generate and copy links**.

 **Note** After a share link is generated, you can turn off **Turn off sharing** if you no longer want to share the SQLConsole tab. Then, the generated share link becomes invalid.

Share a ticket

- Log on to the **DMS console**.
- On the **Workbench** tab of the DMS console, click **Submitted Tickets** in the **My Tickets** section. The **My Tickets** page appears.
- On the **My Tickets** page, find the ticket that you want to share and click the ticket number to go to the **Ticket Details** page.

 **Note** Schema design tickets cannot be shared.

4. On the **Ticket Details** page, click the  icon in the upper-right corner.
5. In the dialog box that appears, click **Open sharing**.
6. Set the parameters as required.

| Parameter | Description |
|----------------|--|
| Share object | The recipients of the share link. You can specify one or more recipients. Valid values: Everyone and Designated . <ul style="list-style-type: none"> ◦ Everyone: all the users of the current DMS tenant. ◦ Designated: one or more specific users of the current DMS tenant. |
| Sharing period | The validity period of the share link. Valid values: 7 Days , 30 Days , 180 Days , and 360 Days . |
| Designated | The specific users to whom you want to send the share link. You can specify one or more users. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note This parameter is displayed only after you set the Share object parameter to Designated.</p> </div> |

7. Click **Generate and copy links**.

 **Note** After a share link is generated, you can turn off **Turn off sharing** if you no longer want to share the ticket. Then, the generated share link becomes invalid.

15.3. Control modes

DMS provides three control modes for you to manage instances: Flexible Management, Stable Change, and Security Collaboration. You can specify a control mode for each instance.

| Control mode | Description | Scenario | Logon method |
|---------------------|--|--|--|
| Flexible management | This control mode allows you to manage the visualized data and schemas of multiple types of databases. It also provides a variety of data management solutions. This simplifies the use of databases and facilitates management. | <ul style="list-style-type: none"> • Database instances do not require strict control. • Database instances are used by a single user. | A database account and the related password. |

| Control mode | Description | Scenario | Logon method |
|------------------------|--|--|--|
| Stable change | <ul style="list-style-type: none"> This control mode provides multiple solutions to ensure database reliability. These solutions allow you to change data without the need to lock the related table or schema. All features that are included in the flexible management control mode are available. | <ul style="list-style-type: none"> Database instances require a high level of availability. This ensures that these database instances function as expected for an extended period of time. Database instances are used by a small-sized group that includes multiple users. | A database account and the related password. |
| Security collaboration | <ul style="list-style-type: none"> This control mode provides multiple solutions to ensure data security. These solutions include fine-grained access control at the database, table, or field level and sensitive data management. This control mode allows you to produce enterprise-specific database DevOps solutions through custom design specifications and approval processes. All features that are included in the flexible management and stable change control modes are available. | <ul style="list-style-type: none"> Ensure the data security of database instances. Implement strict access control over development or change workflows. Manage compliance for enterprises. | Logon-free through authorization. |

 **Note** The instances that are managed in Stable Change mode consume the billing quota of the instances that are managed in Security Collaboration mode.

15.4. Features that are supported by each role

DMS provides the following roles: regular user, DBA, security administrator, and DMS administrator. This topic describes the features that are supported by each role.

| Category | Feature | Regular user | DBA | Security administrator | DMS administrator | Description |
|------------|-----------------------|--------------|-----|------------------------|-------------------|---|
| Permission | Permission management | √ | √ | √ | √ | You can use this feature to apply for permissions on instances, databases, tables, and sensitive fields. You can also view permissions that you have. |
| | Data Changes | √ | √ | √ | √ | You can use this feature to initialize data for a newly published project, clean up historical data, fix bugs, or run a test. |
| | Data Import | √ | √ | √ | √ | You can use this feature to import a large amount of data to your databases at a time. |

| Category | Feature | Regular user | DBA | Security administrator | DMS administrator | Description |
|--------------|----------------------------|--------------|-----|------------------------|-------------------|--|
| Data Plans | Data Export | √ | √ | √ | √ | You can use this feature to export a large amount of data for analysis or export the required data. |
| | Data Tracking | √ | √ | √ | √ | If specific data fails to meet your requirements due to reasons such as misoperation, you can use this feature to restore data to the normal state. |
| | Test Data Generate | √ | √ | √ | √ | Some business scenarios may require frequent data preparation. In this case, you can use this feature to generate test data to ensure data security and discreteness and improve production efficiency. |
| | Data Warehouse Development | √ | √ | √ | √ | DMS uses a database as a computing engine and integrates various tools and services, such as Data Transmission Service (DTS) and Data Lake Analytics (DLA), in the database ecosystem for data warehouse development. You can use this feature to develop and manage data warehouses in DMS with ease. |
| | Data Service | √ | √ | √ | √ | You can use this feature to export data at the field or row level, display data in a visualized manner, and publish API operations to the Alibaba Cloud Marketplace for sale. |
| | Database Clone | √ | √ | √ | √ | You can use this feature to clone MySQL databases. |
| Schemas | Schema Design | √ | √ | √ | √ | When you develop or optimize projects or process new business requirements, you can use this feature to change schemas. For example, you can use this feature to create a table or modify an existing table. |
| | Table Sync | √ | √ | √ | √ | You can use this feature to compare and synchronize the schemas of tables in different environments, such as online and offline environments. This feature helps ensure the consistency of schemas. |
| Optimization | SQL Review | √ | √ | √ | √ | You can use this feature to prevent SQL statements that do not use indexes or do not conform to database development standards. This feature helps protect against SQL injection attacks. |

| Category | Feature | Regular user | DBA | Security administrator | DMS administrator | Description |
|---------------------|--------------------------|--------------|-----|------------------------|-------------------|---|
| SQLConsole | Single Database query | √ | √ | √ | √ | You can write SQL statements to query data in a single database. This feature can be used to verify business code, analyze product effects, and identify issues in an online environment. |
| | Cross-database Query | √ | √ | √ | √ | You can use this feature to perform join queries across online heterogeneous databases that are deployed in different environments. |
| System Management | Instance management | × | √ | × | √ | You can use this feature to manage instances. For example, you can register, view, or edit instances. |
| | User management | × | × | × | √ | You can use this feature to manage users. For example, you can add, view, or edit users as needed. |
| | Task management | × | √ | × | √ | You can use this feature to manage tasks. For example, you can create, start, or stop tasks. |
| | Configuration management | × | × | × | √ | You can use this feature to view and modify system configurations, or view the historical modifications of the configurations. |
| Security management | Security Rules | × | √ | × | √ | You can use this feature to configure security rules. Only SQL statements that conform to the security rules can be executed. |
| | Approval Processes | × | √ | × | √ | Approval processes are associated with security rules. You can configure different approval processes for different types of tickets. |
| | Operation Logs | × | √ | √ | √ | Operations logs record data changes. Each record contains information such as the user who performed the operation, operation details, and time at which the operation was performed. You can use this feature to track historical user operations at any time. |
| | Access IP Whitelists | × | × | × | √ | After you configure an access IP whitelist, only the IP addresses or Classless Inter-Domain Routing (CIDR) blocks in the whitelist can access the resources within your DMS tenant. This effectively enhances data security. |

| Category | Feature | Regular user | DBA | Security administrator | DMS administrator | Description |
|----------|-------------------|--------------|-----|------------------------|-------------------|--|
| | Sensitive Data | × | √ | √ | √ | You can use this feature to manage sensitive data. For example, you can use algorithms to de-identify sensitive data or adjust the security levels of sensitive data. |
| Tickets | Ticket management | √ | √ | √ | √ | You can use this feature to configure notification methods. DMS can notify you of the approval or execution status of tickets by using DingTalk notifications or emails. |

15.5. Apply for permissions

You can apply for the query, change, and export permissions on a database, table, or column. After the database owner approves your application, you can query, change, and export data.

Permissions

- **Query permissions:** the permissions to execute SQL statements in the SQLConsole to query the data of the object on which you want to apply for the permission.
- **Change permissions:** the permissions to submit tickets to change data or synchronize data in a database or table. You cannot change data without approval.
- **Export permissions:** the permissions to submit tickets to export data from the object on which you want to apply for the permission. You cannot export data without approval.

Permission categories that are supported by each control mode

| Permission category | Permissions | Control mode | | |
|-----------------------|---|---------------------|---------------|------------------------|
| | | Flexible management | Stable change | Security collaboration |
| Instance logon | After you obtain the instance logon permission on an instance, you can use the preset database account and password to log on to the instance. | √ | √ | × |
| Database | <p>Database permissions are classified into query, export, and change permissions. After you obtain the database permissions on a database, you can access the following resources of the database: 1. All the insensitive fields. 2. All the tables to which row-level control settings are not applied. 3. All new tables.</p> <ul style="list-style-type: none"> • Query permission: You can execute SQL statements in the SQLConsole to query data. • Change permission: You can submit data change and data import tickets. • Export permission: You can submit data export tickets. | × | × | √ |

| Permission category | Permissions | Control mode | | |
|------------------------|--|---------------------|---------------|------------------------|
| | | Flexible management | Stable change | Security collaboration |
| Table | <p>Table permissions are classified query, export, and change permissions. After you obtain the table permissions on a table, you have full access to all data in the table except sensitive fields.</p> <ul style="list-style-type: none"> • Query permission: You can execute SQL statements in the SQLConsole to query data. • Change permission: You can submit data change and data import tickets. • Export permission: You can submit data export tickets. | x | x | √ |
| Sensitive field | <p>Sensitive field permissions are classified into query, export, and change permissions. After you obtain the sensitive field permissions on sensitive fields in a table, you have full access to all sensitive fields in the table. Before you apply for the sensitive field permissions, you must obtain the database and table permissions to which the sensitive fields belong.</p> <ul style="list-style-type: none"> • Query permission: You can execute SQL statements in the SQLConsole to query data. • Change permission: You can submit data change and data import tickets. • Export permission: You can submit data export tickets. | x | x | √ |
| Database owner | <ul style="list-style-type: none"> • The owner of a database can manage the permissions on the database. For example, the owner of a database can grant or revoke permissions on the database and tables in the database. • The owner of a database can query all data in the database except sensitive or confidential fields. The owner can also submit tickets to perform operations on the data and schemas in the database without the need to apply for permissions. • DMS automatically identifies and assigns database owners to owner nodes in approval processes. | √ | √ | √ |
| Table owner | <ul style="list-style-type: none"> • The owner of a table can manage the permissions on the table. For example, the owner can grant or revoke permissions on the table. • The owner of a table can query all data in the table except sensitive or confidential fields. | √ | √ | √ |

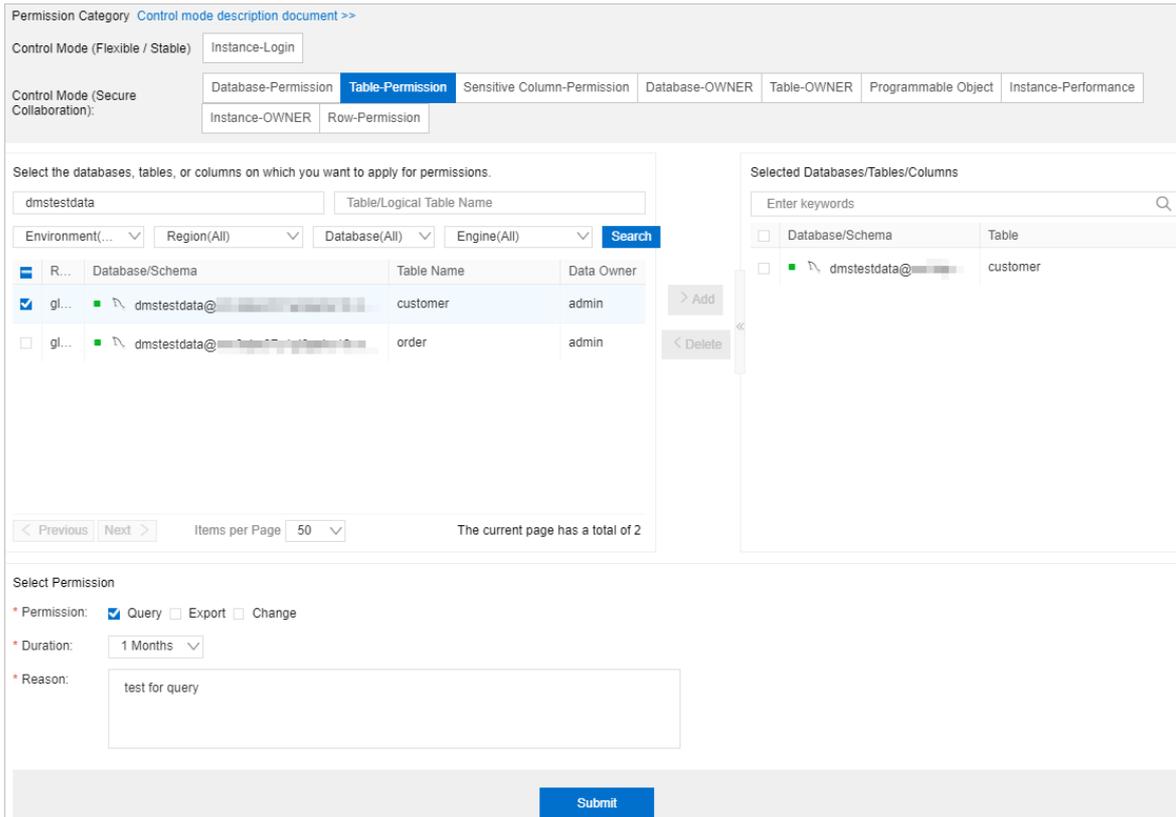
| Permission category | Permissions | Control mode | | |
|-----------------------------|---|---------------------|---------------|------------------------|
| | | Flexible management | Stable change | Security collaboration |
| Programmable object | <p>Programmable object permissions are classified into query, export, and change permissions.</p> <ul style="list-style-type: none"> • Query permission: You can execute SQL statements in the SQLConsole to query data. • Change permission: You can submit data change and data import tickets. • Export permission: You can submit data export tickets. | x | x | x |
| Instance performance | You can apply for permissions to view the performance of instances that are managed in security collaboration mode. | x | x | √ |
| Instance owner | <ul style="list-style-type: none"> • The owner of an instance can manage the permissions on the instance. For example, the owner of an instance can grant or revoke permissions on the instance. • The owner of an instance can query all data in the databases of the instance except sensitive or confidential fields. The owner can also submit tickets to perform operations on the data and schemas of the instance without the need to apply for permissions. | √ | √ | √ |
| Row | <p>Row permissions are classified into query, export, and change permissions. You can apply for permissions on specific values of a managed field in a table. You can also apply for permissions on all values of a managed field in a table.</p> <ul style="list-style-type: none"> • Query permission: You can execute SQL statements in the SQLConsole to query data. • Change permission: You can submit data change and data import tickets. • Export permission: You can submit data export tickets. | x | x | x |

Apply for permissions

1. [Log on to the DMS console.](#)
2. In the top navigation bar, choose **More > Permissions** and select a permission category. For more information about permission categories, see [Permission categories that are supported by each control mode.](#)

 **Note** In the top search box, you can search for databases or tables by name. In the search results, find the required database or table and click **Access apply** in the **Actions** column.

3. Set the required parameters of the permission for which you want to apply.



- i. Select the category of the permission for which you want to apply.
- ii. Select the databases, tables, or columns on which you want to apply for permissions.

Note Enter keywords, specify filter conditions, and then click Search to search for databases or tables. The keywords that you enter can contain percent signs (%) as wildcards. In the search results, select the databases or tables on which you want to apply for permissions. Then, click Add.

- iii. Select the type of permissions for which you want to apply and specify the duration for which you want to have the permission. Then, enter the reason for your application.

4. Click **Submit** and wait for approval.

Note You can view the status of application ticket in the My Tickets section of the Workbench tab.

Manage permissions

| Management type | Action | Description |
|--------------------|---------------------|---|
| Active management | Release permissions | On the Workbench tab, click Effective Permissions in the Permissions section. Select the object on which you want to release permissions and click Release Permission . |
| | Renew permissions | On the Workbench tab, click Expiring Permissions in the Permissions section to view and check the permissions that are about to expire. If you want to renew a permission, submit a ticket to apply for the renewal. |
| Passive management | N/A | The owner of a database can view and assess the rationality of permission assignments at any time and manage the permissions that are granted. |

 **Note** All the operations that you perform to apply for, release, revoke, and grant permissions are recorded in operation logs. To view the operation logs, choose **System Management > Security > Operation Logs** in the top navigation bar.

15.6. SQLConsole

15.6.1. Single database query

The single database query feature allows you to execute various SQL statements in the SQLConsole of the Data Management (DMS) console with ease. You can use this feature to visualize the add, delete, modify, and query operations on the data in a database. This feature applies to scenarios, such as data queries and data development.

Prerequisites

You are granted the query permission on the database or table that you want to query.

Precautions

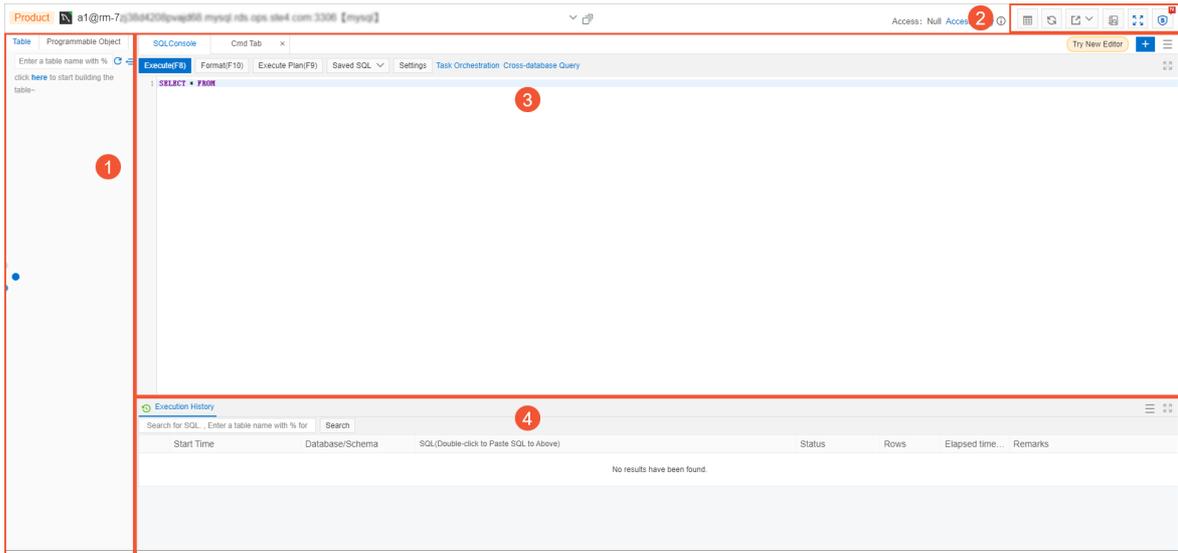
- A table may contain sensitive or confidential fields. You do not have permissions to access these fields. Therefore, the values of these fields are displayed as ********* in the query results. For more information, see [Manage sensitive data](#).
- By default, a maximum of 200 data rows can be returned for each query. If you are an administrator, you can change this value based on your business requirements. To change this value, perform the following steps: 1. Log on to the Data Management (DMS) console. 2. In the top navigation bar, choose **System > Security > Security Rules**.
- A full scan can be performed on a table that does not exceed 10 GB in size. If you are an administrator, you can change this value based on your business requirements. To change this value, perform the following steps: 1. Log on to the DMS console. 2. In the top navigation bar, choose **System > Security > Security Rules**.
- By default, the timeout period to execute a single SQL statement is 60 seconds. If you are an administrator, you can change this value in the **Advanced information** section of the **Edit instance** dialog box.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **More > SQLConsole > Single Database Query**.

 **Note** To go to the SQLConsole tab, you can also double-click the required database in the left-side instance list of the DMS console.

3. Select the database that you want to query from the drop-down list. You can also search for databases by keyword. After you find and select the required database, click **Confirm**.
4. Enter the SQL statement to be executed on the SQLConsole tab and click **Execute**.



GUI of the SQLConsole

| No. | Section | Description |
|-----|--------------------------|--|
| ① | Visual operation section | <p>In this section, you can visually manage your database.</p> <ul style="list-style-type: none"> Tables You can view all tables, fields, and indexes of the current database. You can also right-click a table in the database to modify the table schema, import data to the table, or export data from the table. Programmable objects You can create, view, execute, and manage programmable objects, such as views, stored procedures, functions, triggers, and events. <p>Note A maximum of 1,000 entries can be displayed.</p> <ul style="list-style-type: none"> Key-value pair information The key-value pair information can be displayed only for a NoSQL database. |

| No. | Section | Description |
|-----|--------------------------|---|
| ② | Extended feature section | <p>In this section, shortcuts to extended features are provided. You can click the icons of the features to use the features. The following table describes the icons.</p> <ul style="list-style-type: none"> ◦ : the Tables icon. You can click the  icon to view the details about the table. Then, click the  icon to return to the SQLConsole tab. ◦ : the Sync Metadata icon. After you click this icon, DMS collects most recent metadata information about the database, such as tables, fields, indexes, and programmable objects. This way, you can manage permissions on tables, fields, and programmable objects based on the security level. ◦ : the Export icon. You can click this icon to export the data of the database, table schemas of the database, or table creation statements. ◦ : the Operation audit icon. You can click this icon to view the information about all data query and data change records. For example, you can query the information about an operation, the operator, and the time when the operation is performed. For more information, see View operations logs. |
| ③ | Command running section | <p>In this section, you can write and execute SQL statements to manage the current database. You can also format SQL statements, create execution plans, save commonly used SQL statements, and configure display settings.</p> <p> Note You can click the  icon to add multiple query tabs.</p> |
| ④ | Execution result section | <p>In this section, you can view the execution results after SQL statements are executed. You can also view the details about a single row and add, delete, or modify data.</p> <p> Note You can click the Execution History tab to view the historical execution records. For example, you can view the time at which the execution of an SQL statement started, the affected database, and the details about the SQL statement. You can also export the execution results as required.</p> |

15.6.2. Cmd Tab

Data Management (DMS) provides the Cmd Tab feature as a CLI for you to write and execute SQL statements. The executed SQL statements and execution results are displayed in the upper part of the Cmd Tab tab. This topic describes the GUI of the Cmd Tab tab and how to use the Cmd Tab tab.

Prerequisites

- The database to be queried is a relational database, such as a MySQL, an Oracle, or an SQL Server database.
- You are granted the query permissions on the database or table that you want to query.

Procedure

1. [Log on to the DMS console](#).

2. In the top navigation bar, move the pointer over the **More** icon and choose **SQLConsole > Cmd Tab**.
3. Select the database that you want to query from the drop-down list or enter a keyword in the field to search for the database. After you select the database, click **Confirm**.
4. Enter the SQL statement to be executed in the lower part of the Cmd Tab tab and press **Ctrl+Enter** or click **Execute** to execute the SQL statement.

15.6.3. Super SQL mode

Data Management (DMS) provides the super SQL mode feature. After you enable this feature as a DMS administrator or a database administrator (DBA), all SQL statements that you execute on the SQLConsole tab are executed without being affected by security rules.

Prerequisites

- You are a DMS administrator or a DBA.
- An instance is managed in Security Collaboration mode.

Context

To enhance the stability and security of databases, DMS administrators and DBAs may configure security rules for the databases. For example, a security rule is configured to prevent unauthorized users from executing DML statements in a production database on the SQLConsole tab. They can execute those statements only by submitting a ticket. However, these security rules may cause inconvenience to privileged users, such as DMS administrators and DBAs.

In view of this, DMS provides the **super SQL mode** feature. If you enable this feature as a DMS administrator or a DBA, all SQL statements that you execute on the SQLConsole tab are executed without being affected by security rules.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, move the pointer over the **More** icon and choose **SQLConsole > Single Database query**.

 **Note** To go to the SQLConsole tab, you can also double-click the database that you want to query in the left-side navigation pane of the DMS console.

3. Select the database that you want to query from the drop-down list or enter a keyword in the field to search for the database. After you select the database, click **Confirm**.
4. On the SQLConsole tab, click the  icon in the upper-right corner. In the message that appears, click **OK**.

Then, the outside borders of the SQLConsole tab turn orange. This indicates that the **super SQL mode** feature is enabled. The SQL statements that you enter on the SQLConsole tab are directly executed.

 **Notice** After you enable this feature as a DMS administrator or a DBA, all SQL statements that you execute on the SQLConsole tab are executed without being affected by security rules.

To disable the **super SQL mode** feature, click the  icon in the upper-right corner.

15.6.4. Cross-database query

Data Management (DMS) provides the cross-database query feature that allows you to perform cross-database queries on online heterogeneous data sources that are deployed in different environments. The cross-database query feature allows you to perform cross-database queries on databases and tables in database instances that are added to DMS.

Prerequisites

- The type of the database instance that you want to query is MySQL, SQL Server, PostgreSQL, PolarDB-X, or Redis.
- The cross-database query feature is enabled for each database instance.

Note If the cross-database query feature is not enabled for a database instance, you can enable the feature in the Edit instance dialog box in the DMS console. Then, enter a database link name in the Advanced information section. The name of each database link must be unique.

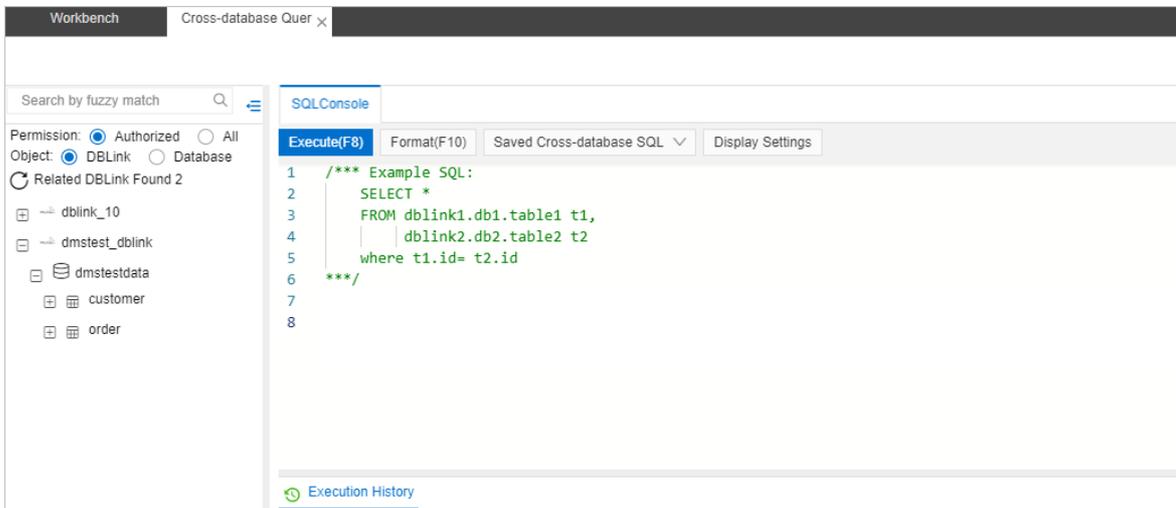
- You are granted the query permission on each database or table that you want to query.

Limits

You can perform cross-database queries only on physical databases. This feature is unavailable for logical databases.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **More > SQLConsole > Cross-database Query**.
3. In the left-side pane of the **Cross-database Query** tab, view the databases on which you have the query permission or all the databases.
4. On the **SQLConsole** tab, enter an SQL statement. You can perform cross-database queries on databases and tables on which you have the query permission in database instances.



Note

- You must specify the table that you want to query in the format of `<DBLinkName>.<databaseName>.<tableName>`, for example, `dmstest_dblink.dmstestdata.customer`.
- In the left-side list, you can double-click a table on which you have the query permission or drag the table to the SQLConsole tab. An SQL statement that is used to query data in the table is automatically generated.

5. Click **Execute(F8)**. You can view the execution results and execution history in the Execution History section.

 **Note** DMS also provides the format, saved cross-database SQL, and display settings features. You can use these features based on your business requirements.

15.7. Data plans

15.7.1. Change data

DMS provides the data change feature that allows you to change data. This topic describes how to use the data change feature to change data.

Context

DMS allows you to submit data change tickets to initialize data for a newly published project, clear historical data, fix bugs, or run a test. The operations that you can perform to change data include insert, update, delete, or truncate operations.

Data change types

| Type | Description |
|-----------------------|--|
| Normal data modify | <p>You can use the normal data modify feature to perform the following data changes:</p> <ul style="list-style-type: none"> • Perform normal data changes. • Perform lock-free schema changes. You can perform this type of operations to change character sets and collations for tables, adjust time zones, and change column data types. Compared with normal data change operations, lock-free schema changes can be performed to achieve the following benefits: <ul style="list-style-type: none"> ◦ Ensures business continuity regardless of whether tables are locked due to database schema changes. ◦ Ensures consistent synchronization between the primary and secondary databases regardless of whether the latency occurs due to native online data definition language (DDL) operations that are performed on the databases. ◦ Reclaims tablespaces and reduces fragmentation rates without the need to lock tables. You no longer need to use the OPTIMIZE TABLE statement that results in table locks. <p> Note You can use this feature only for MySQL databases. Before you use this feature, you must set the OnlineDDL parameter to Open(DMS OnlineDDL first) in the Advanced information section when you add or edit an instance. For more information, see Register database instances with DMS.</p> |
| Lock-free data modify | <p>You can use this feature to change a large amount of data. For example, you can use this feature to delete historical data and update all fields in a table. Multiple SQL statements for data changes are divided and executed at the same time based on the primary key or unique key. This limits the consumption of database performance and space.</p> <p> Note You can use this feature only for MySQL databases.</p> |
| History data clean | <p>You can use this feature to regularly clean historical data. This way, the stability of the production environment is not affected when you obtain historical data.</p> <p> Note You can use this feature only for MySQL databases.</p> |

| Type | Description |
|---------------------|---|
| Programmable object | Databases provide programmable objects such as stored functions and stored procedures. This feature allows you to use the programmable objects to standardize management processes and provide audit records. |

Procedure

1. [Log on to the DMS console.](#)
2. In the top navigation bar, choose **More > Data Plans** and select a data change type.

 **Note** This topic describes how to submit a ticket that is used to **modify normal data**.

3. Specify the required parameters for the data change ticket.

| Parameter | Description |
|------------------------------------|--|
| Reason Category | The reason for the data change. This helps you identify the ticket in the future. |
| Business Background | The purpose or objective of the data import. This reduces communication costs. |
| Change Stakeholder | The stakeholders of the data change. All specified stakeholders can view the details about the ticket and participate in the approval process. Unauthorized users, except for administrators and DBAs, cannot view the details about the ticket. |
| Execution Method | The method that is used to execute the ticket based on your business requirements. |
| Database | The database on which you have the change permission. If you have only query permissions on the database or the change permission on the tables of the database, you cannot submit a data change ticket. |
| Affected Rows | The estimated number of data rows that are affected by the data change. To obtain the actual number of affected rows, you can write an SQL statement that includes the COUNT function on the SQLConsole tab. |
| SQL Statements for Change | The executable SQL statements that are used to export data. You can upload a file to provide the SQL statements. DMS verifies the syntax of the SQL statements when you submit the ticket. If the syntax is invalid, DMS rejects the ticket. |
| SQL Statements for Rollback | The executable SQL statements that are used to roll back the data import operation. You can write the SQL statements in the SQL Text field or upload an SQL that includes the required SQL statements. |
| Attachments | The images or files that are uploaded to add more information about the data change. |

4. After you configure the settings, click **Submit**.
5. After your ticket passes the precheck, click **Submit for Approval**. In the message that appears, click **OK**.
6. After the ticket is approved, click **Execute Change**.
7. Set the Execute Immediately parameter and click **Confirm Execution**.

 **Note** By default, the Execute Immediately switch is turned on. You can turn off the **Execute Immediately** switch and specify a point in time to run the ticket. The system automatically runs the ticket at the specified point in time.

8. Wait until the execution is completed.

15.7.2. Import data

Data Management (DMS) provides the data import feature that allows you to import large amounts of data to a database with ease. This saves manpower and resources.

Supported databases

- Self-managed MySQL databases and ApsaraDB RDS for MySQL databases
- PolarDB-X databases

Supported file formats for data import

- TXT format.
- SQL script. By default, you can use only the INSERT and REPLACE statements to import data to database instances that are managed in Security Collaboration mode. If you want to use other SQL statements to import data, modify the security rules for data import as a database administrator (DBA) or DMS administrator. To modify the security rules, click the **SQL Correct** tab on the **Security Rules** tab and set the Checkpoints parameter to **Batch Data import rules**.
- CSV format. Values in a CSV file must be separated by commas (.). The first row must be field names.

 **Note** The file size cannot exceed 5 GB.

Usage notes

- If you need to change only a small amount of data, we recommend that you submit a Normal Data Modify or Lockless change ticket to ensure stable data change. For more information, see [Change data](#).
- If you submit a Large Data Import ticket to import a large amount of data to a table, the table will be locked even if you set the OnlineDDL parameter to Open (DMS OnlineDDL first) for the database instance.
- We recommend that you use SQL statements with better performance to import a large amount of data, such as the INSERT, UPDATE, and DELETE statements. Indexes of primary keys are used in the UPDATE and DELETE statements.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, move the pointer over the **More** icon and choose **Data Plans > Data Import**.
3. On the Large Data Import tab, set the parameters that are described in the following table.

* Reason ▼

Category:

* Business Background:

Change Stakeholder: ▼

* Database: ▼

* File Encoding: ▼

* SQL SQL Script CSV

Statements for Change:

*
 You can upload only TXT, SQL, and CSV files no greater than 1 GB.

SQL Statements Text Attachment

for Rollback:

Attachments:
 You can upload files in the format of "picture" and "document" to supplement the current work order information.

| Parameter | Description |
|------------------------------------|--|
| Reason Category | The reason for the data import. This helps you identify the ticket in the future. |
| Business Background | The purpose or objective of the data import. This reduces unnecessary communication. |
| Change Stakeholder | The stakeholders involved in the data import. All the specified stakeholders can view the ticket details and assist in the approval process. Irrelevant users other than DMS administrators and DBAs are not allowed to view the ticket details. |
| Database | The database on which you have the change permissions. You cannot submit a data import ticket if you have only the permissions to query data in the database or change data in tables. |
| File Encoding | The encoding algorithm to be used by the database. |
| SQL Statements for Rollback | The executable SQL statements for rolling back the data import. You can write the SQL statements in the SQL Text field or upload an SQL script that includes the required SQL statements. |
| Attachments | The images or files that are uploaded to add more information about the data import. |

4. After you configure the settings, click **Submit**.
5. After your ticket passes the precheck, click **Submit for Approval**. In the message that appears, click **OK**.

6. After the ticket is approved, click **Execute Change**.
7. Set the **Execute Immediately** parameter and click **Confirm Execution**.

Note By default, the **Execute Immediately** switch is turned on. You can turn off the **Execute Immediately** switch and specify a point in time to run the ticket. The system automatically runs the ticket at the specified point in time.

8. Wait until the execution is completed.

15.7.3. Data export

DMS provides the data export feature. You can use this feature to export a database or SQL result sets. Then, you can extract the required data for data analysis.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **More > Data Plans** and select **SQL Result Set Export** or **Database Export**.
3. On the Data Export Ticket Application tab, set the required parameters.
 - o Set the required parameters on the **SQL Result Set Export** tab

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|---------------------|---|
| Reason Category | The reason for the data export. This helps you find the ticket in a timely manner in the future. |
| Business Background | The purpose or objective of the data export. This parameter reduces communication costs. |
| Database Name | The database on which you have the export permission. |
| Affected Rows | The estimated number of data rows that are affected by the data export. To obtain the actual number of affected rows, use the <code>COUNT</code> function in SQL statements on the SQL Editor tab. |
| Skip Validation | Specifies whether to skip validation. If you select Skip Validation , you must enter a reason in the field next to the check box. <div style="border: 1px solid #ccc; background-color: #fff9c4; padding: 5px; margin-top: 10px;"> <p> Warning If you select Skip Validation, DMS does not check the number of rows that are affected by the data export. If a large amount of data is exported, potential risks to your business increase. Proceed with caution.</p> </div> |
| Stakeholder | The stakeholders of the data export. All specified stakeholders can view the details about the ticket and are included in the approval process. Unauthorized users, except for administrators and DBAs, cannot view the details about the ticket. |
| Export Statement | The executable SQL statement that is used to export data. Example: <code>select * from testtable</code> . DMS verifies the syntax of the SQL statement when you submit the ticket. If the syntax is invalid, you cannot submit the ticket. |
| Attachments | The images or files that are uploaded to add more information about the data export. |

o Configure a database export ticket

Data Export Type in Application: SQL Result Set Export Database Export

* Database Name:

Stakeholder:

Export content: Data Structure Data & Structure

Exported Structure: Procedure Function Trigger View Event

Type:

More Options:
> Big data type export options
> SQL script other options

Attachments:
You can upload files in the format of "picture" and "document" to supplement the current work order information.

Tables & Filters (Currently selected 2 item)

| <input checked="" type="checkbox"/> | Table Name | Filter Condition |
|-------------------------------------|------------|------------------|
| <input checked="" type="checkbox"/> | table3 | where 1 = 1 |
| <input checked="" type="checkbox"/> | table4 | where 1 = 1 |

| Parameter | Description |
|--------------------------------|---|
| Database Name | The database on which you have the export permission. After you select the database, you must select the table to which you want to export data and configure filter conditions in the Tables & Filters section. |
| Reason Category | The reason for the data export. This helps you find the ticket in a timely manner in the future. |
| Business Background | The purpose or objective of the data export. This parameter reduces communication costs. |
| Stakeholder | The stakeholders of the data export. All specified stakeholders can view the details about the ticket and are included in the approval process. Unauthorized users, except for DMS administrators and DBAs, cannot view the details about the ticket. |
| Export content | The type of data that you want to export. Valid values: Data , Structure , and Data & Structure . |
| Exported Structure Type | The type of schema that you want to export. |
| More Options | The other objects that you want to export. Click Big data type export options or SQL script other options and select the required options. |
| Attachments | The images or files that are uploaded to add more information about the data export. |

4. After you complete the configurations, click **Submit** and wait for approval.

 **Note** When you export an SQL result set, DMS prechecks the SQL statements. After the SQL statements pass the precheck, click **Submit for Approval**. In the message that appears, click **OK**.

5. After your ticket is approved, go to the **Workbench** tab and click **Submitted Tickets** in the My Tickets section.
6. Find the data export ticket that is submitted and click the ticket number.
7. In the **Execute/Automatic Execution** section, click **Download Exported File**.

15.7.4. Generate test data

DMS provides the test data generation feature that allows you to generate data in a quick manner. You can generate test data for functional or performance tests.

Prerequisites

- A relational database is created to store test data that you generate. The relational database can be a self-managed MySQL, ApsaraDB RDS for MySQL, AnalyticDB for MySQL, or PolarDB-X database.
- A table is created. You can use the schema design feature to create a table. For more information, see [Design a schema](#).

Context

Functional tests or performance tests often require test data. You can use one of the following methods to generate test data:

- Write test data. This method has low efficiency and is not applicable to scenarios in which a large amount of test data is required.
- Use scripts. This method requires high costs and the data that is generated by using this method cannot meet

discreteness requirements.

- Export data from a production environment as test data. This method is not secure and may cause data leak.

DMS provides the test data generation feature that allows you to generate test data in a quick, efficient, and secure manner. You can use this feature to control the discreteness of the data that is generated.

Precautions

- You can use this feature to generate test data in only one table at a time. To generate test data in multiple tables, submit a ticket.
- To prevent database overload due to the instantaneous generation of excessive data, DMS allows you to perform traffic throttling. Check the following examples for your reference.
 - About 1 minute is required to generate one million rows of data in a table that has four fields.
 - About 2 to 3 minutes are required to generate one million rows of data in a table that has 40 fields.

Procedure

1. [Log on to the DMS console.](#)
2. In the top navigation bar, choose **More > Data Plans > Test Data Generation.**
3. In the upper-right corner, click **Test Data Generation.**
4. In the Test data build ticket application dialog box, set the required parameters.

| Parameter | Description |
|------------------|---|
| Task Name | The name of the task. The name can help you identify and manage the task in the future. |

| Parameter | Description |
|---------------------------------|---|
| Database Name | The name of the database that includes the required table. |
| Table Name | <p>The name of the table in which you want to generate test data. You can search for tables by keyword and then select a table from the matched results. After you select a table, the Configure the algorithm parameter that contains the field information of the table is displayed.</p> <p>Note You can use this feature to generate test data only in one table at a time. To generate test data in multiple tables, submit a ticket.</p> |
| Configure the algorithm | <p>The algorithms that are used to generate test data. Find the required field and click the value of the Generation mode parameter next to the field. Then, set the required parameters in Generation mode dialog box based on your business requirements.</p> <p>Note For example, you can use the random, customize, or enumeration algorithm to generate test data of the STRING type. The customize algorithm can be used to generate multiple industry-standard types of data.</p> |
| Number of rows generated | The number of data rows that you want to generate. |
| Conflict Handling | The method that is used to handle conflicts based on your business requirements. |
| Change Stakeholder | The stakeholders of the ticket. All specified stakeholders can view the details about the ticket and are included in the approval process. Unauthorized users, except for DMS administrators and DBAs, cannot view the details about the ticket. |

- After you configure the settings, click **Submit**.
- After your ticket is approved, DMS generates test data.

15.7.5. Clone databases

The database clone feature allows you to replicate data at the database level. This topic describes how to use the database clone feature.

Prerequisites

- A MySQL database is used.
- A database instance is managed in flexible management mode. You have logged on to the database instance in the Data Management (DMS) console.

Scenarios

- Create a full database backup.
- Initialize databases that are deployed in different environments, such as development and test environments.
- Copy data from a database in an online environment to a database in an offline environment for data processing and analysis.

Procedure

- [Log on to the DMS console](#).
- In the top navigation bar, choose **More > Data Plans > Database Clone**.
- In the upper-right corner, click **Database Clone**.

4. In the Apply step, set the required parameters.

| Parameter | Description |
|---|--|
| Task Name | The name of the task. The name is used to identify and manage the task. |
| Source database (Only support MySQL) | The source database that you want to clone. You can search for databases by keyword and then select a database from the matched results. |
| Target database (Only support MySQL) | The destination database to which you want to write the data that is cloned from the source database. You can search for databases by keyword and then select a database from the matched results. <div style="background-color: #e0f2f1; padding: 5px;"> <p>Note The destination database must be different from the source database.</p> </div> |
| Select source table | The tables that you want to clone from the source database. You can search for tables by keyword and then select a table from the matched results. <div style="background-color: #e0f2f1; padding: 5px;"> <p>Note To clone all tables, set this parameter to All Tables.</p> </div> |
| Duplicate objects | The method that is used to handle object conflicts based on your business requirements. <ul style="list-style-type: none"> Skip duplicate name object: The system skips objects that have a duplicate name. Overwrite duplicate name object (warning: the structure and data of the target object will be replaced): If the two objects have the same name, the schema and data of the object in the destination database are overwritten by the schema and data of the object in the source database. |

| Parameter | Description |
|-------------------|---|
| Migration Objects | The objects that you want to clone. In addition to tables, you can simultaneously clone other objects from the source database to the destination database. These objects include views, stored procedures, functions, triggers, and events. |
| Time options | Valid values: Running immediately and Specified time . If you set the Time options parameter to Specified time , you must specify a date and time to run the task. <ul style="list-style-type: none"> ◦ Running immediately: The task is run immediately after the ticket is approved. ◦ Specified time: DMS automatically runs the task to clone data at a specified point in time. |

5. After you configure the settings, click **Submit**.
6. After the ticket is approved, the task is automatically run at a specified point in time.

15.8. Data factory

15.8.1. Task orchestration

The task orchestration feature runs based on a distributed scheduling engine that is developed by Alibaba Cloud. This feature allows you to create task flows and schedule these task flows. This feature also meets requirements for data processing based on powerful task flows that include various types of task nodes. For example, you can use this feature to archive data, integrate data, and transform data.

Scenarios

- Periodically archive and analyze business data.
- Synchronize online data to data warehouses for complex analysis.
- Periodically clean and transform offline data.
- Periodically schedule and orchestrate data definition language (DDL) or data manipulation language (DML) operations on databases.

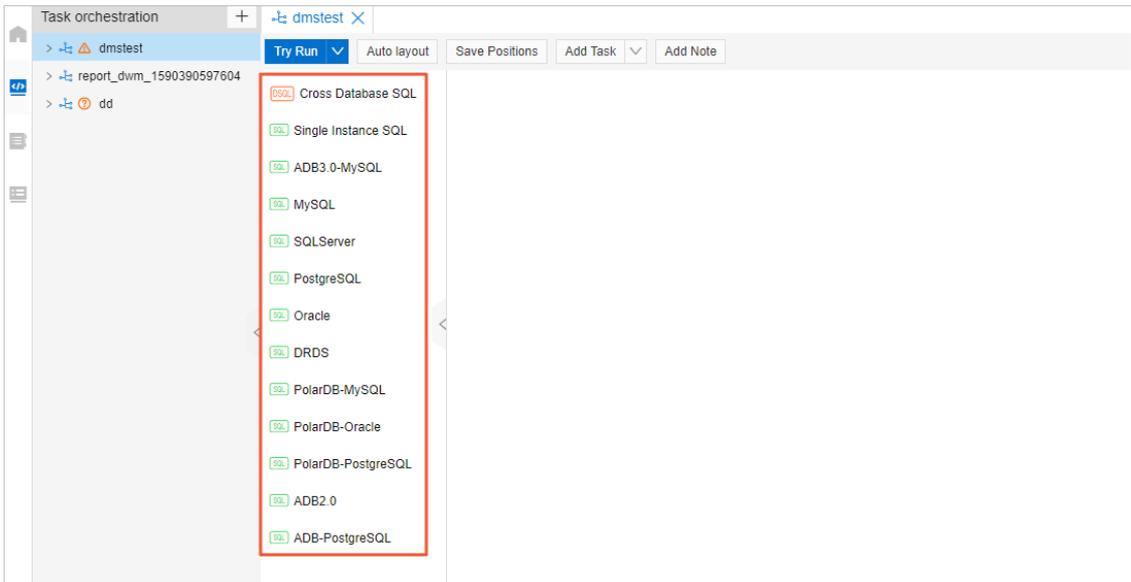
Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **More > Data Factory > Task Orchestration**.
3. In the left-side navigation pane, click the  icon.
4. Configure a task flow.
 - i. Click the  icon next to **Task Orchestration**.
 - ii. In the dialog box that appears, set the Task Flow Name and Description parameters.
 - iii. Click **OK**.
5. Configure a task in the task flow.

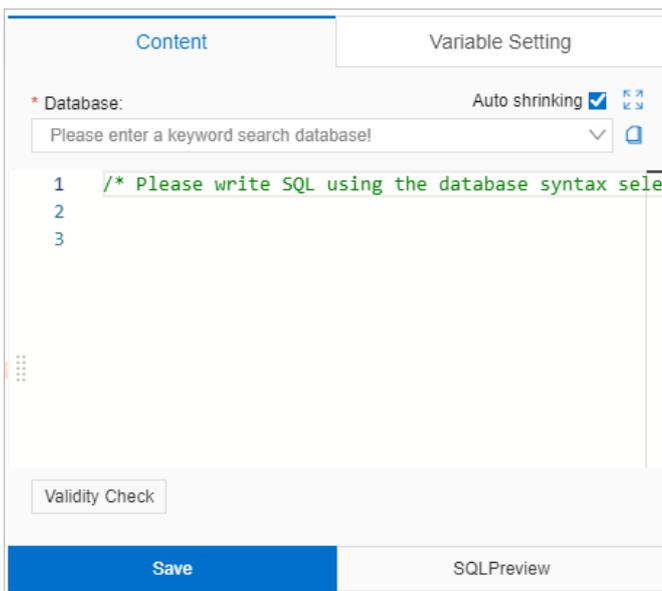
 **Note** You can repeat this step to configure more tasks for the task flow.

- i. In the left-side navigation pane of the Task Orchestration tab, find the new task flow and double-click the task flow name.

- ii. On the Task Orchestration tab, drag a task node from the task node list to the canvas. On the canvas, you can create a directed acyclic graph (DAG) based on your business requirements.



- iii. In the DAG, click the required task node.
- iv. In the right-side pane, click a tab and set the required parameters.



| Tab | Parameter | Description |
|-----|-----------|-------------|
|-----|-----------|-------------|

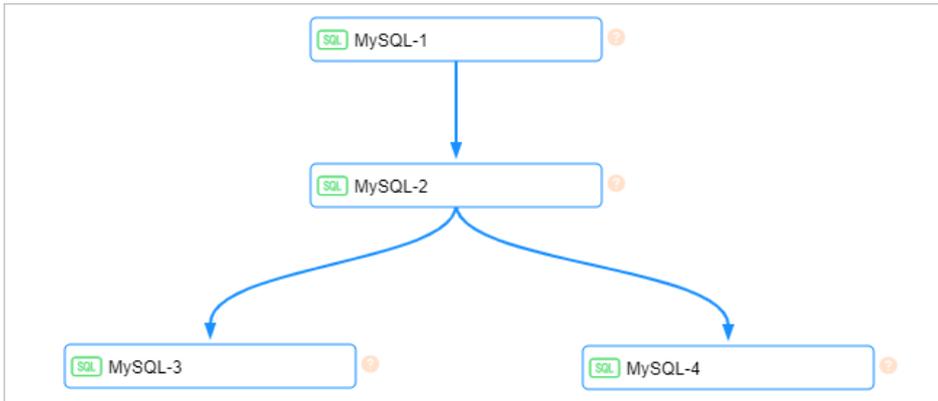
| Tab | Parameter | Description |
|------------------|---------------|---|
| Metadata | Database | <p>a. Select the database that you want to manage from the drop-down list.</p> <p>b. Enter the SQL statements to be executed in the field.</p> <p>c. Click Save.</p> <p>d. In the dialog box that appears, select Existed table or New table to store the query results.</p> <div style="border: 1px solid #add8e6; padding: 5px; margin: 10px 0;"> <p>Note If you select Existed table, you must select a name for the required table from the drop-down list. If you select New table, you need to enter the name of a new table.</p> </div> <p>e. Click OK.</p> |
| Variable Setting | Variable Name | <p>bizdate is the only default system variable. This variable indicates the day before the task is executed. The value of the bizdate variable is in the yyyy-MM-dd format.</p> <p>If the default system variable cannot meet your business requirements, you can create a custom variable and enter the variable name in the Variable Name field.</p> <div style="border: 1px solid #add8e6; padding: 5px; margin: 10px 0;"> <p>Note To delete a variable that you set, click the  icon.</p> </div> |
| | Variable Rule | <p>To configure a variable rule, set the Time Format parameter. Then, set the required operator, integer value, and time unit, and click Save. To create more variables, you can click Increase Variable.</p> <div style="border: 1px solid #add8e6; padding: 5px; margin: 10px 0;"> <p>Note After you configure a variable, you can reference the variable in the <code>\${Variable name}</code> format in SQL statements on the Content tab. You can also click SQLPreview to check whether the variable is configured.</p> </div> |

6. Configure node dependencies in the task flow.

- i. In the DAG, move the pointer over a task node and click and hold the circle, as shown in the following figure. Then, draw a line from the circle to the next node.



- ii. Repeat the previous step to create dependencies between task nodes based on your business requirements. These dependencies indicate the execution order of each task.



Note In the preceding figure, the MySQL-1 node is the first task to be executed and the MySQL-2 node is the second. After that, the MySQL-3 and MySQL-4 nodes are executed at the same time.

7. Specify scheduling information for the task flow.

- i. Click the blank area on the canvas. In the right-side pane, click the **Scheduling** tab.

Note You can also click the **Properties** tab to configure the basic information of the task flow or click the **Operations** tab to view the historical operations that are performed on the task flow.

ii. Specify the scheduling information.

| Parameter | Description |
|--------------|--|
| Turn on/off | Turn on the Turn on/off switch to enable scheduling. |
| Trigger type | <p>Set this parameter based on your business requirements.</p> <ul style="list-style-type: none"> ▪ If you set this parameter to Cyclic scheduling, you must set the Effective Time, Scheduling cycle, and Specific Time parameters. ▪ If you set this parameter to Schedule once, you need to set only the Specific Time parameter. |

iii. Click **Save**.

- After the preceding parameters are configured, click **Try Run** in the upper-left corner of the tab to check whether the task flow can be run as expected.
- In the left-side navigation pane, click the  icon, specify filter conditions, and then view the status of the task flow.

| Name | Status | Trigger Type | Owner | Business Date | Starting Time | End Time | Operating |
|-------------|--------|----------------|-------|---------------------|---------------------|---------------------|----------------------------|
| dmstest(34) | Fail | Manual Trigger | | 2020-06-16 05:19:21 | 2020-06-17 05:19:21 | 2020-06-17 05:19:21 | DAG Executive History More |

Note Click the icon next to the task flow name to view the status of each task.

You can perform the following operations on the task flow:

- **DAG:** View the DAG of the task flow.
- **Executive History:** View the details of the historical operations that are performed on the task flow.
- **More > Exits:** Stop a task flow that is in the running state.
- **More > Rerun:** Rerun a task flow that is executed or fails to be executed.
- **More > Pause:** Pause a task flow.
- **More > Restore:** Resume a paused task flow.
- **More > Set Successfully :** Set the status of a task flow that fails to be executed to Success.

15.8.2. Data warehouse development

15.8.2.1. Overview

Data Management (DMS) provides the data warehouse development feature. This feature uses databases as the computing engine and integrates a variety of tools and services in the database ecosystem. This allows you to develop and manage data warehouses with ease. This feature is designed to provide you with a one-stop development platform for data integration, processing, visualization, and value mining.

Benefits

- A variety of data warehouse engine types

You can choose a data warehouse engine type based on your enterprise scale, data volume, and requirement for real-time performance. For example, you can choose AnalyticDB for MySQL or ApsaraDB RDS for MySQL as the data warehouse engine type.

- Two development modes

The data warehouse development feature of DMS provides two development modes: task orchestration and professional development. The two modes meet different business requirements.

- **Task orchestration:** This development mode allows you to develop a data warehouse by creating task flows and writing SQL scripts for task nodes. You do not need expertise in data warehouse development. You need only to focus on your business logic.
- **Professional development:** This development mode meets the requirements of professional warehouse developers. It provides capabilities such as theme management, hierarchical management, production, release, multi-person collaboration, and data quality control. These capabilities empower professional warehouse development solutions for your enterprise.

Note Some of the capabilities are planned to be supported soon.

- Support for offline and real-time data warehouses

The data warehouse development feature supports offline data synchronization and task scheduling. This allows you to develop offline data warehouses with ease in DMS. In addition, DMS is integrated with Data Transmission Service (DTS) and cloud-native data warehouses. This allows you to build a real-time data warehouse system based on the real-time synchronization feature of DTS and cloud-native data warehouse engines. Then, you can develop data and consume data in real time in DMS.

- Unified management of online and offline data

DMS supports unified database management and permission management. You can manage your online transaction processing (OLTP) databases and online analytical processing (OLAP) databases in a centralized manner in DMS. This avoids security issues that are caused by the isolation between offline and online systems.

15.8.2.2. Create a data warehouse project

Before you can use the data warehouse development feature of Data Management (DMS), you must create a data warehouse project and select a data source for data warehouse development. This topic describes how to create a data warehouse project.

Procedure

1. [Log on to the DMS console.](#)
2. In the top navigation bar, move the pointer over the **More** icon and choose **Data Factory > Data Warehouse Development**.
3. On the left-side navigation submenu, click the  icon.
4. Click the  icon to the right of **Data warehouse**.
5. In the New Warehouse Project dialog box, set the parameters as required.

New Warehouse Project ✕

Basic Information

* Project Name ✔

* Mode ▾

Description

Select Data Development Services

Data integration, data development, data services, and operations management
You can perform data synchronization integration, workflow orchestration, periodic task scheduling, and operations.

Data warehouse engine selection

AnalyticDB for MySQL 3.0 [Go buy](#) AnalyticDB for PostgreSQL [Go buy](#) RDS for MySQL [Go buy](#)
 PolarDB MySQL [Go buy](#)

* Select an existing database ▾ [🔗](#)
Only instances in common mode are supported, and the creator must be the database owner ⓘ

Spark
Waiting.....

| Section | Parameter | Description |
|---|---|---|
| Basic Information | Project Name | The name of the project. Specify a descriptive name for easy identification. |
| | Mode | The mode of the project. Set this parameter to Simple Mode(Single environment) . This way, you can use the same database in a development environment and a production environment. |
| | Description | The description of the project. |
| Select Data Development Services | N/A | DMS automatically completes the configuration in this section. |
| Data warehouse engine selection | Select a type of data warehouse engine. | After you select a data warehouse engine, such as AnalyticDB for MySQL 3.0 , select a database from the Select an existing database drop-down list. <div style="background-color: #e0f2f7; padding: 5px; margin-top: 10px;"> ? ⓘ Note Only the databases of the instances that are managed in Security Collaboration mode are available in the drop-down list. You must be the owner of the selected database. </div> |

6. Click **OK**.

What's next

[Create or import an internal table](#)

15.8.2.3. Create or import an internal table

After you create a data warehouse project in Data Management (DMS), you must create or import an internal table for the project. An internal table refers to a table that exists in the data warehouse engine. This topic describes how to create or import an internal table.

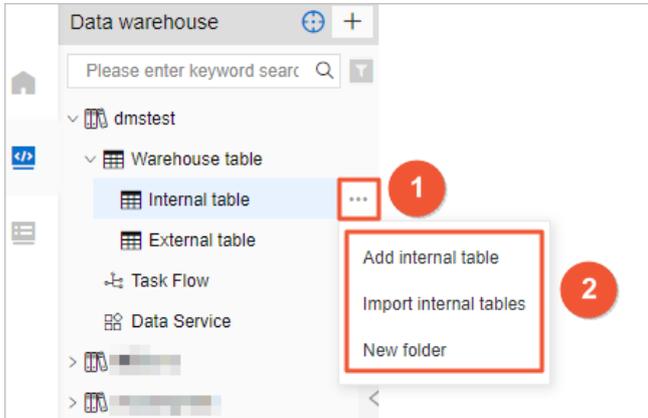
Prerequisites

You have the change permissions on the database for which you want to create or import an internal table. For information about how to apply for permissions, see [Apply for permissions](#).

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **Data Factory > Data Warehouse Development**.
3. On the left-side navigation submenu, click the  icon.
4. In the left-side navigation pane, expand a data warehouse project and then expand Warehouse table. Move the pointer over **Internal table** and then the More icon that appears. Then, select an option from the menu to perform one of the following operations.

 **Note** You cannot configure external tables in DMS.



- o Create an internal table:
 - a. Select **Add internal table**.
 - b. On the tab that appears, enter an SQL statement to create a table.
 - c. Click **Execute**.
- o Import an internal table:

 **Note** The data warehouse development feature does not support real-time synchronization of tables that are created by using other means such as a command-line tool. You can import such tables to data warehouse projects in DMS.

- a. Select **Import internal tables**.

- b. In the Import internal tables dialog box, select the table that you want to import from the **Choose table** drop-down list and enter a description in the Remarks field.
 - c. Click **OK**.
- o Create a folder:

 **Note** If you have a large number of tables, you can use folders to organize and classify the tables.

- a. Select **New folder**.
- b. In the New folder dialog box, enter a folder name.
- c. Click **OK**.

What's next

[Manage task flows](#)

15.8.2.4. Manage task flows

Data Management (DMS) supports task flows and timed scheduling. You can configure a variety of task nodes in task flows. This can meet your requirements for data archiving, integration, and processing.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **Data Factory > Data Warehouse Development**.
3. On the left-side navigation submenu, click the  icon.
4. In the left-side navigation pane, expand a data warehouse project. Move the pointer over **Task Flow** and then the  icon that appears.
5. In the New Task Flow dialog box, enter a name and the description for the task flow.
6. Click **OK**.
7. On the tab that appears, configure task nodes for the task flow.

 **Note** The configurations of a task flow in professional development mode are basically the same as the configurations of a task flow in task orchestration mode. For more information, see [Step 5](#) in the *Task orchestration* topic.

15.8.2.5. Use the data service feature

In Data Management (DMS), the data warehouse development feature is integrated with the data service feature. The data service feature allows you to export the data that is managed by DMS. This feature is applicable to scenarios in which you need to export data at the column or row level, visualize data, or perform complex analysis.

Limits

When you use the data service feature to create an API for a data warehouse, the data source of the API must be a table in the data warehouse project.

Procedure

1. [Log on to the DMS console](#).

- In the top navigation bar, choose **More > Data Factory > Data Warehouse Development**.

Note You can also choose **More > Data Factory > Data Service**.

- On the left-side navigation submenu, click the  icon.
- In the left-side navigation pane, expand the required data warehouse project and double-click **Data Service**.

Configure an API

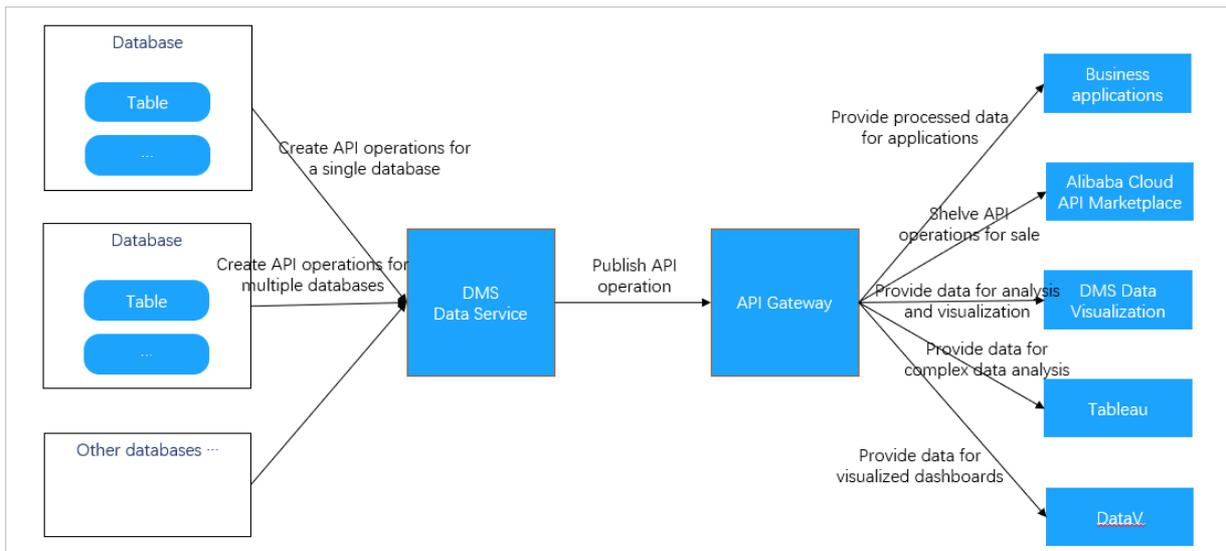
For information about how to configure an API, see [Develop an API](#), [Unpublish or test an API](#), [Test an API](#), and [Call an API](#).

15.8.3. Data service

15.8.3.1. Overview

Data Management (DMS) provides the data service feature, which allows you to export the data that is managed by DMS. This feature is applicable to scenarios where you need to export data at the column or row level, display data in a visualized manner, or perform complex analysis.

Features



- You can use the data service feature to create APIs that can be called to access the data that is managed by DMS. When you create the APIs, you can apply the security control features that are used for SQL execution in the SQLConsole, such as permission control and data de-identification.
- The data service feature works based on a serverless architecture. This feature frees you from the concern about the infrastructure of the runtime environment, such as servers and networks. You need to focus only on how to create APIs and design data query logic. This avoids operations and maintenance (O&M) overheads that are generated by using traditional architectures.
- The data service feature is fully integrated with API Gateway. You can use this feature to publish APIs to API Gateway. This way, you can use all the features that are provided by API Gateway, such as API permission control, IP address-based access control, throttling, metering and billing, and SDKs.

Scenarios

| Scenario | Description |
|--|---|
| Minimize data exposure | Assume that you need to export the data that is managed by DMS to an external environment. In this case, APIs can be called to export the data of specific rows or columns to the external environment. To export the data of specific rows, specify a filter condition in the SQL statement. To export the data of specific columns, specify the columns in the SQL statement. Compared with data export of a whole table, this minimizes data exposure and ensures data security. |
| Connect visualization tools to databases | Most visualization tools can connect to databases by calling APIs. You can connect a visualization tool to your database by calling an API, instead of by using a username and a password. This method is easy to implement and avoids account exposure. |
| Sell APIs in the Alibaba Cloud Marketplace | If you want to provide paid or free data for other users, publish an API to the Alibaba Cloud Marketplace. |
| Provide processed data for applications | After data is processed and summarized by using the data warehouse development feature of DMS, APIs can be created and provided for applications to read the processed data from DMS to meet business needs. To modify the logic of data reading, you need only to modify the query logic of the required API without the need to republish the application. |

15.8.3.2. Develop an API

The data service feature of Data Management (DMS) allows you to develop APIs with ease. This topic describes how to create and manage APIs.

Prerequisites

API Gateway is activated. For more information, see the documentation of *API Gateway*.

Context

The data service feature allows you to export the data that is managed by DMS. This feature is applicable to various scenarios. These scenarios include data export at the column or row level, data visualization, or complex data analysis. For more information, see [Overview](#).

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **More > Data Factory > Data Service**.
3. In the left-side navigation pane, click the **API Development** tab.
4. On the APIManagement tab, click **New API** in the upper-right corner.
5. On the tab that appears, set the required parameters.

i. Set the required parameters on the AttributeConfiguration tab.

| Parameter | Description |
|--|--|
| APIName | The name of the API. The name must be 4 to 100 characters in length, and can contain letters, digits, and underscores (_). The name must start with a letter. |
| Description | Optional. The description of the API. Enter an informative description, for example, a description of the data that you want the API to return or the scenarios in which the API can be called. |
| Path | <p>The path of the API. The path must start with a forward slash (/) and can contain letters, digits, underscores (-), and hyphens (-).</p> <p>The specified path forms a part of the URL that is used to call the API. A URL that is used to call an API must be in the <code>https://{Domain name}{Path}</code> format. For example, if the domain name is <code>xxx-cn-hangzhou.alicloudapi.com</code> and the path is <code>/item/monthly_data</code>, the URL that is used to call the API is <code>https://xxx-cn-hangzhou.alicloudapi.com/item/monthly_data</code>.</p> |
| ReturnFormat | The format in which you want the API to return data. Valid value: JSON . |
| RequestMode | The request method. Valid values: POST and GET . |
| TimeOut (MS) | The maximum period of time that the system can wait until an API request expires. Unit: milliseconds. If the execution time of an API exceeds the specified timeout period, the system returns a timeout error. Maximum value: 30000. |
| Returns the maximum number of records | <p>The maximum number of entries that can be returned for an API request. This parameter limits the number of entries that can be returned for each API query.</p> <div style="background-color: #e6f2ff; padding: 5px;"> <p> Note If the database instance is managed in security collaboration mode, the value of this parameter must be less than the maximum number of entries that is specified in the security rules.</p> </div> |

ii. Click the **ExecuteConfiguration** tab and set the required parameters.

| Parameter | Description |
|---------------------|--|
| Instance query type | <ul style="list-style-type: none"> ▪ Single InstanceQuery: You can call the API to read data from only one database instance. ▪ Cross-instanceQuery: You can write dynamic SQL statements for the API to query data across multiple database instances. <p>Note If you set this parameter to Cross-instanceQuery, you need only to enter dynamic SQL statements in the QuerySQL field.</p> |
| Data source | The database that is queried by the API. You can search for databases on which you have query permissions by keyword and then select a database. |
| ConfigurationMode | <ul style="list-style-type: none"> ▪ Table boot mode: You can configure data query by selecting a table and fields. ▪ Script mode: You must configure a data query by specifying variables and writing SQL statements. <p>Note After you set this parameter to Script mode, you need only to enter SQL statements in the QuerySQL field.</p> |
| SelectTable | The table to be queried. You can search for tables by keyword. |
| FieldList | The fields in the selected table. You can specify the required fields as request parameters or response parameters. |
| Script mode | <p>The mode in which an SQL script is written to define the data query logic.</p> <p>Note You can set the ConfigurationMode parameter to Table boot mode or Script mode.</p> |
| QuerySQL | <p>The SQL statement that is used to query the data in the table. After you enter an SQL statement, click ParsingScript to verify the syntax and to parse the request parameters and response parameters.</p> <p>Note</p> <ul style="list-style-type: none"> ▪ Custom variables are supported. Custom variables can be mapped as request parameters in API requests. The variables of an SQL statement must be specified in the <code>{Variable name}</code> format. For example, the <code>select item_id, item_name from ex_item where category=\${category}</code> SQL statement include a variable named <code>category</code>. ▪ If you set the Instance query type parameter to Cross-instanceQuery, you must use the syntax of cross-database query SQL statements. For more information, see Cross-database query. |

iii. Click the **RequestParameters** tab and set the required parameters.

| Parameter | Description |
|------------------------|---|
| ParametersName | <p>The name of the request parameter.</p> <ul style="list-style-type: none"> ▪ The name can contain letters, digits, hyphens (-), and underscores (_). ▪ The name must start with a letter or an underscore (_). ▪ The name must be 1 to 50 characters in length. |
| FieldName | The name of the field that is specified by the request parameter. The field name is specified on the ExecuteConfiguration tab and cannot be changed. |
| Cannot be empty | Specifies whether the request parameter is required. |
| Description | The description of the request parameter. |
| Data type | <p>The data type of the request parameter. The data type is used to check whether the value of the request parameter in an API request is valid. Valid values: String, Integer, and Floating point. Default value: String.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note This parameter affects the SQL statement that is executed when the API is called.</p> </div> |
| Example value | The sample value of the request parameter. You can use the sample values that are provided in SDKs and documentation as references when you call API operations. |
| Default value | The default value of the request parameter. If the request parameter is optional and not specified in the API request, the default value is used. |

iv. Click the **Return parameter** tab and set the required parameters.

| Parameter | Description |
|-----------------------|---|
| ParametersName | <p>The name of the response parameter.</p> <ul style="list-style-type: none"> ▪ The name can contain letters, digits, hyphens (-), and underscores (_). ▪ The name must start with a letter or an underscore (_). ▪ The name must be 1 to 50 characters in length. |
| FieldName | The name of the field that is returned. The name cannot be changed. |
| Description | The description of the response parameter. |
| Data type | The data type of the response parameter. Valid values: String, Integer, and Floating point. Default value: String. This parameter is used by DMS to convert the type of the data in API responses. This parameter affects the JSON data that is returned. |
| Example value | The sample value of the response parameter. You can use the sample values that are provided in SDKs and documentation as references to help you understand API responses. |

6. Click **Save**.

7. In the left-side navigation pane, click the **API Development** tab.

8. Perform the following operations to manage the API based on your business requirements:

- Publish the API
On the APIManagement tab, find the required API and click **Publish** in the **Operation** column of the API. In the message that appears, click **OK**.
- Modify the API
On the APIManagement tab, find the required API and click **Modify** in the **Operation** column of the API. Modify the configurations of the API based on the descriptions in [Step 5](#) and click **Save**.
- Delete the API:
On the APIManagement tab, find the required API and click **Delete** in the **Operation** column of the API. In the message that appears, click **OK**.

15.8.3.3. Unpublish or test an API

This topic describes how to unpublish or test an API that has been published.

Prerequisites

An API is created. For more information, see [Develop an API](#).

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **More > Data Factory > Data Service**.
3. Click the **API Publish** tab on the left side.
The **APIPublishList** tab displays all the published APIs.
4. Find the API that you want to manage and perform the following operations based on your business requirements:
 - Unpublish the API:
Click **Offline** in the **Operation** column. In the message that appears, click **OK**.
 - Test the API:
Click **Test** in the **Operation** column. For more information, see [Test an API](#).

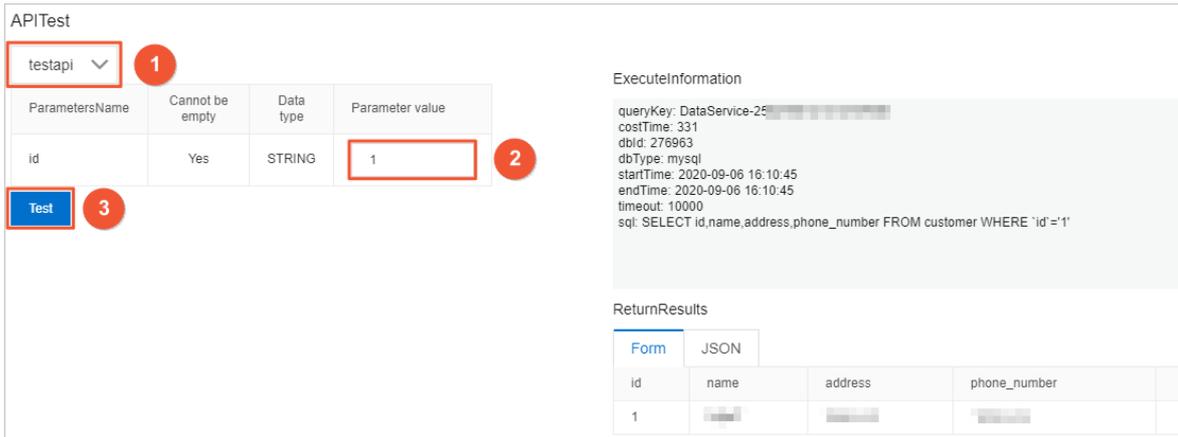
15.8.3.4. Test an API

After you create an API, you can test the API to verify whether the API meets your business requirements.

Prerequisites

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **More > Data Factory > Data Service**.
3. Click the **API Test** tab on the left side.
4. On the **APITest** tab, test an API.



- i. Select the API that you want to test from the drop-down list.
- ii. Enter values in the Parameter value column.
- iii. Click Test.

After the test is complete, the execution information and return results appear on the right side. You can evaluate whether the API meets your business requirements based on the information.

Note You can click the **JSON** tab in the **ReturnResults** section so that the return results are displayed in the JSON format.

15.8.3.5. Call an API

After you create, publish, and test an API, you can call the API in an application by using an SDK.

Prerequisites

- An API is created and published. For more information, see [Develop an API](#).
- API Gateway is activated. For more information, see the documentation of *API Gateway*.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **More > Data Factory > Data Service**.
3. Click the **API Call** tab on the left side.
4. View the API call address and the authentication information.

API calls

API Call Address

Endpoint: [https://\[...\]](#)
 The specific API call address is the path defined by Endpoint + API, such as [https://\[...\]/your_api_path](#)

API call authentication method

| | |
|---|---|
| Authentication Method 1: Simple identity authentication | Authentication Method 2: Encrypted signature identity |
| AppCode: <input type="text" value="d*****f"/> | AppKey: <input type="text" value="*****"/> |
| Display Copy Reset | <input type="text" value="*****"/> Copy |
| For this authentication method, add the AppCode parameter after the API call address. | AppSecret: <input type="text" value="U*****R"/> |

API Call SDK

Note: Please bind an independent domain name to API Gateway. The second-level domain name of API Gateway can only be called up to 1000 times a day. There is no limit on the number of calls after binding an independent domain name.

[Expand](#) ▼

- **Simple identity authentication:** requires only an AppCode. This authentication method is suitable for calling APIs by using URLs. This authentication method has a low security level and is generally used in scenarios in which data visualization is involved, such as calling APIs in DataV.
 - **Encrypted signature identity authentication:** requires an AppKey and an AppSecret, which are used to dynamically generate an encrypted signature for calling an API. This authentication method has a high security level.
5. Call the API in an application by using an SDK.

? **Note** For more information about how to call an API in an application by using an SDK, see the documentation of *API Gateway*.

15.9. Schemas

15.9.1. Schema design

Data Management (DMS) provides the schema design feature. This feature allows you to change schemas with ease. This topic describes how to change schemas.

Prerequisites

The destination database is a MySQL, a PolarDB-X, or an ApsaraDB for OceanBase database.

Context

When you create projects, process new business requirements, or optimize business operations, you may need to change schemas. These schema operations include creating and editing tables. For example, you may need to add or delete fields or indexes, adjust field attributes, or adjust the index composition. In these scenarios, you can use the schema design feature of DMS.

- This feature allows multiple users to simultaneously change a schema in the DMS console at the same time.
- This feature allows you to send verified scripts to other environments. This ensures consistency between schemas in different environments.

Precautions

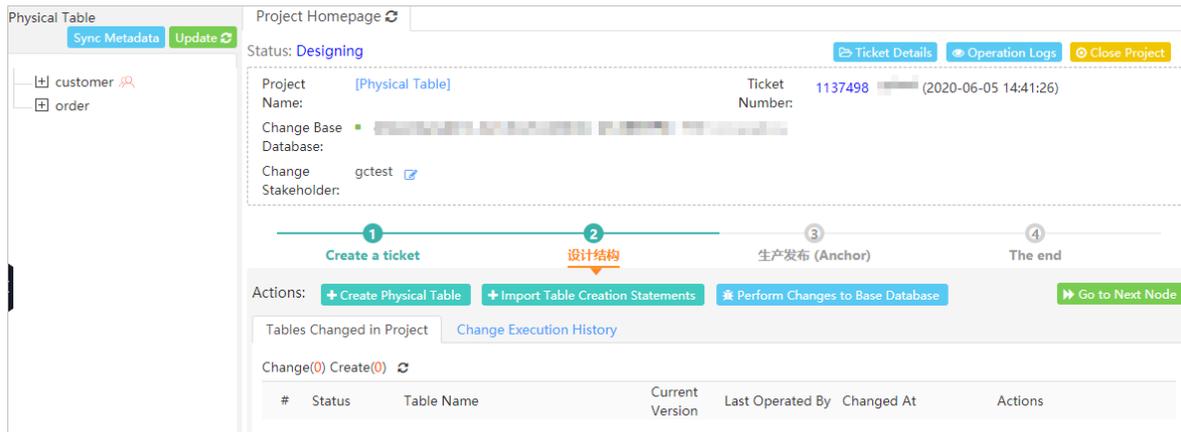
When you submit a schema design ticket to delete a table, make sure that the table is created by using a schema design ticket.

Procedure

1. Log on to the DMS console.
2. In the top navigation bar, choose **More > Schemas > Schema Design**.
3. In the upper-right corner of the Schema Design tab, click **Schema Design**.
4. On the Schema Design tab, specify the required parameters for a schema design ticket.

| Parameter | Description |
|-----------------------------|---|
| Project Name | The name of the project. Specify a name that can help you identify the project. |
| Project description | The background information about the project, such as the purpose or objective of the project. The description is used to reduce communication costs. |
| Change Base Database | The database whose schema you want to change. You can search for databases by keyword. Prefix match is applied. Only databases on which you have permissions in test or development environments are displayed. <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e6f2ff;"> ? Note You must have at least the query, export, or change permissions on the database that you select. </div> |
| Security Rules | No configurations are required. The default setting is specified. |
| Change Stakeholder | The stakeholders of the changes. The specified stakeholders can view the ticket details and are included in the approval process. Unauthorized users, except for administrators and database administrators (DBAs), cannot view the ticket details. |

5. Click **Create Ticket**.
6. Change a schema based on your business requirements.



- o Create a table:

- a. Click **Create Physical Table**.

Note If the destination database is a logical database, click **Create Logical Table**.

- b. On the **Create Physical Table** tab, set the required parameters. The parameters include the table name, character set, fields, and indexes.
- c. Click **Save**.

Note After you click **Save**, DMS verifies the specified information based on design specifications. If the information does not comply with the design specifications, a message appears.

- d. After the information passes the precheck, click **Confirm Changes and Submit to Save**.

- o Change the schema of a table:

- a. In the left-side table list, right-click the name of the required table.
- b. On the menu that appears, select **Design Table**.
- c. Change the schema as required and click **Save**.

Note After you click **Save**, DMS verifies the specified information based on design specifications. If the information does not comply with the design specifications, a message appears.

- d. After the specified information passes the verification, click **Confirm Changes and Submit to Save**.

7. After the schema is changed, click **Perform Changes to Base Database**.
8. In the **Perform Changes to Base Database** dialog box, set the **Execution Strategy** parameter to **Execute Now** or **Schedule**.
9. Click **Submit for Execution** and wait until the ticket is approved.
10. After the ticket is approved, click **Go to Next Node**.

 **Note**

- After the ticket is approved, DMS applies the changes at the specified point in time. If you do not specify the execution time, the changes are automatically applied after the ticket is approved at the last approval node. You can view the execution status and operation logs. After all changes are applied, you can repeat the preceding procedure to change the schema again. If no additional changes are required for the schema, click **Go to Next Node**.
- After the ticket is submitted to the next node, whether you can go back to the previous node is subject to the predefined design specifications.

11. In the **Go to Next Node** message, click **Go to Next Node**.
12. On the **Project Homepage** tab, click **Perform Changes to Target Database**.
13. In the **Perform Changes to Target Database** dialog box, set the **Target Database** and **Execution Strategy** parameters and click **Submit for Execution**.

 **Note** The required database must reside in a production environment.

14. Wait until the ticket is approved and the changes are applied.
15. Click **Go to Next Node**.
The schema design process ends and the ticket is closed.

15.9.2. Schema synchronization

Data Management (DMS) provides the schema synchronization feature. You can use this feature to compare the schemas of two databases, generate a script to synchronize schemas, and then run the script on the destination database. This topic describes the schema synchronization feature and how to synchronize schemas.

Prerequisites

The source databases and destination databases are ApsaraDB for OceanBase or MySQL databases.

Precautions

- You cannot synchronize schemas to a destination database that resides in a production environment.
- The empty database initialization feature allows you to synchronize some or all tables from a physical or logical database.

Scenarios

You can use the schema synchronization feature to synchronize schemas and ensure schema consistency in the following scenarios:

- Synchronize data between a database in a production environment and a database in a test environment.
- Synchronize data between different databases that are deployed in a test environment.
- Synchronize data between different databases that are deployed in a production environment.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **More > Schemas > Schema Synchronization**.
3. Specify the required parameters for a schema synchronization ticket.

Requested Database Table Synchronization Category: **Schema Synchronization** Empty Database Initialization Repair Table Consistency

* Source

Database:

* Target Database:

* Synchronized Partial Tables All Tables

| Table | Seri... | SOURCE table name | Target table name (Do not fill in the same name as t... | Actions |
|-------------|---------|-------------------|---|---------|
| | 1 | customer | customer | Delete |
| + Batch add | | | | |

* Whether to Not Ignore Ignore [What is the result?](#)

Ignore Error:

* Business Background(Remarks):

Submit

| Parameter | Description |
|--------------------------------|---|
| Source Database | The name of the source database from which you want to synchronize schemas. You must have the read permissions on the source database. |
| Target Database | The name of the destination database to which you want to synchronize schemas. You must have the change permissions on the destination database. |
| Synchronized Table | <p>The tables that you want to synchronize. Valid values:</p> <ul style="list-style-type: none"> Partial Tables: Synchronize one or more tables in the source database. You can click Batch Add to add multiple tables. <div style="background-color: #e0f2f1; padding: 5px; border: 1px solid #ccc;"> <p>Note If you do not set this parameter, the names of the destination tables are the same as the names of the source tables.</p> </div> <ul style="list-style-type: none"> All Tables: Synchronize all tables in the source database. |
| Whether to Ignore Error | <ul style="list-style-type: none"> Not Ignore: If an error occurs when SQL scripts are executed in serial mode, the system immediately stops executing the current and remaining SQL scripts. Ignore: If an error occurs when SQL scripts are executed, DMS stops executing the current SQL script and continues to execute the next statement until all remaining SQL scripts are executed. |

4. Click **Submit**. DMS starts to analyze the schemas.
5. Check the comparison results.

Note If the schemas are changed when the system analyzes the schemas, click **Re-analyze** in the Schema Analysis step.

6. Verify the script that is used to synchronize schemas and click **Submit and Synchronize to Target Database**.

Note If the schemas of the source database and destination database are the same, you do not need to submit the script, and the schema synchronization ticket is closed.

15.9.3. Synchronize shadow tables

Data Management (DMS) provides the shadow table synchronization feature to automatically create a shadow table based on the schema of a source table. DMS generates the name of the shadow table by attaching a prefix or suffix to the name of the source table. You can use this feature for end-to-end stress testing.

Procedure

1. Log on to the DMS console.
2. In the top navigation bar, move the pointer over the **More** icon and choose **Schemas > Shadow Table Synchronize**.
3. On the **Table/Database Synchronization Application** page, set the parameters that are described in the following table.

| Parameter | Description |
|------------------------------|---|
| Source Database | The database whose data is to be synchronized. |
| Prefix / Suffix | The prefix or suffix that is used to create a shadow table name. The name can be in the Prefix + Source table name format or Source table name + Suffix format. You can use a custom prefix or suffix as needed. By default, the Prefix + Source table name format is used. Default shadow table name: <code>__test_Source table name</code> . |
| Synchronized Table | The tables whose schemas you want to synchronize. Valid values: <ul style="list-style-type: none"> Partial Tables All Tables |
| Synchronization Policy | The policy that is used for shadow table synchronization. Valid values: <ul style="list-style-type: none"> Synchronize Now: DMS immediately synchronizes the tables after you submit the ticket. In this case, the tables are synchronized only once. Scheduled Synchronization: DMS synchronizes the tables at the specified time on a regular basis. You can use a crontab expression to schedule synchronization based on your requirements. The minimum interval for synchronization is 1 hour. By default, the shadow tables start to be synchronized at 02:00 every day. For more information, see the Crontab expressions section of this topic. |
| Whether to Ignore Error | Specifies whether to skip errors that occur when SQL statements are being executed. Valid values: <ul style="list-style-type: none"> Not Ignore: If an error occurs when SQL statements are being executed, DMS stops executing the current and subsequent SQL statements. Ignore: If an error occurs when SQL statements are being executed, DMS skips the current SQL statement and continues to execute subsequent SQL statements until all remaining statements are executed. |
| Business Background(Remarks) | The business background of the project, such as the purposes and objectives of the project. |

4. Click **Submit**. DMS starts to analyze the schemas.
5. Check the comparison results.

 **Note** If the schemas are changed when the system analyzes the schemas, click **Re-analyze** in the Schema Analysis step.

- Verify the script that is used to synchronize schemas and click **Submit and Synchronize to Target Database**.

 **Note** If the schemas of the source database and destination database are the same, you do not need to submit the script, and the schema synchronization ticket is closed.

Crontab expressions

If you need to schedule the synchronization task to be run in a more precise manner, you can use a crontab expression. The interval for running the task can be specified by using a combination of minutes, hours, days, weeks, or months.

A crontab expression consists of five fields of the NUMERIC type. Valid values of each field:

- Minutes:** 0 to 59 .
- Hours:** 0 to 23 . A value of 0 indicates the midnight.
- Days:** 1 to 31 . A value of this field indicates a specific day of a month.
- Months:** 1 to 12 . A value of 1 indicates January, and a value of 2 indicates February. Similarly, the specific month that is indicated by a specific value can be obtained.
- Weeks:** 1 to 7 . A value of 1 indicates Sunday, and a value of 2 indicates Monday. In other words, the seven week days from Sunday to Saturday are indicated by values 1 to 7.

Usage notes

- Specify the time for running a stress testing task by the day or week. You cannot specify the day and week at the same time. After you specify one of the preceding two values, you must set the other value to `?`. A value of `?` indicates an unspecified value. For example, if you schedule the task to be run on the first and second days of each month, the **Weeks** field must be set to `?`.
- Limit the characters in a crontab expression to English special characters. The special characters can be wildcards such as asterisks (*) and question marks (?).
- Separate multiple values with commas (,).
- Use a hyphen (-) to indicate a value range. For example, if you set the **Days** field to `1-5`, the task is scheduled to be run on the first to fifth days of a month.
- Use a forward slash (/) to indicate an interval for running the task. For example, if you set the **Days** field to `*/2`, the task is scheduled to be run every two days.

Crontab expression examples

- To schedule the task to be run at 23:00 every Saturday and Sunday, use the following crontab expression: `0 23 ? * 7,1`.
- To schedule the task to be run at 09:30 on the fifth, fifteenth, and twenty-fifth days of each month, use the following crontab expression: `30 9 5,15,25 * ?`.
- To schedule the task to be run at 00:00 every two days, use the following crontab expression: `0 0 */2 * ?`.

15.9.4. Initialize empty databases

DMS provides the empty database initialization feature. This feature allows you to compare the schemas of two databases, generate a script that is used to synchronize data from the source database to the destination database, and run the script on the destination database. To use this feature, the destination databases must be empty. This topic describes how to initialize empty databases.

Prerequisites

- The source databases and destination databases are MySQL or ApsaraDB for OceanBase databases.
- The destination databases are empty databases that do not contain tables.

Features

The empty database initialization feature allows you to synchronize some or all tables from a physical or logical database.

Scenario

Synchronize data between databases that are deployed in different regions or units. For example, this feature is applicable to the following scenarios:

- Synchronize data between a database that is deployed in a production environment and a database that is deployed in a test environment.
- Synchronize data between different databases that are deployed in the test environment.
- Synchronize data between different databases that are deployed in the production environment.

Procedure

1. [Log on to the DMS console.](#)
2. In the top navigation bar, choose **More > Schemas > Empty Database Initialization**.
3. On the Table/Database Synchronization Application tab, set the required parameters to create an empty database initialization ticket.

| Parameter | Description |
|-------------------------|--|
| Source Database | The name of the source database from which data is synchronized. You must have the read permissions on the source database. |
| Target Database | The name of the destination database to which data is synchronized. You must have the write permissions on the destination database. <div style="background-color: #e6f2ff; padding: 5px; border: 1px solid #d9e1f2;"> ? Note The type of the destination database must be the same as the type of the source database. </div> |
| Initialized Table | The tables that you want to synchronize. Valid values: <ul style="list-style-type: none"> ◦ Partial Tables: Synchronize one or more tables in the source database. To add a table, click the + icon and specify a name for the source table. <div style="background-color: #e6f2ff; padding: 5px; border: 1px solid #d9e1f2;"> ? Note <ul style="list-style-type: none"> ▪ You can also click Batch Add. In the Batch Add dialog box, select the required tables and click Batch Add. ▪ If you do not set this parameter, the names of the destination tables are the same as the names of the source tables. </div> ◦ All Tables: Synchronize all tables in the source database. |
| Whether to Ignore Error | <ul style="list-style-type: none"> ◦ Not Ignore: If an error occurs when an SQL script is being executed in serial mode, the system immediately stops executing the current and remaining SQL scripts. ◦ Ignore: If an error occurs when an SQL script is being executed, DMS stops executing the current SQL script and continue to execute the next statement until all remaining SQL scripts are executed. |

4. Click **Submit**. DMS starts to analyze the schemas.

5. Check the comparison results.

 **Note** If the schemas are changed when the system analyzes the schemas, click Re-analyze in the Schema Analysis step.

6. Verify the script that is used to synchronize schemas and click **Submit and Synchronize to Target Database**.

 **Note** If the schemas of the source database and destination database are the same, you do not need to submit the script, and the schema synchronization ticket is closed.

15.9.5. Repair table consistency

DMS provides the table consistency repairing feature. This feature is used to compare schemas between tables in databases that are deployed in different environments, provides an efficient way to identify schema differences, and execute SQL statements that are specific to the required environment. This ensures schema consistency between different environments.

Prerequisites

The source databases and destination databases are MySQL or ApsaraDB for OceanBase databases.

Scenarios

- Ensure the schema consistency between physical tables that are deployed in the test environment and the production environment.
- Ensure the schema consistency between physical tables and logical tables in a physical database or a logical database.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **More > Schemas > Table Consistency Repairing**.
3. On the Table/Database Synchronization Application tab, set the required parameters to create a Repair Table Consistency ticket.

| Parameter | Description |
|---|--|
| Base Database(Physical Database) | The source database based on which schema consistency is to be repaired. You must have the query permissions on the source database. |
| Target Database | The destination database whose data is to be modified. You must have the change permissions on the destination database. |
| Repaired Table | The tables between which schema consistency is to be repaired. To add tables, click the + icon and specify the required table names.  Note If you do not specify the destination table name, the system names the destination table after the name of the specified source table. |

| Parameter | Description |
|-------------------------|--|
| Whether to Ignore Error | <ul style="list-style-type: none"> ◦ Not Ignore: If an error occurs when SQL statements are being executed in serial mode, the system immediately stops executing the current and remaining SQL statements. ◦ Ignore: If an error occurs when DMS is executing an SQL statement, DMS stops executing the current SQL statement and continues to execute the remaining SQL statement. |
| Business Background | The business background of the ticket. This parameter reduces communication costs. |

4. Click **Submit**. DMS starts to analyze the schemas.
5. Check the comparison results.

 **Note** If the schemas are changed when the system analyzes the schemas, click **Re-analyze** in the Schema Analysis step.

6. Verify the script that is used to synchronize schemas and click **Submit and Synchronize to Target Database**.

 **Note** If the schemas of the source database and destination database are the same, you do not need to submit the script, and the schema synchronization ticket is closed.

15.10. SQL review

DMS provides the SQL review feature to help you prevent SQL statements that do not use indexes or do not conform to database development standards. This reduces the risk of SQL injection attacks.

Prerequisites

The environment type of the database instance in which you want to use the SQL review feature is **Test** in the DMS console. This is because SQL review is performed before DMS publishes SQL statements to an online environment.

Context

When you develop a project, you need to execute SQL statements on databases to add, delete, modify, and query data so that you can implement business logic and display data. Before the project is published, you must review all SQL statements that are used. This prevents SQL statements that do not conform to database development standards from being published to an online environment and accordingly impeding the business.

If DBAs manually review all SQL statements one by one, excessive human resources are consumed and the R&D efficiency is low. The SQL review feature helps you review SQL statements and also provides optimization suggestions.

Usage notes

- Only XML or TXT files can be uploaded.
- Tables that are involved in SQL statements must exist in the database that you select. Otherwise, DMS cannot review these SQL statements.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **Optimization > SQL Review**.
3. On the tab that appears, click **New** in the upper-right corner.

4. Configure information about an SQL review ticket.

| Parameter | Description |
|-----------------------------|--|
| Project name | Enter a project name based on your business requirements so that the ticket can be distinguished from others in subsequent processing. |
| Data source | Select the database in the test environment that is used in your project. You must have the change permission on the database. |
| Business Description | Enter detailed information about the business scope of the project as required to help relevant users know about the project. |
| Relevant personnel | Enter an at sign (@) and select a user.  Note You can repeat this operation to select multiple users. |
| Upload a file | Click Add , select files, and then click Upload .  Note <ul style="list-style-type: none"> ◦ The iBatis and MyBatis files are in the XML format. ◦ SQL statements are saved as TXT files. Multiple SQL statements are separated by semicolons (;). ◦ To remove an added file, you can select the check box before the file name and click Delete. |

5. Click **Submit application**.

6. View SQL review results.

 **Note**

- If SQL statements in a file conform to database development standards and use indexes, DMS determines that these SQL statements pass the SQL review and offers no suggestion about indexes.
- If SQL statements in a file conform to database development standards but do not use indexes, DMS determines that these SQL statements pass the SQL review and offers suggestions about indexes.
- If SQL statements in a file do not conform to database development standards, DMS determines that these SQL statements fail the SQL review.

7. Find a failed SQL review result and click **View reason** to check the reason. You can also click **Details**, **Adjust SQL**, or **More** in the **Operation** column to perform other operations.

 **Note** After you optimize SQL statements in a file and click **Confirm**, DMS reviews the SQL statements again. For dynamic SQL statements in XML files, you must optimize each SQL statement combination.

8. When all SQL statements pass the SQL review, click **Inspection results**.9. In the dialog box that appears, click **Submit for approval** and wait for approval.

Note The approval process of the ticket is based on the security rules that are configured for the current database instance.

15.11. System management

15.11.1. Manage instances

Data Management (DMS) allows you to manage database instances. For example, you can export the information about instance configurations.

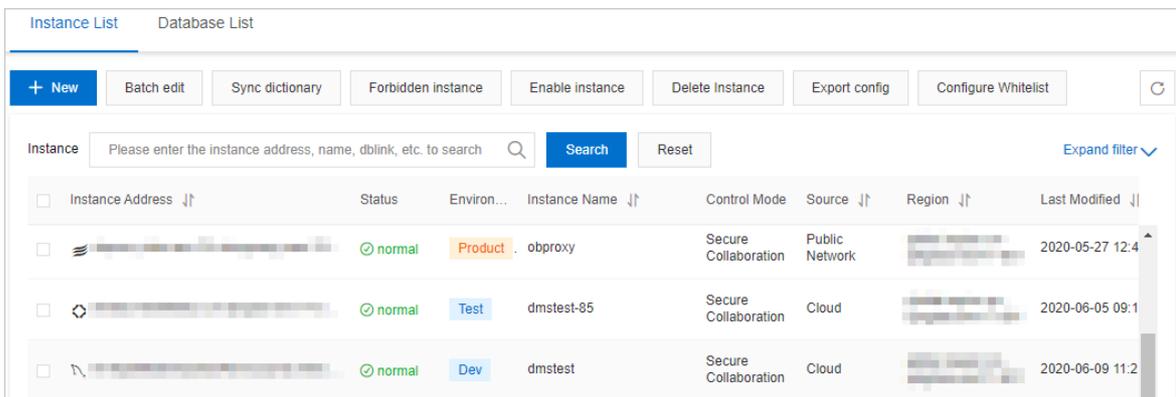
Prerequisites

You are a database administrator (DBA) or a DMS administrator.

Procedure

1. [Log on to the DMS console.](#)
2. In the top navigation bar, move the pointer over the **More** icon and choose **System > Instance**.
3. On the **Instance List** tab, select one or more database instances that you want to manage. Then, you can perform the following operations based on your business requirements:

Note You can click **Expand filter** to show more filter conditions.



- o Add an instance
Click **New** and register a database instance with DMS. For more information, see [Register database instances with DMS.](#)

- o Edit multiple instances at a time
Click **Batch edit**. In the dialog box that appears, modify the instance information and click **OK**.

Note The database instances that you select must be of the same database type, such as MySQL.

- o Synchronize the data dictionary
Click **Sync dictionary**. In the message that appears, click **OK**.

Note

- If you change schemas for a database instance by using DMS, DMS automatically synchronizes the data dictionary of the instance.
- If you change schemas for a database instance by using a service other than DMS, you must manually synchronize the data dictionary of the instance.

- Disable or enable one or more instances

Click **Forbidden instance** or **Enable instance**. In the message that appears, click **OK**.

Note

- After you disable a database instance, the instance is removed from the left-side instance list. DMS users can no longer find databases or tables in this instance in the DMS console.
- After you enable a database instance, the instance appears in the left-side instance list. Databases in this instance become available. Relevant permissions that have been granted to DMS users on this instance also become valid.

- Remove one or more instances

Click **Delete Instance**. In the message that appears, click **OK**. After you remove a database instance, the instance is removed from the left-side instance list. DMS users can no longer use databases in this instance in the DMS console. Relevant permissions that have been granted to DMS users on this instance also become invalid and are revoked.

Note On the **Instance List** tab, you can find database instances in the Delete state and enable these instances to recover them.

- Export configuration information

Click **Export config**. The browser automatically downloads a CSV file named *instances*. You can use Excel or a text editor to view this file.

- Configure a whitelist

Click **Configure Whitelist**. In the message that appears, click **OK**. The IP addresses of DMS servers are automatically added to the whitelists of the specified database instances.

Note The destination database instances must be ApsaraDB instances.

- Other operations

You can find a database instance and click **Details** in the **Actions** column to view the details about databases and tables in this instance. You can also move the pointer over **More** and perform other operations. For example, you can log on to the instance or edit the instance.

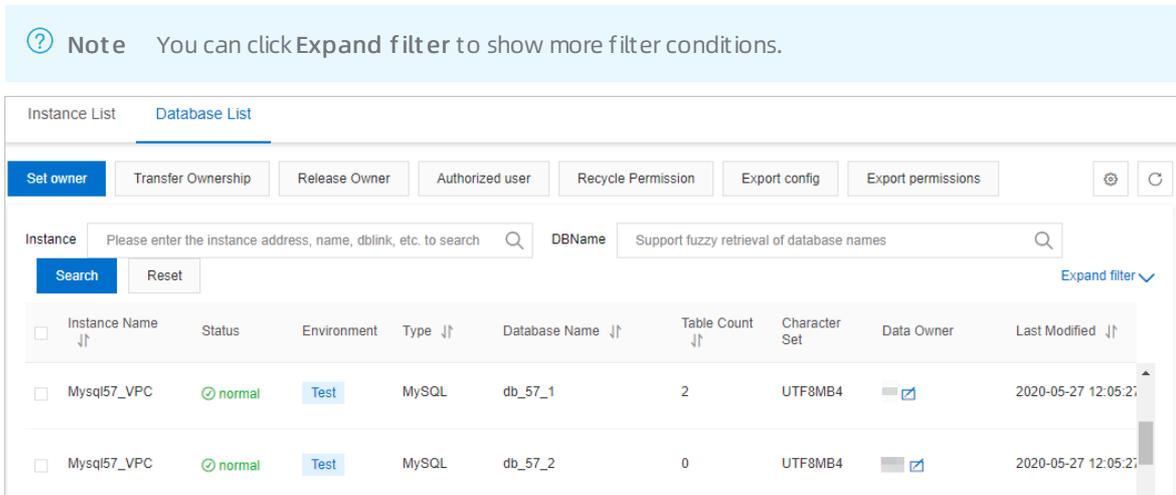
15.11.2. Database management

On the **Database List** tab, you can manage databases. For example, you can specify the database owner, transfer the ownership, revoke the owner permission, grant and revoke user permissions, and export the information about database configurations or permissions.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, move the pointer over the **More** icon and choose **System > Instance**.
3. In the top navigation bar, choose **System > Instance**.

4. Click the **Database List** tab.
5. Set filter conditions and select one or more databases that you want to manage. Then, you can perform the following operations based on your business requirements:



- o Specify an owner

Specify an owner for the selected databases. You can specify multiple owners for multiple databases at a time.
- o Transfer the ownership

Transfer the ownership of the selected databases to a user. If you transfer the ownership of multiple databases at a time, you can select only a user that assumes the ownership of all of the databases as the original owner.
- o Revoke owner permissions

Revoke the owner permission from the owners of the selected databases.
- o Grant permissions

Grant the query, export, or change permission on the selected databases to one or more users. You can also specify an expiration time for the permission.
- o Revoke permissions

Revoke the query, export, or change permission on the selected databases from one or more users. If a user does not have the related permissions, the following message appears: **No corresponding permissions. You do not need to recycle or release permissions**.
- o Export configurations

Export the configurations of the selected databases to an Excel file. The configurations include the instance status, environment, DBA, and owner.
- o Export permission information

Export the permission information about the selected databases to an Excel file. The permission information includes the database information, users, permissions, and users who grant the permissions.
- o Other operations

You can click Tables in the Actions column of a database to view the details about tables in the database. You can also move the pointer over More and select the required operation that you want to perform. For example, you can query data in the database, manage permissions, view the details about the instance to which the database belongs, and find the instance on the Instance List tab.

15.11.3. Manage users

Data Management (DMS) allows you to manage users. For example, you can manage user permissions and roles.

Prerequisites

You are a DMS administrator.

Procedure

1. [Log on to the DMS console.](#)
2. In the top navigation bar, move the pointer over the **More** icon and choose **System > User**.
3. Select the user that you want to manage. Then, you can perform the following operations based on your business requirements:
 - o Add a user
Click **New** and set parameters to add a user. For more information, see [Add a user](#).
 - o Edit a user
Click **Change** in the **Actions** column. In the dialog box that appears, modify the user configurations and click **Confirm Change**.
 - o Disable or enable a user
Move the pointer over **More** in the **Actions** column and select **Disable** or **Enable**.
 - o Remove a user
Move the pointer over **More** in the **Actions** column and select **Delete**. In the message that appears, click **OK**.
 - o Grant permissions
Move the pointer over **Authorize** in the **Actions** column and select the object on which you want to grant permissions to the user. For example, you can select **Authorize instance**, **Authorize database**, **Authorize table**, **Authorize Line**, or **Authorize sensitive column**. In the dialog box that appears, enter a keyword to filter objects, set the Permission and Expire Date parameters, and then click **OK**.
 - o Release permissions
Move the pointer over **More** in the **Actions** column and select **Permission Details**. You can set conditions to filter the permissions that are granted to the user, select the permission type that you want to manage, and then click **Release Permission** to release the permissions.

15.11.4. Enable metadata access control

Data Management (DMS) provides the metadata access control feature. You can use this feature to allow users to view the information about and access a database or database instance on which they have permissions. Before this feature is enabled, regular users can query all databases and database instances within the current tenant account. After this feature is enabled, you can allow specific users to view the information about and access the databases or database instances on which they have permissions. This further enhances the data security of your enterprise.

Background information

As a centralized data management service, DMS provides different roles that are assigned different permissions. This helps you manage data in your enterprise in a secure manner. After you enable metadata access control for a database instance or database, only users who have permissions on the instance or database can view the information about and access the instance or database. This way, users can view the information about and access only databases on which they have permissions. This further enhances data security.

Note In DMS, permissions on a database include the query, export, and change permissions. If you have one of these permissions on a database, you can view the following information about the database:

- Information about the database. You can search for the database in the search box in the upper part of the left-side navigation pane or in the top navigation bar of the DMS console. Alternatively, you can search for the database in the search box of the Select the databases, tables, or columns on which you want to apply for permissions field on the Permission Application Ticket page. Whether you can query the data in the database depends on whether you have the query permissions on the database.
- Information about the instance to which the database belongs. Whether you can view the information about other databases in this instance depends on whether you have permissions on the other databases.

You can enable metadata control access for the following objects:

- A user: The user can view the information about and access only databases on which the user has permissions.
- A database: Only users who have permissions on the database can view the information about and access the database.
- A database instance: Only users who have permissions on the database instance can view the information about and access the database instance. If a user has permissions on a database in this database instance, the user can view the information about and access this database.

Enable metadata access control for a user

1. [Log on to the DMS console.](#)
2. In the top navigation bar, move the pointer over the **More** icon and choose **System > User**.

Note You are a DMS administrator.

3. Find the user for whom you want to enable metadata access control, move the pointer over **More** in the **Actions** column, and then select **Access control**.
4. In the User access control dialog box, turn on **Metadata access control** and click **OK**.

Enable metadata access control for a database instance

1. [Log on to the DMS console.](#)
2. In the top navigation bar, move the pointer over the **More** icon and choose **System > Instance**.

Note You are a database administrator (DBA) or a DMS administrator.

3. On the **Instance List** tab, find the instance for which you want to enable metadata access control. Then, select the instance and click **Access control** in the upper part of this tab.

Note You can also enable metadata access control for multiple instances at a time. Select multiple instances and click **Access control** in the upper part of this tab.

4. In the dialog box that appears, turn on **Metadata access control** and click **OK**.

Enable metadata access control for a database

1. [Log on to the DMS console.](#)
2. In the top navigation bar, move the pointer over the **More** icon and choose **System > Instance**.

Note You are a DBA or a DMS administrator.

3. On the **Database List** tab, find the database for which you want to enable metadata access control. Move

the pointer over **More** in the **Actions** column and select **Access control**.

 **Note** You can also enable metadata access control for multiple databases at a time. Select multiple databases and click **Access control** in the upper part of this tab.

4. In the dialog box that appears, turn on **Metadata access control** and click **OK**.

15.11.5. Manage tasks

The task management feature allows you to manage various tasks that are created by using tickets. You can also use this feature to directly create or manage tasks.

Prerequisites

You are a database administrator (DBA) or an administrator.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **More > System > Task**.
3. On the **Task** tab, view and manage tasks that are created by using tickets.
4. Find a task and perform one of the following operations based on your business requirements: For example, you can **pause**, **retry**, or **delete** a task.
 - o **Pause a task**
Click **Pause** to pause a task.
 - o **Retry a task**
If a task is in the **Failure** state, click **Retry** to run the task again.
 - o **Delete a task**
Click **Delete** to delete a task. After a task is deleted, the status of the task changes to **Delete** and the task cannot be run.
 - o **Create a task**
Click **Add SQL task**. In the Add SQL task dialog box, enter the task description, the database that you want to manage, and the SQL statements that you want to execute. Then, click **Submit Task**.

15.11.6. Configuration management

DMS allows you to manage system configurations. To implement flexible management, you must log on to the DMS console as an administrator to modify the required system configurations.

Prerequisites

An administrator account is required to perform the operation.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **More > System > Configuration**.
3. Find the required parameter and click **Change** in the **Actions** column of the parameter.

 **Note** You can click **Change History** to view the change history of the parameter.

4. In the Change Parameter Configuration dialog box, enter the required value.

5. Click **Confirm Change**.

Types of data changes

| key | value | Description |
|---------------------------------------|---------------------------------------|--|
| config_correct | Modify Config | Modifies configurations. |
| project_init_data | Init Project Data | Initializes the data for a project. |
| program_bug | Program Bug | Fixes a bug. |
| require_deal_without_backend_function | Requirements Without Backend Function | Manages the data of an application that does not support backend management. |
| history_data_clear | History Data Clean | Clears historical data. |
| test | Test | Runs a test. |
| mis_operation | Mis Operation | Restores data after a misoperation. |
| others | Others | Changes data for other reasons. |

15.11.7. Database grouping

This topic describes how to create a database group in Data Management (DMS). You can use this feature to apply a data change or a schema change to all of the databases in a database group with ease.

Prerequisites

The databases that you want to add to a database group meet the following conditions:

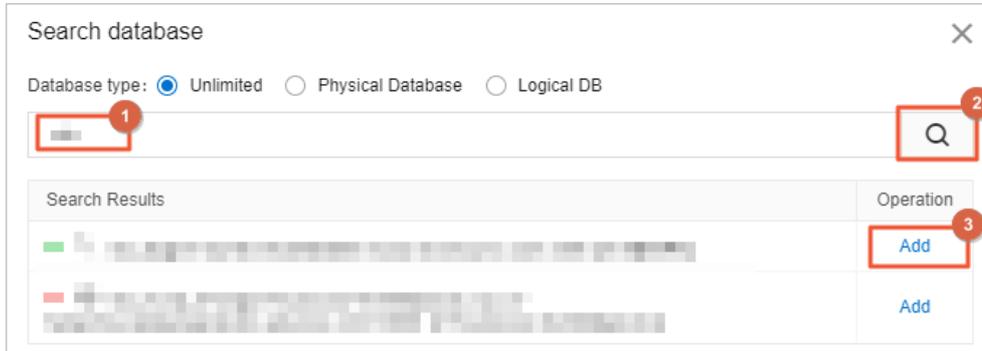
- All of the instances to which the databases belong are managed in Security Collaboration mode.
- All of the databases are physical databases or logical databases.
- All of the databases are deployed in the same environment, such as the development environment.
- The engines of the databases are of the same type. For example, all of the databases are MySQL databases.

Create a database group

1. [Log on to the DMS console](#).
2. In the top navigation bar, move the pointer over the **More** icon and choose **System > Database grouping**.
3. Click **New Group**.
4. In the **NewGrouping** dialog box, perform the following steps:
 - i. Enter a group name in the **Group name** field.
 - ii. Set the **Grouping type** parameter to **General grouping**.

 **Note** You cannot set this parameter to **Remote live**. This feature will be available soon.

- iii. Click **Add database**. In the **Search database** dialog box, enter a prefix in the search box to search for databases. Select one or more databases to be grouped from the matched results and click **Add** in the Operation column.

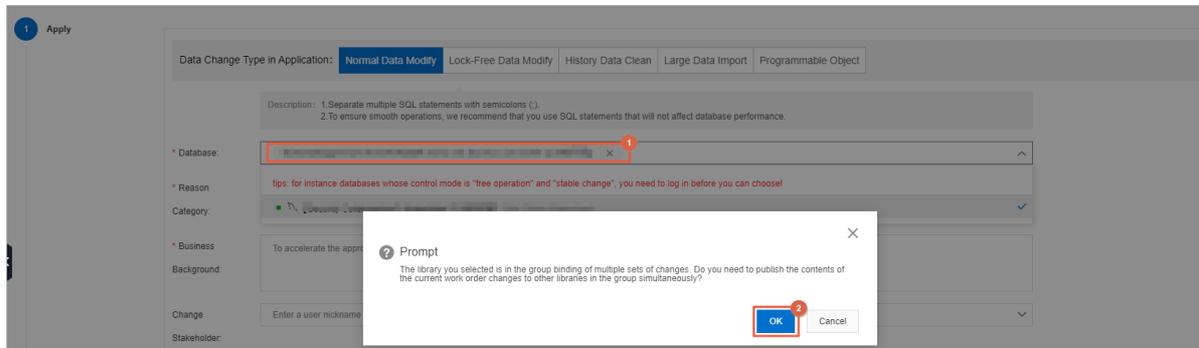


- iv. After you add all the databases to be grouped, click the **X** icon in the upper-right corner to close the Search database dialog box.
- 5. After you complete the configurations, click **Save**.

Scenarios

- Data change

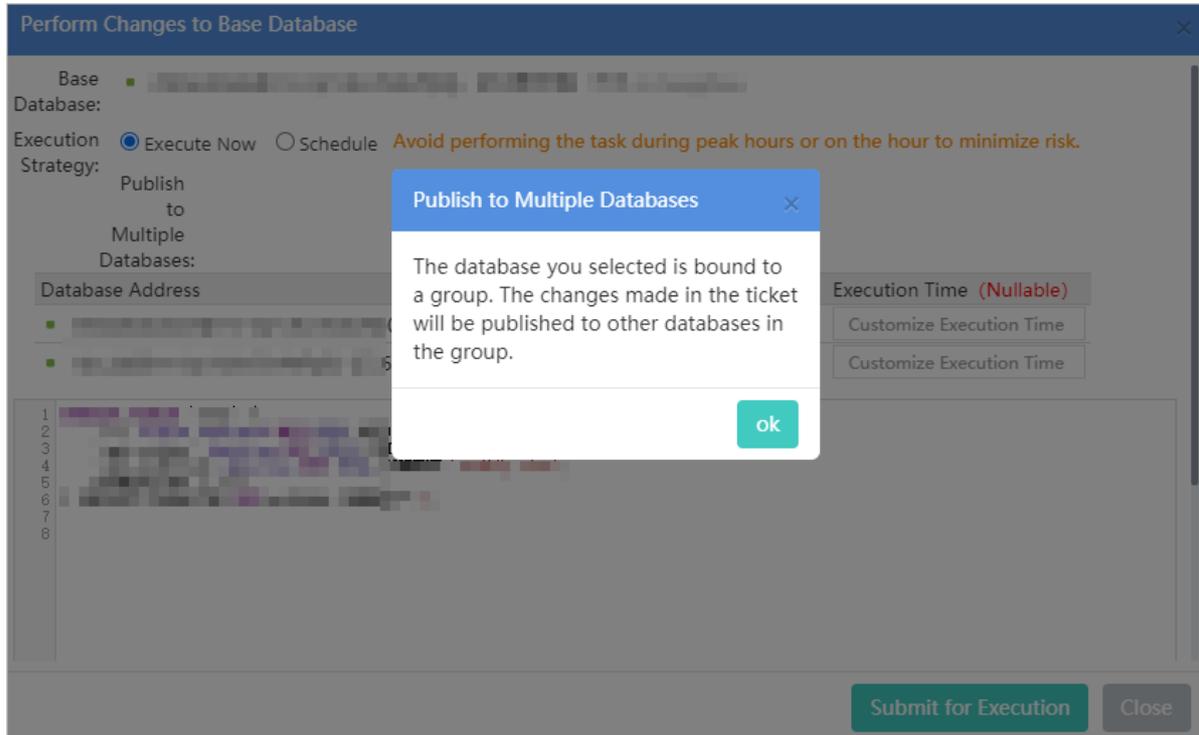
For example, you want to create a ticket to perform a data change on a database, and the database belongs to a database group. After you select the database, DMS displays a message to remind you that the selected database belongs to a database group. If you click **OK**, DMS adds all the other databases in the group as the databases on which the data change will be performed. This saves your effort in selecting databases one by one. If you click **Cancel**, the other databases in the group will not be selected. The following figure shows the message.



This feature applies to the data change and data import tickets that are supported by DMS. For more information about how to create a ticket, see [Change data](#) and [Import data](#).

- Schema design

For example, you want to create a schema design ticket and select a database that belongs to a database group as a base database. After you click **Perform Changes to Base Database**, DMS displays a message. This message is used to remind you that the base database belongs to a database group and the current operation will apply to all the other databases in the group. The following figure shows the message.



For more information about how to use the schema design feature, see [Design a schema](#).

15.11.8. Security management

15.11.8.1. Manage security rules

Security rules are implemented by using a collection of domain-specific languages (DSLs) to control user access to databases based on several factors. These factors include the type of databases, the syntax of database operations, and the number of affected rows. You can use security rules to standardize database operations, development processes, and approval processes as required. This topic describes how to manage security rules.

Prerequisites

You are a database administrator (DBA) or an administrator.

Procedure

1. [Log on to the DMS console](#).
2. In the top navigation bar, choose **More > System > Security Rules**.
3. Perform one of the following operations based on your business requirements:
 - o Create a rule set
Click **Create Rule Set**. In the Create Rule Set dialog box, set the Engine Type, Rule Set Name, and Remarks parameters, and click **Submit**.
 - o Edit a rule set
 - a. Find the required rule set and click **Edit** in the **Actions** column of the rule set.

- b. In the left-side pane, click the required rule subset, for example, **SQLConsole**. In the right-side pane, select a checkpoint.
- c. Find the required rule and click **Edit** next to the rule. For more information about the related syntax, see [DSL syntax for security rules](#).

 **Note** You can disable or delete a rule.

- o Create a similar rule set
 - a. Find the required rule set and click **Create As** in the **Actions** column of the rule set.
 - b. In the dialog box that appears, enter a name and a description for the new rule set.
 - c. Click **Submit**. The system copies the configurations of the original rule set to the new rule set.

- o Delete a rule set

Find the required rule set and click **Delete** in the **Actions** column of the rule set. In the message that appears, click **OK**.

 **Note**

- A deleted rule set cannot be recovered. Proceed with caution.
- You can delete only custom rule sets. You cannot delete built-in rule sets.

- o Set a rule set as the default rule set

Find the required rule set and click **Set as Default** next to the rule set. In the message that appears, click **OK**. The rule set is used as the default rule set for the related database engine.

15.11.8.2. DSL syntax for security rules

DMS provides a domain-specific language (DSL) to describe security rules. You can use the DSL syntax to define security rules. This allows you to define database development standards based on your business requirements.

Overview

The DSL syntax can include one or more conditions and related actions that are specified by an IF-ELSE statement.

 **Note** The if clause is required. Zero or more elseif clauses can be specified. Zero or one else clause can be specified.

Example 1: If Condition 1 is met, DMS performs Action 1.

```
if
  Condition 1
then
  Action 1
end
```

Example 2: If Condition 1 is met, DMS performs Action 1. If Condition 2 is met, DMS performs Action 2. If Condition 1 and Condition 2 are not met, DMS performs Action 3.

 **Note** If the `else Action 3` clause is removed and Condition 1 and Condition 2 are not met, DMS performs no action.

```

if
  Condition 1
then
  Action 1
elseif
  Condition 2
then
  Action 2
[else Action 3]
end
    
```

DSL syntax

- Conditional clauses

DMS uses conditional clauses to evaluate whether to perform actions. The result of a conditional clause is true or false. A conditional clause consists of one or more connectors, operators, and factors. Connectors include AND and OR. Factors are predefined system variables. The following examples are valid conditional clauses:

```

1. true           // This is the simplest conditional clause. The result is true.
2. 1 > 0
3. 1 > 0 and 2 > 1
4. 1 <= 0 or 1 == 1
    
```

- Connectors

Connectors include AND and OR. The AND connector has higher priority than the OR connector. The two connectors have lower priority than operators. For example, a conditional clause is `1 <= 0 or 1 == 1`. DMS evaluates the result of the `1 <= 0` expression and the result of the `1 == 1` expression. Then, DMS evaluates the result of the OR expression based on the preceding results.

- Operators

Operators are used to connect factors and constants to perform logical operations. The following table describes the operators that are supported by DMS.

| Operator | Description | Examples |
|----------|--|----------------------------|
| == | Evaluates whether a value is equal to another value. | 1 == 1 |
| != | Evaluates whether a value is not equal to another value. | 1 != 2 |
| > | Evaluates whether a value is greater than another value. | 1 > 2 |
| >= | Evaluates whether a value is greater than or equal to another value. | 1 >= 2 |
| < | Evaluates whether a value is less than another value. | 1 < 2 |
| <= | Evaluates whether a value is less than or equal to another value. | 1 <= 2 |
| in | Evaluates whether a value belongs to an array of values. | 'a' in ['a', 'b', 'c'] |
| not in | Evaluates whether a value does not belong to an array of values. | 'a' not in ['a', 'b', 'c'] |

| Operator | Description | Examples |
|-------------|---|------------------------------|
| matches | Evaluates whether a string matches a regular expression. | 'idxaa' matches 'idx\w+' |
| not matches | Evaluates whether a string does not match a regular expression. | 'idxaa' not matches 'idx\w+' |
| isBlank | Evaluates whether a value is empty. | " isBlank |
| isNotBlank | Evaluates whether a value is not empty. | " isNotBlank |

Note

- If you need to use a backslash (\) in a regular expression, you must add another backslash (\) as an escape character before the backslash that you want to use. For example, if you want to write the `idx_\w+` expression, you must enter `idx_\\w+`.
- If a conditional clause includes nested expressions, we recommend that you enclose the required expressions in parentheses (). For example, a conditional clause is `1 <= 2 == true`. To specify the priority, you can change the clause to `(1 <= 2) == true`. DMS first evaluates the result of the `1 <= 2` expression in the parentheses.

• Factors

A factor is a predefined variable in DMS. You can use factors to obtain the context to be validated by security rules. The context includes command categories and the number of affected rows. A factor name is prefixed by `@fac.`. Each tab of the Security Rules tab includes different factors for different checkpoints. The following table describes the factors that are supported by DMS.

| Factor | Description |
|---|--|
| <code>@fac.env_type</code> | The type of the environment. The value is the display name of the environment type, such as DEV or PRODUCT. |
| <code>@fac.sql_type</code> | The type of the SQL statement, for example, UPDATE or INSERT. For more information, see the SQL subcategories that are described in the "SQLConsole for relational databases" topic. |
| <code>@fac.detail_type</code> | The type of the data change. Valid values: <ul style="list-style-type: none"> ○ COMMON: a normal data modify ticket ○ CHUNK_DML: a lock-free data modify ticket ○ PROCEDURE: a programmable object ticket ○ CRON_CLEAR_DATA: a history data clean ticket ○ BIG_FILE: a large data import ticket |
| <code>@fac.is_logic</code> | A Boolean value that indicates whether the affected database is a logical database. |
| <code>@fac.extra_info</code> | Other information about the ticket. This factor is not in use. |
| <code>@fac.is_ignore_affect_rows</code> | A Boolean value that indicates whether to skip the validation. |
| <code>@fac.insert_rows</code> | The number of data rows to be inserted. |

| Factor | Description |
|-----------------------------|--|
| @fac.update_delete_rows | The number of data rows to be updated. |
| @fac.max_alter_table_size | The size of the largest tablespace in which the table to be modified is stored. |
| @fac.is_has_security_column | A Boolean value that indicates whether sensitive fields are specified in the SQL statement to be executed. |
| @fac.security_column_list | The sensitive fields that are specified in the SQL statement to be executed. |
| @fac.risk_level | The risk level that is identified. |
| @fac.risk_reason | The reason based on which the operation is identified as this risk level. |

 **Note** You can use factors in conditional clauses. For example, you can write `@fac.sql_type == 'DML'` to evaluate whether an SQL statement is a DML statement.

- Action clauses

An action indicates an operation that is performed when the if clause evaluates to true. For example, DMS can disable the submission of a ticket, select an approval process, approve a ticket, or reject a ticket. An action indicates the usage of a security rule. An action name is prefixed by `@act.`. Each tab of the Security Rules tab includes different actions for different checkpoints. The following table describes the actions that are supported by DMS.

| Action | Description |
|--|---|
| @act.allow_submit | Requires the submission of SQL statements to be executed in a ticket. |
| @act.allow_execute_direct | Allows the execution of SQL statements in the SQLConsole. |
| @act.forbid_execute | Disables the execution of SQL statements. |
| @act.mark_risk | Marks the risk level of a data change. Example: <code>@act.mark_risk 'medium-level risk: online environment'</code> . |
| @act.do_not_approve | Specifies the ID of an approval template. |
| @act.choose_approve_template | |
| @act.choose_approve_template_with_reason | |

- Predefined functions

DMS provides predefined functions that can be used in conditional clauses and action clauses. A function name is prefixed by `@fun.`.

| Function | Description | Format |
|----------|-------------|--------|
|----------|-------------|--------|

| Function | Description | Format |
|-----------------------|--|--|
| @fun.concat | Connects strings to form a single string. Output: a string. Input: multiple strings. | @fun.concat('d', 'm', 's') // The output is the string 'dms'. @fun.concat('[Development standards] The [', @fac.column_name, '] You must enter remarks.') // The output is a prompt that reminds the user who submits the ticket to enter a value in the field. |
| @fun.char_length | Calculates the length of a string. Output: an integer. Input: a string. | @fun.char_length('dms') // The output is 3. @fun.char_length(@fac.table_name) // The output is the length of the table name. |
| @fun.is_char_lower | Evaluates whether all the letters in a string are lowercase letters. Output: true or false. Input: a string. | @fun.is_char_lower('dms') // The output is true. @fun.is_char_lower(@fac.table_name) // If the output is true, it indicates that all the letters in the table name are lowercase. |
| @fun.is_char_upper | Evaluates whether all the letters in a string are uppercase letters. Output: true or false. Input: a string. | @fun.is_char_upper('dms') // The output is false. @fun.is_char_upper(@fac.table_name) // If all the letters in the table name are uppercase letters, the output is true. |
| @fun.array_size | Counts the number of values in an array. Output: an integer. Input: an array of values. | @fun.array_size([1, 2, 3]) // The output is 3. @fun.array_size(@fac.table_index_array) // The output is the number of indexes of the table. |
| @fun.add | Adds multiple numeric values. Output: a numeric value. Input: multiple numeric values. | @fun.add(1, 2, 3) // 6 |
| @fun.sub | Deducts a numeric value from another numeric value. Output: a numeric value. Input: two numeric values. | @fun.sub(6, 1) // 5 |
| @fun.between | Evaluates whether a value belongs to a specific closed range. The supported data types are NUMERIC, DATE, and TIME. Output: true or false. Input: three values. The first value is the value to be evaluated. The second value indicates the lower limit. The third value indicates the upper limit. | @fun.between(1, 1, 3) // The output is true because the value 1 belongs to [1, 3]. @fun.between(2, 1, 3) // The output is true because the value 2 belongs to [1, 3]. @fun.between(7, 1, 3) // The output is false because the value 7 does not belong to [1, 3]. @fun.between(@fac.export_rows, 2001, 100000) // If the number of exported rows belongs to [2001, 100000], the output is true. @fun.between(@fun.current_datetime(), '2019-10-31 00:00:00', '2019-11-04 00:00:00') // If the current date and time belong to [2019-10-31 00:00:00, 2019-11-04 00:00:00], the output is true. @fun.between(@fun.current_date(), '2019-10-31', '2019-11-04') // If the current date belongs to [2019-10-31, 2019-11-04], the output is true. @fun.between(@fun.current_time(), '13:30:00', '23:59:59') // If the current time belongs to [13:30:00, 23:59:59], the output is true. |
| @fun.current_datetime | Returns the current date and time, in the format of yyyy-MM-dd HH:mm:ss. Output: a string. Input: none. | @fun.current_datetime() // For example, the output is 2019-10-31 00:00:00. |
| @fun.current_date | Returns the current date, in the format of yyyy-MM-dd. Output: a string. Input: none. | @fun.current_date() // For example, the output is 2020-01-13. |

| Function | Description | Format |
|-------------------|---|---|
| @fun.current_time | Returns the current time, in the format of HH:mm:ss. Output: a string. Input: none. | @fun.current_time() // For example, the output is 19:43:20. |

DSL configuration examples

Limit the number of SQL statements in a ticket: If the number of SQL statements in a ticket exceeds 1,000, DMS rejects the ticket and returns the related message.

```
if
  @fac.sql_count > 1000
then
  @act.reject_execute 'The number of SQL statements in a ticket cannot exceed 1,000.'
else
  @act.allow_execute
end
```

Allows the submission of only data manipulation language (DML) statements: If the SQL statements in a ticket are DML statements such as the UPDATE, DELETE, and INSERT statements, DMS allows the execution of the statements.

```
if
  @fac.sql_type in ['UPDATE','DELETE','INSERT','INSERT_SELECT']
then
  @act.allow_submit
end
```

15.11.8.3. Configure security rules for a database instance

This topic describes how to configure security rules for a database instance.

Prerequisites

- You are a database administrator (DBA) or an administrator.
- The control mode of the database instance is **security collaboration**.

Procedure

- Log on to the [DMS console](#).
- In the left-side navigation pane, right-click the required database instance.
- In the shortcut menu that appears, choose **Control Mode > Security Collaboration** and select the required security rule set.

 **Note** You can also change the security rule set of the database instance on the Instance tab. For more information, see [Instance management](#).

15.11.8.4. Customize approval processes

Data Management (DMS) allows you to configure instance-level security rules so that you can customize different approval processes for different database instances or database operations. However, instance-level security rules have some limits in the production environment. This topic describes how to customize an approval process.

Prerequisites

You are a database administrator (DBA) or an administrator.

Context

- Each database instance has only one DBA. However, multiple DBAs are included in an approval process to ensure business continuity regardless of whether one of the DBAs is unavailable.
- If multiple business units share the same database in a database instance, each business unit must approve the tickets for their respective business operations in an approval process.

To resolve the issues in the preceding scenarios, you can customize approval processes.

Precautions

- Do not assign only one approver to an approval node. We recommend that you assign at least two approvers to each approval node and at least two data owners to a database.
- You can assign a maximum of three data owners to a database. If multiple business units share the same database, you can specify these business units in an approval process by performing the following steps: Create an approval node and add the data owners of the business units as approvers. Then, add the new node instead of the system node Owner to an approval template.

Procedure

This topic describes how to customize an approval process and specify multiple DBAs in the approval process. You can perform similar steps to customize an approval process in other scenarios.

1. [Log on to the DMS console.](#)
2. In the top navigation bar, choose **More > System > Approval Processes**.
3. Create an approval node.
 - i. Click the **Approval Node** tab. Then, click **Create Approval Node**.
 - ii. The following table describes the parameters that you can specify for the approval node.

| Parameter | Description |
|------------------|---|
| Node Name | The name of the approval node. The name must be globally unique. |
| Remarks | The description of the approval node. This parameter distinguishes the approval node from other approval nodes. |
| Approver | The Apsara Stack tenant accounts of the approvers for the approval node. You can search for approvers by keyword. Prefix match is used. <div style="background-color: #e0f2f7; padding: 5px; margin-top: 10px;"> ? Note In this example, three approvers are selected. </div> |

- iii. Click **Submit**.
4. Create an approval template.
 - i. Click the **Approval Template** tab. Then, click **Create Approval Template**.

ii. The following table describes the parameters that you can specify for the approval template.

| Parameter | Description |
|----------------------|---|
| Template Name | The name of the approval template. The name must be globally unique. |
| Remarks | The description of the approval template. This parameter distinguishes the approval template from other approval templates. |
| Approval Node | <p>Click Add Node and select the required approval nodes. In this example, the system node Owner and the approval node that is created in Step 3 are selected to allow multiple DBAs to participate in the approval process.</p> <div style="border: 1px solid #add8e6; padding: 5px;"> <p>Note The approval process is implemented based on the values of the Approval Order parameter in ascending order.</p> </div> |

iii. Click **Submit**.

After the approval template is created, you can view the template ID. In this example, the template ID is 9.

The screenshot shows the 'Create Approval Template' page. At the top, there is a search bar with 'dmstest' and a 'Create Approval Template' button. Below the search bar is a note: "Note: When the template ID is -1, it is free of approval, that is, the approval process with the approval template of -1 is selected, and the approval is automatically passed." A table below lists the template details:

| Templ... ID | Template Name | Template Type | Created By | Approval Node | Remarks | Actions |
|-------------|---------------|---------------|------------|---------------|---------|---------------|
| 9 | dmstest | Custom | [Redacted] | 1 | dmstest | Edit Delete |

5. Apply a new approval process.

This example shows how to edit a rule that is applied to medium-level risk approval processes under the **Risk Approval Rules** checkpoint. You can perform similar steps to apply a rule to other scenarios.

- i. In the top navigation bar, choose **System > Security > Security Rules**.
- ii. Find the required rule set that you want to edit and click **Edit** in the **Actions** column of the rule set.
- iii. In the left-side navigation pane, click the **SQL Correct** tab.
- iv. Select **Risk Approval Rules** as the checkpoint.
- v. Find the rule that is related to the medium-level risk approval process and click **Edit**.
- vi. In the **Rule DSL** field, change the template ID.

The screenshot shows the 'Change Rule - SQL Correct' dialog box. The 'Checkpoints' dropdown is set to 'Risk Approval Rules'. The 'Template' is 'Load from Template Database'. The 'Rule Name' is '中风险审批流程'. The 'Rule DSL' field contains the following code:

```

1  if
2    @fac.risk_level=='middle'
3  then
4    @act.choose_approve_template 3
5  end
    
```

Note In this example, change 3 to 9, as shown in the preceding figure. The ID 9 is the ID of the approval template that is created in Step 4.

vii. Click **Submit**.

Result

If the data change tickets that you submit match the rule, all specified DBAs receive ticket approval notifications and can participate in the approval process.

15.11.8.5. Operation audit

Data Management (DMS) provides the operation audit feature in addition to the basic features of operation log management. You can use this feature to troubleshoot database issues with ease and audit the operations that are performed on databases. You can also use this feature to view and manage the SQL statements that are used in the SQLConsole, tickets, logon information, and operation logs.

Features

The following table describes the two modules of the operation audit feature in DMS: Operation Logs and Operation audit.

| Module | Description | Item |
|-----------------|--|--|
| Operation Logs | Displays the logs of all the operations that are performed in DMS. | Includes the logs of management and configuration operations, SQL statements that are used in the SQLConsole, tickets, and logon information. |
| Operation audit | Displays all the operations that are performed on databases in DMS.  Note This module provides a user interface (UI) for you to audit operations in a centralized manner. This also helps you troubleshoot database issues with ease. | Includes SQL statements that are used in the SQLConsole, tickets, and logon information.  Note Only a DMS administrator, a database administrator (DBA), a ticket submitter, and stakeholders involved in the ticket approval process are allowed to view the ticket details. |

Log data is permanently retained in DMS. You can access and view the log data of the instances that are managed in **Stable Change** or **Security Collaboration** mode at any time.

 **Note** You can view the log data of the instances that are managed in **Flexible Management** mode only for the last seven days. To view all log data, change the control mode of the instances.

Links and supported roles

The following table describes the roles that you can assume to use the operation audit feature. It also shows you how to go to the Operation audit tab in the DMS console.

| Auditing dimension | Limit | Link to operation audit | Supported role |
|--------------------|-------|-------------------------|----------------|
| | | | |

| Auditing dimension | Limit | Link to operation audit | Supported role |
|--------------------|--|---|--|
| Database | You can view and audit only the operations that are performed on the current database. | <ul style="list-style-type: none"> On the SQLConsole tab of the database that you want to audit, click the  icon in the upper-right corner. In the left-side navigation pane of the DMS console, click the instance in which the database you want to audit resides, right-click the database, and then select Operation audit. | <p>You can be a DMS administrator, a security administrator, a DBA, an instance owner, or a regular user.</p> <p>Note If you are a regular user, you can view and audit only the operations that you performed on the current database.</p> |
| Instance | You can view and audit only the operations that are performed on the current instance. | In the left-side navigation pane of the DMS console, right-click the instance that you want to audit and select Operation audit . | <p>You can be a DMS administrator, a security administrator, a DBA, an instance owner, or a regular user.</p> <p>Note If you are a regular user, you can view and audit only the operations that you performed on the current instance.</p> |
| Global | You can view and audit all the operations that are performed in DMS. | In the top navigation bar, move the pointer over the More icon and choose System > Operation audit . | You can be a DMS administrator, a security administrator, or a DBA. |

View and download operation records

This example shows you how to view and download all the SQL statements that are used in the SQLConsole in the last month.

1. [Log on to the DMS console](#).
2. In the top navigation bar, move the pointer over the **More** icon and choose **System > Operation audit**.
By default, a list of SQL statements appears.
3. Set the Time parameter to **Last One Month** and click **Search**.
Then, DMS returns the search results.
4. Click the  icon to download the results.

Then, DMS exports an XLSX file that contains the search results on the current page.

The screenshot shows the 'Operation audit' interface with the following details:

- Function:** SQL window list (selected), Tickets, Logon list
- Classification:** All (dropdown), Search (button), Enter the user (input), Enter the instance name and database name (input), Search (button), Reset (button)
- Time:** Last One Day, Last Three Days, Last One Week, Last One Month (selected), Custom (dropdown), Download (button), Refresh (button)
- Table Columns:** Operated At, SystemUse..., Operation instance, Database/scha..., SQL Type, SQL, Status, Rows, Elapsed time(ms), Remarks
- Table Data:**

| Operated At | SystemUse... | Operation instance | Database/scha... | SQL Type | SQL | Status | Rows | Elapsed time(ms) | Remarks |
|---------------------|--------------|--------------------|------------------|----------|--|---------|------|------------------|---------|
| 2021-01-16 14:33:36 | | | | SELECT | SELECT 'student_id' as 'student id','course id' as | Success | 0 | 6 | |
| 2021-01-16 14:32:54 | | | | SELECT | | Success | 0 | 14 | |
| 2021-01-16 14:14:33 | | | | SELECT | | Success | 3 | 5 | |
| 2021-01-16 14:14:04 | | | | SELECT | | Success | 3 | 4 | |

Note To preview and export more results, you can set the **Items per page** parameter to 100.

15.11.8.6. Configure IP whitelists

DMS allows you to configure IP whitelists to control the service scope of DMS. You can allow user access to DMS only from specific trusted network environments.

Prerequisites

You are an administrator.

Procedure

1. Log on to the DMS console.
2. In the top navigation, choose **More > System > Access IP Whitelists**.
3. Perform the following operations based on your business requirements:
 - o Enable or disable the whitelist control feature
Click **Click to Open** or **Click to Close** to enable or disable the whitelist control feature.
 - o Create a whitelist
 - a. Click **Create Whitelist**.
 - b. In the dialog box that appears, enter the IP addresses and description.

Note

- Separate IP addresses with semicolons (;). Make sure that each IP address in a whitelist is unique.
- You can specify IP addresses such as 10.23.12.24 or CIDR blocks such as 10.23.12.24/24, where /24 indicates the length of the IP address prefix in the CIDR block. The IP address prefix can be 1 to 32 bits in size.
- The 0.0.0.0/0 value indicates that all IP addresses are allowed.

- c. Click **Submit**.
- o Edit a whitelist
 - a. Find the required IP address whitelist and click **Edit** in the **Actions** column of the IP address.

- b. In the dialog box that appears, modify the IP address information.
- c. Click **Submit**.
- o Delete a whitelist
 - a. Find the required IP address whitelist and click **Delete** in the **Actions** column of the IP address.
 - b. In the message that appears, click **OK**.

 **Note** You cannot delete all IP address whitelists. At least one IP address whitelist must be retained.

15.11.8.7. Row-level control

In some cases, different users may access different rows in the same table, which can be achieved by using views. Data Management (DMS) provides an alternative solution that is called the row-level control feature to control access at the row level.

Prerequisites

You are a security administrator, a database administrator (DBA), or an administrator.

Context

Row-level control is used to provide horizontal data protection for tables. All the rows in a table are distinguished by one or more specified values. These values are called control values. To access a row that corresponds to a control value in the DMS console, you must have permissions on the row.

 **Note** A control value may correspond to multiple rows. If a user has permissions on a control value that corresponds to multiple rows, the user has permissions on all the rows that correspond to the control value.



Limits

- The sensitive data management feature applies only to relational databases, such as MySQL. However, this feature is unavailable for NoSQL databases.
- You can use the row-level control feature only on database instances managed in security collaboration mode.
- This feature applies only to physical databases. However, this feature is unavailable for logical databases.
- When you execute SQL statements to query, modify, or delete the data of a row-level control table, the following limits are set on filter conditions.
 - i. The control field must be specified in SQL statements to filter data.
 - ii. The system controls access to all the rows of a row-level control table. Users who do not have permissions on all rows can use only the `=` and `IN` operators to specify a control field. The control value that is specified in an SQL statement must be one of the control values for the table.
 - iii. Users who do not have permissions on all rows cannot use some operators, such as OR, XOR, and logical NOT.

Terms

| Term | Description |
|-------------------------|---|
| row permission | You can apply for permissions on a control value to access rows that correspond to the control value. Permissions on the rows of a table are defined as row permissions and are incorporated into the existing permissions of DMS. Permissions that can be controlled in security collaboration mode include permissions on databases, tables, columns (fields), and rows. |
| single control value | When a user applies for permissions on the rows of a row-level control table, the user can select Single to apply for permissions on a single control value. <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e6f2ff;"> <p> Note A control value may correspond to multiple rows. If a user has permissions on a control value that corresponds to multiple rows, the user has permissions on all the rows that correspond to this control value.</p> </div> |
| all control values | When a user applies for permissions on the rows of a row-level control table, the user can select ALL to apply for permissions on all control values. After the application is approved, the user has permissions on all the rows of the table. In this case, the user can access the entire row-level control table without limits. Even if the control values are changed or more control values are added, the user still has permissions on all the rows of the table. |
| row-level control table | A table that requires row-level control is called a row-level control table. |
| control field | A control field is a field to which the control values of a row-level control table are added. |
| control group | A control group is a group of row-level control tables that have the same control values. For example, if Table A and Table B have the same control values, you can add the two tables to a control group. This way, you can manage the two tables at the same time by using one set of control values. |

Procedure

1. [Log on to the DMS console.](#)
2. In the top navigation bar, choose **More > System > Sensitive Data**.
3. Click the **Row Level Security** tab.
4. Add a control group.
 - i. Click **Add control group**.

ii. In the Add control group dialog box, set the required parameters.

| Parameter | Description |
|--------------------------|---|
| Control Group | Enter a name for the control group. |
| Row Configuration | Click Add to add a row configuration in which you can specify a database, table, and field. <div style="background-color: #e0f2f1; padding: 5px;"> ? Note You can repeat this step to add multiple row configurations. </div> |
| DB Table Column | Search for databases by keyword and select a database. Then, select a table and a field from the drop-down lists. <div style="background-color: #e0f2f1; padding: 5px;"> ? Note The selected field is the control field. </div> |

iii. Click **Add**.

5. Add control values.

i. Find the new control group and click **Details** in the **Actions** column of the control group.

ii. Click **Add Row Value**.

iii. In the Import Row Value dialog box, specify whether to append row values and enter the required row values.

? **Note** Separate multiple row values with commas (,).

iv. Click **Import**.

What to do next

After you configure row-level control settings for a table, a user may still have no permission on a control value that corresponds to one or more rows in the table. In this case, an error appears when the user queries row data. The error indicates that the user does not have permissions to access the row. The user can apply for permissions on the control value to access the rows. For more information, see [Apply for permissions](#).

15.11.8.8. Manage sensitive data

Data Management (DMS) allows you to manage all classified sensitive and confidential fields in a unified manner. You can configure encryption algorithms for sensitive and confidential fields. This improves the control over the data masking feature.

Prerequisites

You are a security administrator, a database administrator (DBA), or an administrator.

Context

When you query a table that contains sensitive or confidential fields on which and you do not have permissions on the fields, the values of the fields are displayed as ********* in the query results. In this case, sensitive data is fully masked. In some scenarios, developers or test engineers may need to view a part of sensitive data for troubleshooting. To meet this requirement, you can configure masking algorithms to show some sensitive data.

Limits

- The sensitive data management feature applies only to relational databases such as MySQL. However, this feature is unavailable for NoSQL databases.
- To use this feature, the required database instance must be managed in security collaboration mode.

Procedure

1. [Log on to the DMS console.](#)
2. Specify security levels for fields in the required table.

 **Note** If security levels are specified for the fields, skip this step.

- i. In the left-side database instance list, click the  icon next to the required database instance to show the databases in the instance.
- ii. Find the required database, right-click the database, and then select **Tables**.
- iii. Click the  icon next to the table name to show the table details.
- iv. Click **Adjust**.
- v. In the Adjust Security Level dialog box, change the security levels of fields.

Adjust Security Level ✕

Table Name: customer Security Level Description

| | Field Name | Description | Original Level | New Level(Adjust Only Changed Fields) | Operation Status |
|---|------------|-------------|----------------|--|------------------|
| 1 | id | | Internal | <input checked="" type="radio"/> Internal <input type="radio"/> Sensitive <input type="radio"/> Confidential | |
| 2 | name | | Internal | <input type="radio"/> Internal <input checked="" type="radio"/> Sensitive <input type="radio"/> Confidential | promote |
| 3 | address | | Internal | <input type="radio"/> Internal <input type="radio"/> Sensitive <input checked="" type="radio"/> Confidential | promote |

Submit for Security Department Approval
Cancel

- vi. Click **Submit for Security Department Approval**.

 **Note** The application to increase the security level of a field is automatically approved. The application to decrease the security level of a field is approved based on the approval process specified by an administrator or DBA.

- vii. In the message that appears, click **OK**.
3. In the top navigation bar, choose **More > System > Sensitive Data**.

4. Find the required field and click **Add Algorithm** in the **Actions** column of the field.
5. In the dialog box that appears, configure a masking algorithm.

Add Algorithm
✕

Basic dmstestdata.customer.name

Information:

Algorithm Fixed Position ▾

Type:

Algorithm Masking String

Configuration

Item:

Algorithm Masking Position

Configuration

Item:

Algorithm

Description:

Add
Cancel

| Parameter | Description |
|-----------------------|--|
| Algorithm Type | The type of the algorithm. You can select an algorithm type based on your business requirements. |

| Parameter | Description |
|------------------------------|---|
| Algorithm Configuration Item | <p>The algorithm configuration items vary based on the specified algorithm type.</p> <ul style="list-style-type: none"> ◦ Fixed Position algorithm type <p>You must set the Masking String and Masking Position parameters. For example, you can set the Masking String parameter to ***.</p> <p>The Masking Position parameter specifies the positions of the characters to be masked in the field values. The positions are in the format of coordinates. Examples:</p> <ul style="list-style-type: none"> ▪ (1, 4): masks the first four characters. You can also enter (4) to simplify the format. ▪ (-4): masks the last four characters. <div style="background-color: #e0f2f1; padding: 5px; border: 1px solid #ccc;"> <p> Note You can specify a maximum of three positions. For example, (1, 4), (8, 10), (-4) indicates to mask the first four characters, the eighth to tenth characters, and the last four characters.</p> </div> ◦ Fixed Character algorithm type <p>You must set the Masking String and Character to Be Replaced parameters.</p> <div style="background-color: #e0f2f1; padding: 5px; border: 1px solid #ccc;"> <p> Note The Character to Be Replaced parameter specifies the characters that you want to mask in the format of a string. You can specify a maximum of three strings.</p> </div> ◦ Full Masking algorithm type <p>You need to set only the Masking String parameter.</p> |
| Algorithm Description | Enter a description that can help you identify the algorithm. |

6. Click **Add**.

15.11.8.9. Data protection

Data Management (DMS) provides the data protection feature. This topic describes how to use the data protection feature.

Context

Data protection is one of the greatest challenges in the process of data application. Data is the core of an enterprise. Most critical and sensitive information is stored as structured data. The information includes ID numbers, bank account numbers, phone numbers, customer records, medical records, transaction records, and salary records. Security events may cause enormous economic loss and damage the reputation of enterprises. These security events include data tampering, data theft, and data misuse.

To ensure data security, you need to understand how to protect your data from these security risks. The data protection feature helps the security management team achieve the following goals:

- Intelligently recognizes and classifies sensitive data. Then, you can group the data based on the security level of fields for DMS.
- Audits databases and prevents against data loss.
- Provides an efficient method to identify data application mode. Then, you can use the user and entity behavior analytics identification model and big data security expert rules of Ant Financial to identify and manage risks.
- Provides unified masking SDKs to protect sensitive information. These SDKs are used to intelligently identify sensitive data that exists in the system content that is displayed and mask the data. These operations are performed based on the definition of sensitive information and masking policies that are specified in the data

protection feature. These SDKs are also used to provide a unified method to manage masking rules within an enterprise, and increase the efficiency of security management.

Control access

Data protection is a tool that is provided by DMS to data security administrators. To enable and use the data protection feature, you must log on to the DMS console with security administrator permissions.

1. [Log on to the DMS console.](#)

 **Note** You must use an account with administrator permissions to log on to the DMS console.

2. In the top navigation bar, choose **More > System > User**.
3. Click **Change**.
4. Select **Security Administrator** and click **Confirm Change**.
5. Use the account to re-log on to the [Log on to the DMS console](#) again.
6. In the top navigation bar, choose **More > Security > Data Protection**.

 **Note** To enable a user to use the data protection feature, you must log on to the DMS console with security administrator permission to enable the feature and authorize the user to use the feature.

Data classification

The data protection feature classifies fields based on the custom automatic classification setting that is specified in the metadata of these fields. The classification results are used to update the levels of these DMS fields. This feature facilitates the access control over DMS fields.

1. [Log on to the DMS console.](#)

 **Note** You must use an account with security administrator permissions to log on to the DMS console.

2. In the top navigation bar, choose **More > Security > Data Protection**.
3. In the left-side navigation pane, choose **Rule Configuration > Data Identification Rules**.
4. In the upper-right corner, click **Create Rule**. Then, set the **Data Type**, **Data Name**, **Owner** and **Remarks** parameters.
5. Click **Next**.
6. Set the **Level** parameter. Valid values: Internal, Sensitive, and Confidential. Select **Field Scanning** in the **Data Recognition Rules** field.
7. Click **Next**. After the rule is configured, click **Save and Enable**.

 **Note**

- You can view all rules on the **Data Identification Rules** page. You can also modify, disable, or enable a rule on this page.
- After a rule is applied, the system identifies and classifies data based on metadata every hour on the hour. The security level of a field in DMS Enterprise is updated based on the classification result. This enables field-level access control over DMS operations.

Manually correct data

The Manual Data Correction page shows all identified fields. This way, you can verify these fields. If some fields are incorrectly identified, you can remove these fields or change the related types. After data is corrected, the results are immediately synchronized to the DMS Enterprise console.

Data detection

The data detection feature is used to collect statistics based on the results of data identification from multiple dimensions and show the details of fields that are identified in the field details list. These dimensions include the security level and related instance.

15.11.9. Security rules

15.11.9.1. Overview of security rule sets

Security rule sets are implemented by using a collection of domain-specific languages (DSLs) to control user access to databases based on several factors. These factors include the type of databases, the syntax of database operations, and the number of affected rows. You can use security rule sets to standardize database operations, development processes, and approval processes as required.

Engine Type: MYSQL (ID: 4)

Rule Set Name: mysql default [Edit](#) Last Changed At: 2020-05-09 12:39:26

Rule Set Description: mysql default auto create triggered by [REDACTED]

Checkpoints: Basic Configuration Item | SQL Execution Quantity Criteria | DQL SQL Criteria | DML SQL specification (obsolete) | DDL SQL specification (obsolete) | DCL SQL specification (discarded) | Other SQL Criteria | SQL Permission Criteria | SQL Execution Performance Criteria | Exception Recognition Criteria of Database and Table Column Permissions | SQL Execution Criteria in Logical Databases

Actions: [Create Rule](#)

| ID | Configuration/Rule Name | Last Changed At | Configuration Value/Rule Status | Actions |
|----|---|---------------------|---------------------------------|----------------------|
| 15 | Maximum number of returned rows per query | 2020-05-09 12:39:26 | 200 | Edit |
| | Maximum number of rows returned for a | | | |

This topic describes the features that are supported by security rule sets. You can click the link of a feature to view the information about the feature. The information includes the basic configuration items, checkpoints, factors, actions, and supported statements or commands.

- [SQLConsole for relational databases](#)
- [SQLConsole for MongoDB](#)
- [SQLConsole for Redis](#)
- [Data change](#)
- [Permission application](#)
- [Data export](#)
- [Schema design](#)
- [Database and table synchronization](#)
- [Sensitive field change](#)
- [Test data generation](#)
- [Database cloning](#)

15.11.9.2. Manage security rules under checkpoints

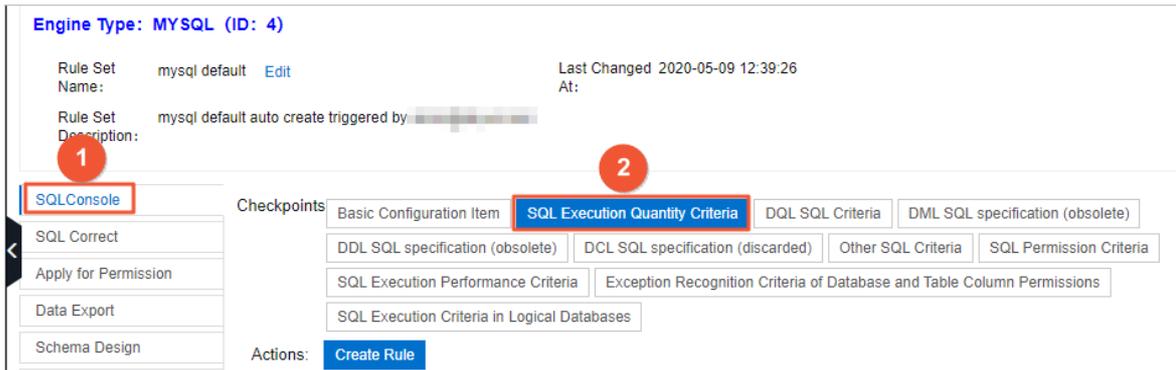
This topic describes how to configure a security rule under the SQL Execution Quantity Criteria checkpoint on the SQL Editor tab. You can configure security rules under other checkpoints by using a method similar to that described in this topic.

Procedure

1. Log on to the DMS console.
2. In the top navigation bar, choose **More > System > Security Rules**.
3. Find the required security rule set and click **Edit** in the **Actions** column of the security rule.

Note In this example, the security rule that is configured for MySQL databases is used.

4. Click a tab and then click a checkpoint based on your business requirements. In this example, click the **SQL Editor** tab and then click the **SQL Execution Quantity Criteria** checkpoint.



Note

- For more information about the tabs and checkpoints, see [Overview of security rule sets](#).
- You can click **Create Rule** to create a security rule. For more information about the syntax, see [DSL syntax for security rules](#).

5. Find the required security rule and click **Edit** in the **Actions** column of the security rule.

Note You can click **Disable** to disable a security rule or click **Delete** to delete a security rule.

6. In the Change Rule - SQL Editor dialog box, modify the DSL statements of the security rule based on your business requirements. For more information about the syntax, see [DSL syntax for security rules](#). In this example, change the maximum number of SQL statements that can be executed at a time from 1,000 to 500.

Note

- A large number of security rule templates are provided for each checkpoint. You can click **Load from Template Database** to use a template.
- For more information about factors and actions, see [Overview of security rule sets](#).

7. Click **Submit**.

15.11.9.3. SQLConsole for relational databases

DMS allows you to manage relational and non-relational databases on the SQLConsole tab. The definition and classification of security rules are different for relational and non-relational databases. This topic describes the security rules for relational databases on the SQLConsole tab, such as MySQL databases.

Default security rules

- Constraints on SQL statement categories: No constraints are imposed on data query language (DQL) statements. By default, DML statements, DDL statements, data control language (DCL) statements, and SQL statements that

cannot be identified by DMS are all blocked. To execute DML, DDL, or DCL statements on the SQLConsole tab, you must configure and enable corresponding security rules.

- Constraints on permissions on databases, tables, and fields: By default, users can perform operations on databases, tables, and fields without permission validation. To enable permission validation, you must configure and enable security rules under the **SQL Permission Criteria** checkpoint. For more information, see [Supported checkpoints](#).

Basic configuration items

| Configuration item | Description |
|---|--|
| Maximum number of returned rows per query | The maximum number of rows that can be returned for a query. |
| Maximum number of rows returned for a single query with sensitive column conditions | The maximum number of rows that can be returned for a query that contains query conditions for sensitive fields. |
| Limit the maximum allowed SQL full table scan (MB) | The maximum size of data that can be scanned. Before an SQL statement is executed, DMS checks the execution plan. If the size of the data to be scanned exceeds the specified threshold, the SQL statement fails to be executed. Note This item can be configured only for MySQL and Oracle databases. |
| Turn off the execution of change SQL validation affects the number of rows and prompts | Specifies whether to check the number of rows to be affected and display a prompt before DMS executes an SQL statement to change data. By default, this item is disabled. |
| How many rows does result set page support | The maximum number of rows that can be returned in the query result set on the SQLConsole tab. |
| Does the result set support paging | Specifies whether the query result set can be displayed on multiple pages on the SQLConsole tab. |
| Does the result set support editing | Specifies whether the query result set can be edited on the SQLConsole tab. |

Supported checkpoints

Note Various security rule templates are built in each checkpoint. You can configure security rules based on these templates or create custom security rules as needed.

| Checkpoint | Description |
|--|---|
| SQL Execution Quantity Criteria | Allows you to limit the number of SQL statements that can be submitted at a time. |
| DQL SQL Criteria | Allows you to set constraints on DQL statements. |

| Checkpoint | Description |
|---|--|
| Other SQL Criteria | <p>Allows you to set constraints on multiple categories of SQL statements. Different enterprises may define different high-risk SQL statements, which may include specific subcategories of DML, DCL, and DDL statements.</p> <p> Note You can also set constraints on SQL statements that cannot be identified by DMS.</p> |
| SQL Permission Criteria | <p>Allows you to set constraints on the execution of SQL statements from the aspect of permissions. For example, DMS checks whether a user has the required permissions on the corresponding databases, tables, and fields.</p> |
| SQL Execution Performance Criteria | <p>Allows you to set constraints on the execution of SQL statements from the aspect of performance. For example, you can specify that a DML statement is not executed if the number of rows to be affected by the statement exceeds the specified threshold, or that a DDL statement is not executed if the size of the table involved exceeds the specified threshold.</p> |
| Exception Recognition Criteria of Database and Table Column Permissions | <p>After a user submits SQL statements on the SQLConsole tab, DMS parses the SQL statements and checks whether the user has the required permissions on the corresponding databases, tables, and fields. You can configure security rules under this checkpoint to ensure that if exceptions occur when DMS parses complex SQL statements, these statements can be executed.</p> <p> Note If you configure and enable security rules under the Exception Recognition Criteria of Database and Table Column Permissions checkpoint, security rules under the SQL Permission Criteria, DQL SQL Criteria, Other SQL Criteria, and SQL Execution Performance Criteria checkpoints are automatically disabled.</p> |
| SQL Execution Criteria in Logical Databases | <p>This checkpoint is reserved for logical databases and not suitable for physical databases.</p> |

Supported factors

A factor is a built-in variable in DMS. You can use factors to obtain the context to be validated by security rules, such as SQL statement categories and the number of rows to be affected. A factor name starts with the prefix `@fac.`. The following table describes the supported factors for relational databases on the SQLConsole tab.

| Factor | Description |
|-----------------------|---|
| @fac.sql_count | The number of SQL statements that are submitted at a time. |
| @fac.select_sql_count | The number of DQL statements among the SQL statements that are submitted at a time. |
| @fac.dml_sql_count | The number of DML statements among the SQL statements that are submitted at a time. |
| @fac.sql_type | The category and subcategory of the SQL statement. For more information, see Supported SQL statements . |
| @fac.sql_sub_type | |
| @fac.env_type | The type of the environment. The value is the display name of the environment type, such as DEV or PRODUCT. |

| Factor | Description |
|----------------------------------|---|
| @fac.fulltable_delete | A Boolean value that indicates whether the current SQL statement deletes a full table. Valid values: <i>true</i> and <i>false</i> . |
| @fac.fulltable_update | A Boolean value that indicates whether the current SQL statement updates a full table. Valid values: <i>true</i> and <i>false</i> . |
| @fac.current_sql | The current SQL statement. |
| @fac.user_is_admin | A Boolean value that indicates whether the current user is a DMS administrator. Valid values: <i>true</i> and <i>false</i> . |
| @fac.user_is_dba | A Boolean value that indicates whether the current user is a DBA. Valid values: <i>true</i> and <i>false</i> . |
| @fac.user_is_inst_dba | A Boolean value that indicates whether the current user is the DBA of the current instance. Valid values: <i>true</i> and <i>false</i> . |
| @fac.user_is_sec_admin | A Boolean value that indicates whether the current user is a security administrator. Valid values: <i>true</i> and <i>false</i> . |
| @fac.sql_affected_rows | The number of rows to be affected by the current SQL statement.  Warning This factor triggers COUNT operations, which may affect the database performance. Use this factor with caution. |
| @fac.sql_relate_table_store_size | The estimated total size of the table to be accessed by the current SQL statement. Unit: MB.  Note This value is estimated based on the metadata that is obtained by DMS. It is not an actual value. |

Supported actions

An action in a security rule is an operation that DMS performs when the if condition in the rule is met. For example, DMS can forbid the submission of a ticket, approve a ticket, or reject a ticket. An action in a security rule denotes the purpose of the security rule. An action name starts with the prefix `@act.`. The following table describes the supported actions for relational databases on the SQLConsole tab.

| Action | Description |
|------------------------------|---|
| @act.reject_execute | Rejects the request to execute the current SQL statement. |
| @act.allow_execute | Allows the current SQL statement to be executed. |
| @act.reject_sql_type_execute | Rejects the request to execute a specific subcategory of SQL statements. You must specify an SQL statement subcategory after the action name. Example: <code>@act.reject_sql_type_execute 'UPDATE'</code> . |
| @act.allow_sql_type_execute | Allows a specific subcategory of SQL statements to be executed. You must specify an SQL statement subcategory after the action name. Example: <code>@act.allow_sql_type_execute 'UPDATE'</code> . |

| Action | Description |
|--|--|
| @act.check_dml_sec_column_permission | Checks whether a user has the required permissions on sensitive fields. If the user does not have the permissions, the DML statement for data change is not executed. |
| @act.uncheck_dml_sec_column_permission | Does not check whether a user has the required permissions on sensitive fields. |
| @act.check_sql_access_permission | Checks whether a user has the required permissions, such as query and change permissions, on the databases, tables, and fields that are involved in the SQL statements to be executed. |
| @act.uncheck_sql_access_permission | Does not check whether a user has the required permissions on the objects that are involved in the SQL statements to be executed. |
| @act.enable_sec_column_mask | De-identifies sensitive fields in query result sets that are returned for SQL statements that are submitted by users who do not have permissions on the sensitive fields. |
| @act.disable_sec_column_mask | Does not de-identify sensitive fields in query result sets that are returned for SQL statements that are submitted by users who do not have permissions on the sensitive fields. |

Supported SQL statements

| Category | Subcategory |
|----------|--|
| DQL | <ul style="list-style-type: none"> • SELECT • DESC • EXPLAIN • SHOW |
| DML | <ul style="list-style-type: none"> • INSERT • INSERT_SELECT • REPLACE • REPLACE_INT O • UPDATE • DELETE • MERGE |

| Category | Subcategory |
|----------|---|
| DDL | <ul style="list-style-type: none"> • DATABASE_OP • CREATE • CREATE_INDEX • CREATE_VIEW • CREATE_SEQUENCE • CREATE_TABLE • CREATE_SELECT • TRUNCATE • DROP_INDEX • DROP_VIEW • DROP_TABLE • RENAME • ALTER • ALTER_INDEX • ALTER_VIEW • ALTER_TABLE • ALTER_SEQUENCE • CREATE_FUNCTION • CREATE_PROCEDURE • ALTER_FUNCTION • ALTER_PROCEDURE • DROP_FUNCTION • DROP_PROCEDURE |
| DCL | <ul style="list-style-type: none"> • GRANT • DECLARE • SET • ANALYZE • FLUSH • OPTIMIZE • KILL |

15.11.9.4. SQLConsole for MongoDB

DMS allows you to manage relational and non-relational databases on the SQLConsole tab. The definition and classification of security rules are different for relational and non-relational databases. This topic describes the security rules for MongoDB databases on the SQLConsole tab.

Basic configuration items

Maximum number of returned rows per query: the maximum number of rows that can be returned for a query.

Supported checkpoints

 **Note** Various security rule templates are built in each checkpoint. You can configure security rules based on these templates or create custom security rules as needed.

| Checkpoint | Description |
|------------------------------------|---|
| User Permission Validation | Allows you to specify whether to check the permissions of specific users when they submit commands. |
| Collection Statement Criteria | Allows you to specify whether to allow DMS to run a specific category of commands. |
| DB Statement Criteria | |
| Cache Query Statement Criteria | |
| User Management Statement Criteria | |
| Role Management Statement Criteria | |
| Replication Set Statement Criteria | |
| Sharding Statement Criteria | |

Supported factors

A factor is a built-in variable in DMS. You can use factors to obtain the context to be validated by security rules, such as command categories and the number of rows to be affected. A factor name starts with the prefix `@fac.`. The following table describes the supported factors for MongoDB databases on the SQLConsole tab.

| Factor | Description |
|------------------------|---|
| @fac.sql_sub_type | The subcategory of the current command. For more information about the supported commands, see Supported MongoDB commands . |
| @fac.env_type | The type of the environment. The value is the display name of the environment type, such as <code>DEV</code> or <code>PRODUCT</code> . |
| @fac.current_sql | The current command. |
| @fac.user_is_admin | A Boolean value that indicates whether the current user is a DMS administrator. Valid values: <i>true</i> and <i>false</i> . |
| @fac.user_is_dba | A Boolean value that indicates whether the current user is a DBA. Valid values: <i>true</i> and <i>false</i> . |
| @fac.user_is_inst_dba | A Boolean value that indicates whether the current user is the DBA of the current instance. Valid values: <i>true</i> and <i>false</i> . |
| @fac.user_is_sec_admin | A Boolean value that indicates whether the current user is a security administrator. Valid values: <i>true</i> and <i>false</i> . |

Supported actions

An action in a security rule is an operation that DMS performs when the if condition in the rule is met. For example, DMS can forbid the submission of a ticket, approve a ticket, or reject a ticket. An action in a security rule denotes the purpose of the security rule. An action name starts with the prefix `@act.`. The following table describes the supported actions for MongoDB databases on the SQLConsole tab.

| Action | Description |
|------------------------------|---|
| @act.reject_execute | Rejects the request to run the current command. |
| @act.allow_execute | Allows the current command to be run. |
| @act.reject_sql_type_execute | Rejects the request to run a specific subcategory of commands. You must specify a subcategory after the action name. Example: <code>@act.reject_sql_type_execute 'UPDATE'</code> . |
| @act.allow_sql_type_execute | Allows a specific subcategory of commands to be run. You must specify a subcategory after the action name. |

Supported MongoDB commands

| Category | Subcategory | Command |
|---------------------|----------------------------------|--|
| Collection commands | Query commands | <ul style="list-style-type: none"> aggregate find findOne count distinct getIndex getShardDistribution isCapped stats dataSize storageSize totalIndexSize totalSize |
| | Data update commands | <ul style="list-style-type: none"> insert save findAndModify remove update |
| | Collection modification commands | <ul style="list-style-type: none"> drop renameCollection |
| | Index modification commands | <ul style="list-style-type: none"> createIndex createIndexes dropIndexes reIndex |
| | Other commands | validate |

| Category | Subcategory | Command |
|--|------------------------------|---|
| Database commands | Database query commands | <ul style="list-style-type: none"> • commandHelp • currentOp • getCollectionInfos • getCollectionNames • getLastError • getLastErrorObj • getLogComponents • getPrevError • getProfilingStatus • getReplicationInfo • getSiblingDB • help • isMaster • listCommands • printCollectionStats • printReplicationInfo • version • serverBuildInfo • serverStatus,stats |
| | Collection creation commands | createCollection |
| | High-risk commands | <ul style="list-style-type: none"> • dropDatabase • fsyncLock • fsyncUnlock • killOp • repairDatabase • resetError • runCommand |
| Commands related to the query plan cache | Read commands | <ul style="list-style-type: none"> • getPlanCache • getPlansByQuery • listQueryShapes |
| | Write commands | clearPlansByQuery |
| User management commands | User query commands | <ul style="list-style-type: none"> • getUser • getUsers |
| | User modification commands | <ul style="list-style-type: none"> • createUser • changeUserPassword • dropUser • dropAllUsers • grantRolesToUser • revokeRolesFromUser • updateUser |

| Category | Subcategory | Command |
|--------------------------|----------------------------|--|
| Role management commands | Role query commands | <ul style="list-style-type: none"> • getRole • getRoles |
| | Role modification commands | <ul style="list-style-type: none"> • createRole • dropRole • dropAllRoles • grantPrivilegesToRole • revokePrivilegesFromRole • revokeRolesFromRole • updateRole |
| Replica set commands | N/A | <ul style="list-style-type: none"> • help • printReplicationInfo • status • conf |
| Sharding commands | N/A | <ul style="list-style-type: none"> • getBalancerState • isBalancerRunning |

15.11.9.5. SQLConsole for Redis

DMS allows you to manage relational and non-relational databases on the SQLConsole tab. The definition and classification of security rules are different for relational and non-relational databases. This topic describes the security rules for Redis databases on the SQLConsole tab.

Supported checkpoints

 **Note** Various security rule templates are built in each checkpoint. You can configure security rules based on these templates or create custom security rules as needed.

| Checkpoint | Description |
|--|---|
| Permission Execution Statement Criteria | Allows you to set constraints on the permissions for command execution. |
| Statement Criteria: Keys | Allows you to specify whether to check the permissions of specific users when they submit commands. |
| Statement Criteria: String | Allows you to specify whether to allow the execution of various Redis commands. |
| Statement Criteria: List | |
| Statement Criteria: SET | |
| Statement Criteria: SortedSet | |
| Statement Criteria: Hash | |

| Checkpoint | Description |
|------------------------------|-------------|
| Statement Criteria: Other | |

Supported factors

A factor is a built-in variable in DMS. You can use factors to obtain the context to be validated by security rules, such as SQL statement categories and the number of rows to be affected. A factor name starts with the prefix `@fac.`. The following table describes the supported factors for Redis databases on the SQLConsole tab.

| Factor | Description |
|-----------------------|--|
| @fac.cmd_type | The type of the Redis command. For more information about valid values, see Supported Redis commands . |
| @fac.env_type | The type of the environment. The value is the display name of the environment type, such as <code>DEV</code> or <code>PRODUCT</code> . |
| @fac.is_read | A Boolean value that indicates whether the current command is a read command. Valid values: <i>true</i> and <i>false</i> . |
| @fac.is_write | A Boolean value that indicates whether the current command is a write command. Valid values: <i>true</i> and <i>false</i> . |
| @fac.current_sql | The current command. |
| @fac.user_is_admin | A Boolean value that indicates whether the current user is a DMS administrator. Valid values: <i>true</i> and <i>false</i> . |
| @fac.user_is_dba | A Boolean value that indicates whether the current user is a DBA. Valid values: <i>true</i> and <i>false</i> . |
| @fac.user_is_inst_dba | A Boolean value that indicates whether the current user is the DBA of the current instance. Valid values: <i>true</i> and <i>false</i> . |

Supported actions

An action in a security rule is an operation that DMS performs when the if condition in the rule is met. For example, DMS can forbid the submission of a ticket, approve a ticket, or reject a ticket. An action in a security rule denotes the purpose of the security rule. An action name starts with the prefix `@act.`. The following table describes the supported actions for Redis databases on the SQLConsole tab.

| Action | Description |
|---------------------|---|
| @act.reject_execute | Rejects the request to run the current command. |
| @act.allow_execute | Allows the current command to be run. |

Supported Redis commands

| Category | Subcategory | Command |
|----------|-------------|---------|
|----------|-------------|---------|

| Category | Subcategory | Command |
|-------------------------|------------------------------|--|
| Key-related commands | Key-related read commands | <ul style="list-style-type: none"> • EXISTS • TTL • PTTL • RANDOMKEY • TYPE • SCAN • OBJECTS |
| | Key-related write commands | <ul style="list-style-type: none"> • DEL • DUMP • EXPIRE • EXPIREART • MOVE • PERSIST • PEXPIRE • PEXPIREAT • RENAME • RENAMENX • RESTORE • SORT • TOUCH • UNLINK • WAIT • MIGRATE |
| String-related commands | String-related read commands | <ul style="list-style-type: none"> • GET • GETRANGE • BITCOUNT • GETBIT • MGET • STRLEN • BITOPS |
| | | |

| Category | Subcategory | Command |
|-----------------------|-------------------------------|--|
| | String-related write commands | <ul style="list-style-type: none"> • APPEND • BITFIELD • BITOP • DECR • DECRBY • GETSET • INCR • INCRBY • INCRBYFLOAT • MSET • MSETNX • PSETEX • SET • SETNX |
| List-related commands | List-related read commands | <ul style="list-style-type: none"> • LINDEX • LLEN • LRANGE |
| | List-related write commands | <ul style="list-style-type: none"> • BLPOP • BRPOP • BRPOPLPUSH • LINSERT • LPOP • LPUSH • LPUSHX • LREM • LSET • LTRIM • RTOP • RPOPLPUSH • RPUSH • RPUSHX |
| Set-related commands | Set-related read commands | <ul style="list-style-type: none"> • SCARD • SISMEMBER • SRANDMEMBER • SSCAN |
| | Set-related write commands | <ul style="list-style-type: none"> • SADD • SMOVE • SPOP • SREM |

| Category | Subcategory | Command |
|-----------------------------|-----------------------------------|--|
| Sorted set-related commands | Sorted set-related read commands | <ul style="list-style-type: none"> • ZCARD • ZCOUNT • ZLEXCOUNT • ZRANGE • ZRANGEBYLEX • ZRANGEBYSCORE • ZRANK • ZREVRNGE • ZREVRANGEBYLEX • ZREVRANGEBYSCORE • ZREVRANK • ZSCAN • ZSCORE |
| | Sorted set-related write commands | <ul style="list-style-type: none"> • ZADD • ZINCRBY • ZINTERSTORE • ZPOPMAX • ZPOPMIN • ZREM • ZUNIONSTORE • BZPOPMIN • BZPOPMAX |
| Hash-related commands | Hash-related read commands | <ul style="list-style-type: none"> • HEXISTS • HGET • HLEN • HMGET • HSCAN • HSTRLEN |
| | Hash-related write commands | <ul style="list-style-type: none"> • HDEL • HINCRBY • HINCRBYFLOAT • HMESET • HSET • HSETNX |

15.11.9.6. Data change

In DMS, you can execute SQL statements for data changes. However, the execution requires a high level of security. DMS allows you to configure security rules on the SQL Correct tab to validate the submission and approval of tickets for data changes. Only the SQL statements that are validated by the security rules can be executed.

Background information

Based on a DSL, new security rules are flexible to use. You can apply new security rules to define risk levels for tickets so that a ticket can be submitted to the approval process that is designed for the specified risk level. For more information, see [DSL syntax for security rules](#).

Basic configuration items

| Configuration item | Description |
|--|---|
| Data change default approval Template | By default, this approval template takes effect if you do not configure different approval rules for data changes at different risk levels under the Risk Approval Rules checkpoint. In the Switch Approval Template dialog box, you can change the approval process of the default approval template. For more information, see Customize approval processes . |
| Data Change risk level list | <p>This risk level list defines risk levels that are used in the Risk Identification Rules and Risk Approval Rules checkpoints to identify and classify risks in data changes. You can set risk levels based on the type and scenario of data changes. Data changes at different risk levels are submitted to different approval processes. DMS allows you to set the following four risk levels:</p> <ul style="list-style-type: none"> • <i>low</i>: a low risk level. • <i>middle</i>: a medium risk level. • <i>high</i>: a high risk level. • <i>highest</i>: a critical risk level. |

Supported checkpoints

 **Note** Various security rule templates are built in each checkpoint. You can configure security rules based on these templates or create custom security rules as needed.

| Checkpoint | Description |
|----------------------------|---|
| SQL execution rules | <p>SQL execution rules are used to limit the SQL statements that can be submitted for execution. If you do not enable SQL execution rules, all SQL statements that are used for data changes cannot be executed. Assume that you want to use DML statements to change the data of a database in an online environment. You can create the following SQL execution rule:</p> <p>Example:</p> <pre> if @fac.env_type not in ['product'] and @fac.sql_type in ['UPDATE','DELETE','INSERT'] then @act.allow_submit end </pre> <p>Note:</p> <p>The preceding rule specifies that you can only submit data change tickets to execute UPDATE, DELETE, and INSERT statements on a database that is deployed in an online environment.</p> |

| Checkpoint | Description |
|----------------------------------|---|
| Risk Identification Rules | <p>If a ticket conforms to the preset SQL execution rules, DMS continues to validate the ticket based on the risk identification rules. Risk identification rules are used to identify and classify risks in data changes. You can create risk identification rules based on your database environment, the number of rows in which data is affected, and the categories and subcategories of SQL statements.</p> <p>Note Different risk identification rules apply to different check items. DMS automatically identifies the highest risk level for a data change. For example, if the risk level of a data change is identified as high, medium, and low by one, three, and five risk identification rules, DMS assumes that the data change is at high risk.</p> <p>Example:</p> <pre>if @fac.env_type not in ['product','pre'] then @act.mark_risk 'low 'Low risk level: offline environment' end</pre> <p>Note: The preceding rule specifies that if the destination database is deployed in an offline environment, data changes are at low risk.</p> |
| Risk Approval Rules | <p>After the risk level of a data change is identified by the risk identification rules, DMS processes the ticket based on the risk approval rules. You can customize risk approval rules under the Risk Approval Rules checkpoint.</p> <p>Note</p> <ul style="list-style-type: none"> If a data change does not hit risk approval rules, DMS uses the default approval template that is specified under the Basic Configuration Item checkpoint to process the ticket. By default, an offline environment is identified as a factor at low risk and requires no approval. |
| Batch Data import rules | <p>These rules apply only to the validation of data import tickets. You can use the default rules that are provided in templates, or configure rules based on your actual needs.</p> |

Supported factors

A factor is a built-in variable in DMS. You can use factors to obtain the context to be validated by security rules, such as SQL statement categories and the number of rows to be affected. A factor name starts with the prefix `@fac.` The following table describes the supported factors on the SQL Correct tab.

| Factor | Description |
|---------------|--|
| @fac.env_type | The type of the environment. The value is the display name of the environment type, such as DEV or PRODUCT. |
| @fac.sql_type | The type of the SQL statement. The value is the subcategory of the SQL statement, such as UPDATE or INSERT. For more information, see Supported SQL statements . |

| Factor | Description |
|-----------------------------|--|
| @fac.detail_type | The type of the data change. Valid values: <ul style="list-style-type: none"> • <i>COMMON</i>: a Normal Data Modify ticket. • <i>CHUNK_DML</i>: a Lock-Free Data Modify ticket. • <i>PROCEDURE</i>: a Programmable Object ticket. • <i>CRON_CLEAR_DATA</i>: a History Data Clean ticket. • <i>BIG_FILE</i>: a Large Data Import ticket. |
| @fac.is_logic | A Boolean value that indicates whether the database to be affected is a logical database. |
| @fac.extra_info | The additional information about the data change. This factor is not in use. |
| @fac.is_ignore_affect_rows | A Boolean value that indicates whether to skip the validation. |
| @fac.insert_rows | The number of rows of data to be inserted. |
| @fac.update_delete_rows | The number of rows of data to be updated. |
| @fac.max_alter_table_size | The size of the largest tablespace where the table to be modified is stored. |
| @fac.is_has_security_column | A Boolean value that indicates whether the SQL statement to be executed involves sensitive fields. |
| @fac.security_column_list | A list of sensitive fields that the SQL statement to be executed involves. |
| @fac.risk_level | The risk level of the operation that is to be performed by the SQL statement. |
| @fac.risk_reason | The reason for identifying the operation to be performed as at the risk level. |

Supported actions

An action in a security rule is an operation that DMS performs when the if condition in the rule is met. For example, DMS can forbid the submission of a ticket, approve a ticket, or reject a ticket. An action in a security rule denotes the purpose of the security rule. An action name starts with the prefix `@act.`. The following table describes the supported action on the SQL Correct tab.

| Action | Description |
|--|---|
| @act.allow_submit | Requires the submission of SQL statements to be executed in a ticket. |
| @act.allow_execute_direct | Allows the execution of SQL statements in the SQLConsole. |
| @act.forbid_execute | Forbids the execution of SQL statements. |
| @act.mark_risk | Marks the risk level of a data change. Example: <code>@act.mark_risk 'Medium risk level: online environment'</code> . |
| @act.do_not_approve | Specifies the ID of an approval template. |
| @act.choose_approve_template | |
| @act.choose_approve_template_with_reason | |

15.11.9.7. Permission application

DMS allows you to configure security rules on the Access apply tab to validate applications for permissions, including permissions on instances, databases, and tables.

Background information

In DMS, security rules are flexible to use. You can apply security rules to define risk levels for tickets so that a ticket can be submitted to the approval process that is designed for the specified risk level. For more information about the DSL syntax, see [DSL syntax for security rules](#).

Basic configuration items

The following table describes the basic configuration items that are supported on the Access apply tab.

| Configuration item | Description |
|--|---|
| [Instance-permission application] default approval Template | By default, this approval template takes effect if you do not set different approval processes for instance permission applications at different risk levels under the Validation for Instance Permission Application checkpoint. <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e0f0ff;"> <p> Note In the Switch Approval Template dialog box, you can change the approval process of the default approval template.</p> </div> |
| [DB-permission application] default approval Template | By default, this approval template takes effect if you do not set different approval processes for database permission applications at different risk levels under the Database Permission Application Validation checkpoint. |
| Table-permission request default approval Template | By default, this approval template takes effect if you do not set different approval processes for table permission applications at different risk levels under the Table Permission Application Validation checkpoint. |
| [Programmable object-permission application] default approval Template | By default, this approval template takes effect if you do not set different approval processes for programmable object permission applications at different risk levels under the Programmable object verification checkpoint. |
| [Field-permission application] default approval Template | By default, this approval template takes effect if you do not set different approval processes for sensitive field permission applications at different risk levels under the Sensitive Field Application Validation checkpoint. |
| Line-permission application default approval Template | By default, this approval template takes effect if you do not set different approval processes for row permission applications at different risk levels under the Line permission application verification checkpoint. |
| [Owner-application] default approval template (when the resource has no Owner) | By default, this approval template takes effect if you do not set different approval processes for data ownership applications at different risk levels under the Owner Application Validation checkpoint and the data that is involved in the application has no owner. |
| [Owner-application] default approval template (when the resource has an Owner) | By default, this approval template takes effect if you do not set different approval processes for data ownership applications at different risk levels under the Owner Application Validation checkpoint and the data that is involved in the application has one or more owners. |

Supported checkpoints

When a user submits a ticket to apply for permissions, DMS checks whether the ticket conforms to rules that are specified under checkpoints. The ticket can be submitted only after DMS determines that the ticket conforms to all rules that are specified under checkpoints.

Note Various security rule templates are built in each checkpoint. You can configure security rules based on these templates or create custom security rules as needed.

| Checkpoint | Description |
|--|--|
| Owner Application Validation | Allows you to set approval processes or constraints for Instance-OWNER, Table-OWNER, and Database-OWNER tickets. |
| Validation for Instance Permission Application | Allows you to set approval processes or constraints for Instance-Performance and Instance-Login tickets. |
| Database Permission Application Validation | Allows you to set approval processes or constraints for Database-Permission tickets. |
| Table Permission Application Validation | Allows you to set approval processes or constraints for Table-Permission tickets. |
| Programmable object verification | Allows you to set approval processes or constraints for Programmable Object tickets. |
| Sensitive Field Application Validation | Allows you to set approval processes or constraints for Sensitive Column-Permission tickets. |
| Line permission application verification | Allows you to set approval processes or constraints for Row-Permission tickets. |

Supported factors

A factor is a built-in variable in DMS. You can use factors to obtain the context to be validated by security rules, such as SQL statement categories and database names. A factor name starts with the prefix `@fac.`. The following table describes the supported factors on the Access apply tab.

| Factor | Description |
|----------------------------|--|
| @fac.env_type | The type of the environment. The value is the display name of the environment type, such as DEV or PRODUCT. |
| @fac.schema_name | The name of the database. |
| @fac.perm_apply_duration | The period of time during which the applicant needs the permission. Unit: hours. |
| @fac.column_security_level | The security level of the field. Valid values: <ul style="list-style-type: none"> <i>sensitive</i> <i>confidential</i> <i>inner</i> |

Supported actions

An action in a security rule is an operation that DMS performs when the if condition in the rule is met. For example, DMS can forbid the submission of a ticket, approve a ticket, or reject a ticket. An action in a security rule denotes the purpose of the security rule. An action name starts with the prefix `@act.`. The following table describes the supported action on the Access apply tab.

| Action | Description |
|--|---|
| @act.forbid_submit_order | Forbids a ticket from being submitted. |
| @act.do_not_approve | Specifies the ID of an approval template. |
| @act.choose_approve_template | |
| @act.choose_approve_template_with_reason | |

15.11.9.8. Data export

DMS allows you to manage security rules on the Data Export tab to validate the permissions of applicants on involved databases, tables, sensitive fields, and rows during the submission and approval of data export tickets. This helps ensure data security.

Basic configuration items

Data export default approval Template: the default approval template that takes effect if you do not set different approval processes for data export tickets at different risk levels under the Approval Rule Validation checkpoint. In the Switch Approval Template dialog box, you can change the approval process of the default approval template. For more information, see [Customize approval processes](#).

Supported checkpoints

 **Note** Various security rule templates are built in each checkpoint. You can configure security rules based on these templates or create custom security rules as needed.

| Checkpoint | Description |
|---------------------------------|---|
| Pre-check Validation | Allows you to specify whether to validate the permissions of applicants on involved databases, tables, sensitive fields, and rows by configuring security rules. |
| Approval Rule Validation | Allows you to submit data export tickets to different approval processes by configuring security rules. For example, you can submit tickets for exporting more than a specific number of rows to an approval process and other tickets to another approval process. |

Supported factors

A factor is a built-in variable in DMS. You can use factors to obtain the context to be validated by security rules, such as SQL statement categories and the number of rows to be affected. A factor name starts with the prefix @fac. The following table describes the supported factors on the Data Export tab.

| Factor | Description |
|----------------------------------|---|
| @fac.env_type | The type of the environment. The value is the display name of the environment type, such as DEV or PRODUCT. |
| @fac.is_ignore_export_rows_check | A Boolean value that indicates whether to skip the check on the number of rows to be affected. |
| @fac.export_rows | The number of rows to be exported. |
| @fac.include_sec_columns | A Boolean value that indicates whether the data to be exported contains sensitive fields. |

| Factor | Description |
|------------------------|---|
| @fac.sec_columns_list | The sensitive fields that require or do not require approval before data is exported. The sensitive fields are displayed in the format of <code>Table name.Field name,[Table name.Field name, ...]</code> . |
| @fac.user_is_admin | A Boolean value that indicates whether the applicant is a DMS administrator. |
| @fac.user_is_dba | A Boolean value that indicates whether the applicant is a DBA. |
| @fac.user_is_inst_dba | A Boolean value that indicates whether the applicant is the DBA of the current instance. |
| @fac.user_is_sec_admin | A Boolean value that indicates whether the applicant is a security administrator. |

Supported actions

An action in a security rule is an operation that DMS performs when the if condition in the rule is met. For example, DMS can forbid the submission of a ticket, approve a ticket, or reject a ticket. An action in a security rule denotes the purpose of the security rule. An action name starts with the prefix `@act.`. The following table describes the supported actions on the Data Export tab.

| Action | Description |
|--|---|
| @act.do_not_approve | Allows a ticket to be processed without approval. |
| @act.choose_approve_template | Specifies an approval template. |
| @act.choose_approve_template_with_reason | Specifies an approval template with a reason provided. |
| @act.forbid_submit_order | Forbids a ticket from being submitted. |
| @act.enable_check_permission | Validates the permissions of an applicant on involved databases and tables. |
| @act.disable_check_permission | Does not validate the permissions of an applicant on involved databases and tables. |
| @act.enable_check_sec_column | Validates the permissions of an applicant on involved sensitive fields. |
| @act.disable_check_sec_column | Does not validate the permissions of an applicant on involved sensitive fields. |

15.11.9.9. Schema design

DMS allows you to configure security rules on the Schema Design tab to check the design rules and risk identification rules that apply to schema design tickets. This helps ensure data security.

Basic configuration items

| Configuration item | Description |
|---|---|
| Enable non-peer Publishing | <p>Specifies whether to enable non-peer publishing. By default, data changes to a table can be published only to a table with the same name in another database. After you enable non-peer publishing, you can perform data changes on all tables.</p> <p> Warning This feature may bring high risks. We recommend that you proceed with caution and enable this feature only for special requirements.</p> |
| R & D process | The whole process of a schema design ticket. It is the most important configuration item on the Schema Design tab. For more information about the parameters of the configuration item, see Parameters involved in the R&D process . |
| Field type configuration | The supported data types of fields to be added. |
| Index type configuration | The supported data types of indexes to be added. |
| It is forbidden to modify the original field data type | Specifies whether to prohibit the data types of the original fields from being modified when the original table is to be modified. |
| Prohibit deleting original fields | <p>Specifies whether to prohibit the existing fields from being deleted when the original table is to be modified.</p> <p> Note We recommend that you enable this feature because deleting existing fields may bring high risks.</p> |
| Prohibit renaming original fields | <p>Specifies whether to prohibit the existing fields from being renamed when the original table is to be modified.</p> <p> Note We recommend that you enable this feature because renaming existing fields may bring high risks.</p> |
| Table character set license configuration | The range of character sets that are allowed to be used when you create a table. For example, you can specify utf8 and utf8mb4. |
| Default approval template for Structural design | The default approval template that is used for a schema design ticket if you do not configure the Approval Rule Validation checkpoint. In the Switch Approval Template dialog box, you can change the approval process of the default approval template. |
| When published, the ticket will automatically advance to the end state | <p>The point that is used to stop the schema change process. If you enable this feature, after the node that is set as the anchor in the R&D process is run, DMS automatically turns the ticket to the Finished state.</p> <p> Note To use this feature, you must set the last node in the R&D process as the anchor.</p> |

Supported checkpoints

 **Note** Various security rule templates are built in each checkpoint. You can configure security rules based on these templates or create custom security rules as needed.

The Schema Design tab contains the following two processes:

- Process of saving changes: DMS provides the following three checkpoints for this process. The checkpoints validate the table headings, fields, and indexes.
 - Save Changes and Validate Header
 - Save Changes and Validate Field
 - Save Changes and Validate Index
- Process of applying changes: DMS provides the following five checkpoints for this process. The first four checkpoints identify the risks that arise from changing schemas without locking tables, and the last checkpoint assigns an approval process to each type of risk.
 - Table Creation Risk Control
 - Field Change Risk Control
 - Index Change Risk Control
 - SQL Execution Risk Control
 - Approval Rule Validation

Supported factors

A factor is a built-in variable in DMS. You can use factors to obtain the context to be validated by security rules, such as SQL statement categories and the number of rows to be affected. A factor name starts with the prefix `@fac.`. The following table describes the supported factors on the Schema Design tab.

| Factor | Description |
|----------------------------------|--|
| @fac.table_kind | The type of the table whose schema is to be changed. Valid values: <ul style="list-style-type: none"> • <i>new</i>: a newly created table. • <i>old</i>: an existing table. |
| @fac.column_kind | The type of the field to be changed. Valid values: <ul style="list-style-type: none"> • <i>new</i>: a newly created field. • <i>old</i>: an existing field. |
| @fac.xxxx_old | The value of an existing field or index that is used for comparison. |
| @fac.column_is_primary | A Boolean value that indicates whether the current field serves as a primary key. Valid values: <i>true</i> and <i>false</i> . |
| @fac.column_type_support_default | A Boolean value that indicates whether the data type of the current field supports a default value. Valid values: <i>true</i> and <i>false</i> . <div style="background-color: #e0f2f1; padding: 5px; margin-top: 10px;"> <p> Note For example, a field of the CHAR type supports a default value, whereas a field of the TEXT type does not.</p> </div> |
| @fac.index_kind | The type of the index to be changed. Valid values: <ul style="list-style-type: none"> • <i>new</i>: a newly created index. • <i>old</i>: an existing index. |
| @fac.index_column_count | The number of fields in the index. |

| Factor | Description |
|-------------------------|---|
| @fac.change_type | The type of the schema change to be performed by DDL statements. Valid values: <ul style="list-style-type: none"> <i>add</i>: adds one or more fields or indexes. <i>modify</i>: modifies one or more fields or indexes. <i>delete</i>: deletes one or more fields or indexes. |
| @fac.altered_table_size | The size of the table whose schema is to be changed. Unit: MB. |
| @fac.online_execute | A Boolean value that indicates whether the schema change can be performed in an online environment. Valid values: <i>true</i> and <i>false</i> . |
| @fac.change_risk_level | The risk level of the schema change. Valid values: <ul style="list-style-type: none"> <i>high</i>: a high risk level. <i>middle</i>: a medium risk level. <i>low</i>: a low risk level. |
| @fac.env_type | The type of the environment. The value is the display name of the environment type, such as DEV or PRODUCT. |

Supported actions

An action in a security rule is an operation that DMS performs when the if condition in the rule is met. For example, DMS can forbid the submission of a ticket, approve a ticket, or reject a ticket. An action in a security rule denotes the purpose of the security rule. An action name starts with the prefix `@act.`. The following table describes the supported actions on the Schema Design tab.

| Action | Description | Format |
|------------------------------|--|--|
| @act.block_submit | Blocks the submission of the schema change and displays the error message. This action can be used in the process of saving changes. | @act.block_submit 'Reason for blocking the submission' |
| @act.show_warning | Displays the error message without blocking the submission of the schema change. This action can be used in the process of saving changes. | @act.show_warning 'Error message' |
| @act.mark_middle_risk | Specifies that the schema change is at medium risk. This action can be used in the process of identifying the risk level. | @act.mark_middle_risk 'Reason for the identification' |
| @act.mark_high_risk | Specifies that the schema change is at high risk. This action can be used in the process of identifying the risk level. | @act.mark_high_risk 'Reason for the identification' |
| @act.forbid_submit_publish | Rejects the ticket. This action can be used in the process of setting the approval process. | @act.forbid_submit_publish 'Reason for the rejection' |
| @act.do_not_approve | Specifies the ID of an approval | |
| @act.choose_approve_template | | |

| Action | Specifies the ID of an approval template | N/A Format |
|--|--|------------|
| @act.choose_approve_template_with_reason | | |

Parameters involved in the R&D process

| Parameter | Description |
|----------------------|--|
| Step | <ul style="list-style-type: none"> The type of the node. Valid values: Design: The design node in the R&D process is generated by default and cannot be removed. It determines the environment where the schema change is designed. Publish: A publish node in the R&D process is used to publish the schema change after the change is designed. You can set multiple publish nodes. |
| Node Name | The name of the node. The node name can be up to 10 characters in length. |
| Database Environment | The environment where the node is run. |
| Execution Strategy | <ul style="list-style-type: none"> The way in which the node is run. Valid values: Immediately: The node is run immediately after it is approved. Periodically: The node is run at the time that you specify. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p> Note If a node is approved before the specified point in time, it is run as scheduled. Otherwise, the node is interrupted and not run.</p> </div> |
| Can Go Back | Specifies whether a publish node can be rolled back to the design node. |
| Can Skip | Specifies whether the current node can be skipped. |
| Anchor | The point that is used to stop the schema change process. If you set a node as the anchor, after the node is published, the nodes that follow the anchor cannot be run and the schema change process ends. At this time, the ticket enters the Published state. |
| Actions | The operation that you can perform on a publish node. You can remove a publish node as required. |

15.11.9.10. Database and table synchronization

DMS allows you to configure security rules on the Table Sync tab to validate operations that are related to schema synchronization, empty database initialization, and table consistency repair.

Basic configuration items

| Configuration item | Description |
|--|--|
| Enable execution capability | Specifies whether to enable SQL-based synchronization. If this configuration item is set to OFF, applicants can compare table schemas but cannot execute SQL statements to synchronize databases and tables. Other configuration items and security rules you set under checkpoints on the Table Sync tab also become invalid. |
| Database table synchronization default approval Template | The default approval template for database and table synchronization applications. You can use the default approval template or click Switch Approval Template and select another template. For more information, see Customize approval processes . |

| Configuration item | Description |
|--|--|
| Analysis phase script Expiration Time (unit: hours) | The timeout period of the analysis phase. You can set an appropriate timeout period in which synchronization can be canceled if schemas are changed in the destination database. |

Supported checkpoints

The Table Sync tab contains three checkpoints that are corresponding to the three features that are supported by the tab. The three checkpoints are unrelated to each other. For example, when you submit a Schema Synchronization ticket, only the basic configuration items and the security rules that are specified under the Schema Synchronization Validation checkpoint are used to validate the ticket.

 **Note** Various security rule templates are built in each checkpoint. You can configure security rules based on these templates or create custom security rules as needed.

| Checkpoint | Description |
|---|--|
| Schema Synchronization Validation | Allows you to set approval processes or constraints for Schema Synchronization tickets. |
| Empty Database Initialization Validation | Allows you to set approval processes or constraints for Empty Database Initialization tickets. |
| Table Consistency Repair Validation | Allows you to set approval processes or constraints for Repair Table Consistency tickets. |

Supported factors

A factor is a built-in variable in DMS. You can use factors to obtain the context to be validated by security rules, such as SQL statement categories and the number of rows to be affected. A factor name starts with the prefix `@fac.`. The following table describes the supported factors on the Table Sync tab.

| Factor | Description |
|-------------------------------|---|
| <code>@fac.env_type</code> | The type of the environment. The value is the display name of the environment type, such as DEV or PRODUCT. |
| <code>@fac.schema_name</code> | The name of the schema. |

Supported actions

An action in a security rule is an operation that DMS performs when the if condition in the rule is met. For example, DMS can forbid the submission of a ticket, approve a ticket, or reject a ticket. An action in a security rule denotes the purpose of the security rule. An action name starts with the prefix `@act.`. The following table describes the supported actions on the Table Sync tab.

| Action | Description |
|---|---|
| <code>@act.forbid_submit_order</code> | Forbids a ticket from being submitted. The statement is in the following format: <code>@act.forbid_submit_order 'Reason for forbidding the submission of the ticket'</code> . |
| <code>@act.do_not_approve</code> | Specifies the ID of an approval template. |
| <code>@act.choose_approve_template</code> | |

| Action | Description |
|--|-------------|
| @act.choose_approve_template_with_reason | |

15.11.9.11. Sensitive field change

The topic describes the security rules on the Sensitive Column Change tab.

Basic configuration items

Sensitive column default approval Template: the default approval template that takes effect if you do not set approval processes for tickets that apply to change the security levels of sensitive fields under the Approval Rule Validation checkpoint. In the Switch Approval Template dialog box, you can change the approval process of the default approval template.

Supported checkpoints

Approval Rule Validation: When a user submits a ticket to change the security level of a sensitive field, DMS checks whether the ticket conforms to the rules that are specified under the Approval Rule Validation checkpoint.

 **Note** Various security rule templates are built in each checkpoint. You can configure security rules based on these templates or create custom security rules as needed.

Supported factors

A factor is a built-in variable in DMS. You can use factors to obtain the context to be validated by security rules, such as SQL statement categories and the number of rows to be affected. A factor name starts with the prefix @fac. The following table describes the supported factors on the Sensitive Column Change tab.

| Factor | Description |
|-------------------------------|--|
| @fac.column_level_change_type | <p>The type of security level change that the applicant wants to perform on a sensitive field. Valid values:</p> <ul style="list-style-type: none"> • <i>upper</i>: raises the current security level, including the following three cases: <ul style="list-style-type: none"> ◦ From inner to sensitive ◦ From inner to confidential ◦ From sensitive to confidential • <i>sensitive_to_inner</i>: lowers the security level from sensitive to inner. • <i>confidential_to_sensitive</i>: lowers the security level from confidential to sensitive. • <i>confidential_to_inner</i>: lowers the security level from confidential to inner. |

Supported actions

An action in a security rule is an operation that DMS performs when the if condition in the rule is met. For example, DMS can forbid the submission of a ticket, approve a ticket, or reject a ticket. An action in a security rule denotes the purpose of the security rule. An action name starts with the prefix @act. The following table describes the supported actions on the Sensitive Column Change tab.

| Action | Description |
|--------------------------|--|
| @act.forbid_submit_order | Forbids a ticket from being submitted. The statement is in the following format: @act.forbid submit order 'Reason for forbidding the submission of the ticket' . |

| Action | Description |
|--|---|
| @act.do_not_approve | Specifies the ID of an approval template. |
| @act.choose_approve_template | |
| @act.choose_approve_template_with_reason | |

15.11.9.12. Test data generation

This topic describes the security rules on the Test Data Generate tab.

Supported checkpoints

Approval rule validation: When a user submits a ticket to generate test data, DMS checks whether the ticket conforms to the rules that are specified under the Approval rule validation checkpoint.

 **Note** Various security rule templates are built in each checkpoint. You can configure security rules based on these templates or create custom security rules as needed.

Supported factors

A factor is a built-in variable in DMS. You can use factors to obtain the context to be validated by security rules, such as SQL statement categories and the number of rows to be affected. A factor name starts with the prefix @fac. The following table describes the supported factors on the Test Data Generate tab.

| Factor | Description |
|------------------|---|
| @fac.env_type | The type of the environment. The value is the display name of the environment type, such as DEV or PRODUCT. |
| @fac.schema_name | The name of the schema. |

Supported actions

An action in a security rule is an operation that DMS performs when the if condition in the rule is met. For example, DMS can forbid the submission of a ticket, approve a ticket, or reject a ticket. An action in a security rule denotes the purpose of the security rule. An action name starts with the prefix @act. The following table describes the supported actions on the Test Data Generate tab.

| Action | Description |
|--|--|
| @act.forbid_submit_order | Forbids a ticket from being submitted. The statement is in the following format: @act.forbid submit order 'Reason for forbidding the submission of the ticket' . |
| @act.do_not_approve | Specifies the ID of an approval template. |
| @act.choose_approve_template | |
| @act.choose_approve_template_with_reason | |

15.11.9.13. Database cloning

This topic describes the security rules on the Database Clone tab.

Basic configuration items

Database clone default approval Template: the default approval template that takes effect if you do not set approval processes for database clone tickets under the Approval rule validation checkpoint. In the Switch Approval Template dialog box, you can change the approval process of the default approval template.

Supported checkpoints

Approval rule validation: When a user submits a database clone ticket, DMS checks whether the ticket conforms to the rules that are specified under the Approval rule validation checkpoint.

 **Note** Various security rule templates are built in each checkpoint. You can configure security rules based on these templates or create custom security rules as needed.

Supported actions

An action in a security rule is an operation that DMS performs when the if condition in the rule is met. For example, DMS can forbid the submission of a ticket, approve a ticket, or reject a ticket. An action in a security rule denotes the purpose of the security rule. An action name starts with the prefix `@act.`. The following table describes the supported actions on the Database Clone tab.

| Action | Description |
|---|---|
| <code>@act.forbid_submit_order</code> | Forbids a ticket from being submitted. The statement is in the following format: <code>@act.forbid_submit_order 'Reason for forbidding the submission of the ticket'</code> . |
| <code>@act.do_not_approve</code> | Specifies the ID of an approval template. |
| <code>@act.choose_approve_template</code> | |
| <code>@act.choose_approve_template_with_reason</code> | |

16. Server Load Balancer (SLB)

16.1. What is SLB?

This topic provides an overview of Server Load Balancer (SLB). SLB distributes inbound network traffic across multiple Elastic Compute Service (ECS) instances that act as backend servers based on forwarding rules. You can use SLB to improve the responsiveness and availability of your applications.

Overview

After you add ECS instances that are deployed in the same region to a SLB instance, SLB uses virtual IP addresses (VIPs) to virtualize these ECS instances into backend servers in a high-performance server pool that ensures high availability. Client requests are distributed to the ECS instances based on forwarding rules.

SLB checks the health status of the ECS instances and automatically removes unhealthy ones from the server pool to eliminate single points of failure (SPOFs). This enhances the resilience of your applications.

Components

SLB consists of three components:

- SLB instances

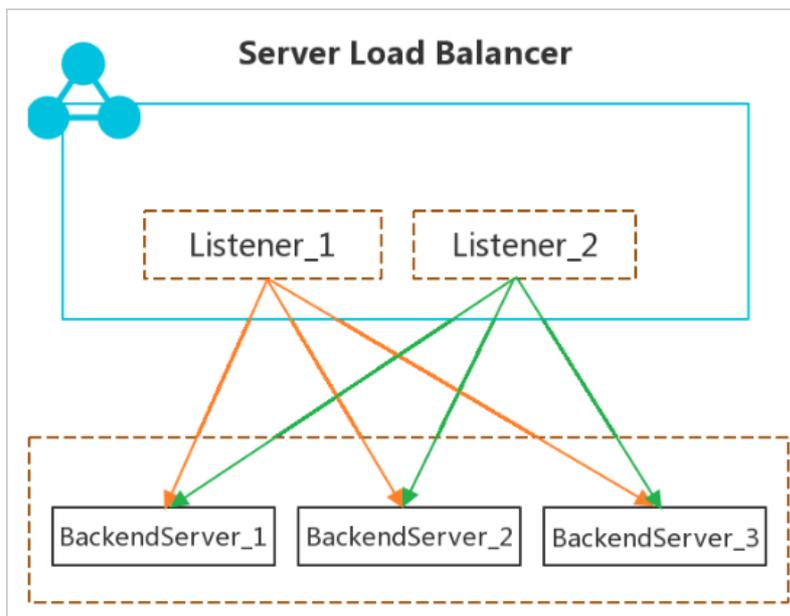
A SLB instance is a running SLB service entity that receives traffic and distributes traffic to backend servers. To get started with SLB, you must create a SLB instance and add at least one listener and two ECS instances to the SLB instance.

- Listeners

A listener checks client requests and forwards them to backend servers. It also performs health checks on backend servers.

- Backend servers

ECS instances are used as backend servers to receive distributed requests. You can separately add ECS instances to the server pool, or use vServer groups or primary/secondary server groups to add and manage ECS instances in batches.



Benefits

- High availability

SLB is designed with full redundancy that avoids SPOFs and supports zone-disaster recovery.

SLB can be scaled based on application loads and can provide continuous service during traffic fluctuations.

- High scalability

You can increase or decrease the number of backend servers to adjust the load balancing capability of your applications.

- Cost-effectiveness

SLB can save 60% of load balancing costs compared with using traditional hardware solutions.

- Security

You can use SLB with Apsara Stack Security to defend your applications against up to 5 Gbit/s DDoS attacks.

- High concurrency

A SLB cluster supports hundreds of millions of concurrent connections and a single SLB instance supports tens of millions of concurrent connections.

16.2. Log on to the SLB console

This topic describes how to go to the Server Load Balancer (SLB) console after you log on to the Apsara Uni-manager Management Console by using the Chrome browser.

Prerequisites

- The domain name of the Apsara Uni-manager Management Console is obtained from the engineer that deploys the service before you log on to the Apsara Uni-manager Management Console.
- We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

 **Note** When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.

4. In the top navigation bar, choose **Products > Networking > Server Load Balancer**.

16.3. Quick start

16.3.1. Overview

This topic describes how to create an Internet SLB instance and forward client requests to two ECS instances that act as backend servers.

 **Note** Before you create an SLB instance, you must plan its configurations such as the instance type, region, and billing method. For more information, see [Make preparations](#).

Perform the following operations:

1. [Create a CLB instance](#)

Create an SLB instance. An SLB instance is an entity running the SLB service.

2. [Add listeners and backend servers](#)

After you create an SLB instance, you must add listeners and backend servers.

3. [Release an SLB instance](#)

If an SLB instance is no longer needed, delete the instance to avoid extra fees.

16.3.2. Make preparations

This topic describes the considerations for configuring a Server Load Balancer (SLB) instance. Before you create an SLB instance, you must determine the listener type and network type of the SLB instance.

Select the region of the SLB instance

When you select a region, take note of the following issues:

- To reduce latency and increase the download speed, we recommend that you select a region that is closest to your customers.
- To provide more stable and reliable load balancing services, SLB supports primary/secondary zone deployment in most regions. This implements disaster recovery across data centers in the same region. We recommend that you select a region that supports primary/secondary zone deployment.
- SLB does not support cross-region deployment. Therefore, you must deploy the SLB instance and backend Elastic Compute Service (ECS) instances in the same region.

Select the network type of the SLB instance (Internet-facing or internal-facing)

SLB provides load balancing services for the public network and internal network:

- If you want to use SLB to distribute requests from the Internet, create an Internet-facing SLB instance.
An Internet-facing SLB instance is assigned a public IP address to receive requests from the Internet.
- If you want to use SLB to distribute requests from the internal network, create an internal-facing SLB instance.
An internal-facing SLB instance is assigned a private IP address and is only accessible from the internal network.

Select an instance type

Apsara Stack released guaranteed-performance SLB instances on April 1, 2018. When you create an SLB instance, you can choose a guaranteed-performance instance that provides exclusive resources and higher service availability. SLB provides six types of instances:

- For a pay-as-you-go instance, we recommend that you select the instance type that provides the highest specifications. This guarantees flexible load balancing services at no extra costs. However, if the capacity of Super I (slb.s3.large), the highest-performance instance type, far exceeds the demand of your business, you can select a more appropriate type such as Higher II (slb.s3.medium).

Select a listener protocol

SLB supports Layer 4 (TCP and UDP) and Layer 7 (HTTP and HTTPS) load balancing:

- A Layer 4 listener directly distributes requests to backend servers without modifying packet headers. After a client request reaches a Layer 4 listener, SLB uses the backend port that is configured for the listener to establish

a TCP connection with a backend ECS instance.

- A Layer 7 listener functions as a reverse proxy. After a client request reaches a Layer 7 listener, SLB establishes a new TCP connection over HTTP with a backend server, instead of directly forwarding the request to the backend server (ECS instance).

Compared with Layer 4 listeners, Layer 7 listeners require an additional step of T Engine processing. Therefore, Layer 4 listeners provide better performance than Layer 7 listeners. In addition, the performance of Layer 7 listeners may be affected by factors such as insufficient client ports or excessive backend server connections. We recommend that you use Layer 4 listeners for high performance purposes.

Create backend servers

Before you use the SLB service, you must create ECS instances, deploy applications on the ECS instances, and then add the ECS instances to your SLB instance to process client requests.

When you create and configure an ECS instance, take note of the following issues:

- Region and zones of the ECS instance

Make sure that the ECS instance is deployed in the same region as the SLB instance. We recommend that you deploy ECS instances in different zones to improve availability.

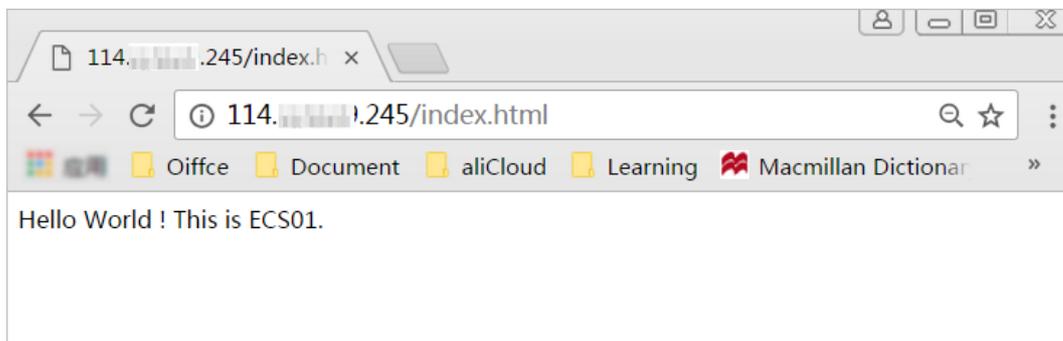
In this example, two ECS instances named ECS01 and ECS02 are created in the **China (Hangzhou)** region. The following figure shows their basic configurations.

| Instance ID/Name | Tag | Monitoring | Zone | IP Address | Status | Network Type | Specifications | Billing Method | Actions |
|------------------|-----|------------|-----------------|--|---------|--------------|---|---|---|
| ECS01 | | | Hangzhou Zone I | 1...internet) 1...(Private) | Running | VPC | 2 vCPU 8 GiB (I/O Optimized) ecs.g6.large 50Mbps (Peak Value) | Subscription 16 September 2021, 23:59 Expire | Manage Upgrade/Downgrade Renew More |
| ECS02 | | | Hangzhou Zone F | 4...internet) 1...internal Network | Running | VPC | 1 vCPU 1 GiB (I/O Optimized) ecs.xn4.small 53Mbps | Subscription 21 December 2020, 23:59 Expire | Manage Upgrade/Downgrade Renew More |

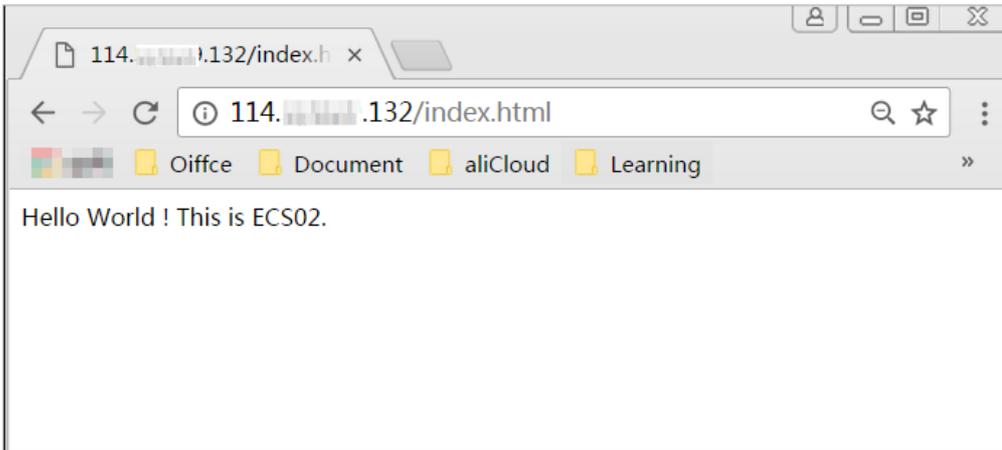
- Configurations

In this example, two static web pages are hosted on ECS01 and ECS02 by using Apache.

- Enter the elastic IP address (EIP) that is associated with ECS01 in the address bar of your browser.



- Enter the EIP that is associated with ECS02 in the address bar of your browser.



No additional configurations are required after you deploy applications on the ECS instances. However, if you want to use a Layer 4 (TCP or UDP) listener and the ECS instances run Linux, make sure that the following parameters in the `net.ipv4.conf` file under `/etc/sysctl.conf` are set to 0:

```
net.ipv4.conf.default.rp_filter = 0
net.ipv4.conf.all.rp_filter = 0
net.ipv4.conf.eth0.rp_filter = 0
```

16.3.3. Create an SLB instance

This topic describes how to create a Server Load Balancer (SLB) instance. An SLB instance is a running entity of the SLB service. You can add multiple listeners and backend servers to an SLB instance.

Prerequisites

- Elastic Compute Service (ECS) instances are created and applications are deployed on the ECS instances.
- The ECS instances and the SLB instance belong to the same organization. In addition, the security group rules configured for the ECS instances allow access from port 80 (HTTP) and port 443 (HTTPS).

Procedure

1. [Log on to the SLB console](#).
2. In the left-side navigation pane, choose **Instances > Instances**.
3. On the **Instances** page, click **Create Instance**.
4. Configure the SLB instance and click **Submit**.

| Parameter | Description |
|--------------|--|
| Organization | Select the organization to which the SLB instance belongs. Note Make sure that the SLB instance and its backend servers belong to the same organization. |
| Resource Set | Select the resource group to which the SLB instance belongs. |
| Region | Select the region where you want to deploy the SLB instance. |
| Zone | Select the zone where you want to deploy the SLB instance. |

| Parameter | Description |
|------------------|--|
| Quantity | Select the number of SLB instances that you want to purchase. |
| Instance Name | Enter a name for the SLB instance. If you set Quantity to a value greater than 1, the system automatically assigns names to the SLB instances. |
| Instance Edition | Select one of the following options: <ul style="list-style-type: none"> ◦ Shared-resource: Shared-resource SLB instances share resources with each other. The performance of shared-resource SLB instances is not guaranteed. ◦ High-performance: High-performance SLB instances use exclusive resources. The performance of high-performance SLB instances varies by specification. |
| Instance Type | Select the type of network traffic that you want to distribute. Valid values: Internal Network and Internet. Internal Network is selected in this example. |
| Network Type | Select the network type of the SLB instance. Valid values: Classic Network and VPC. VPC is selected in this example. |
| IP Version | Select an IP version. |
| VPC | Select a virtual private cloud (VPC). |
| vSwitch | Select a vSwitch. |
| Billing Method | Set the billing method. |

What's next

[Configure an SLB instance](#)

16.3.4. Configure an SLB instance

This topic describes how to configure a Server Load Balancer (SLB) instance. Before an SLB instance can forward traffic, you must add at least one listener and one group of backend servers to the SLB instance. The following example shows how to add a TCP listener and Elastic Compute Service (ECS) instances to an SLB instance. The ECS instances are ECS 01 and ECS 02. These ECS instances function as backend servers that host static web pages.

Procedure

1. [Log on to the SLB console](#).
2. On the **Instances** page, find the SLB instance that you want to manage and click **Configure Listener** in the Actions column.
3. On the **Protocol and Listener** wizard page, set the following parameters to configure the listener. Use the default settings for other parameters. Click **Next**.
 - **Select Listener Protocol**: Select a listener protocol. **TCP** is selected in this example.
 - **Listening Port**: Specify a frontend port to receive and distribute requests to backend servers. In this example, the port is set to **80**.

The SLB instance uses this port to provide external services. In most cases, port 80 is set for HTTP listeners and port 443 is set for HTTPS listeners.

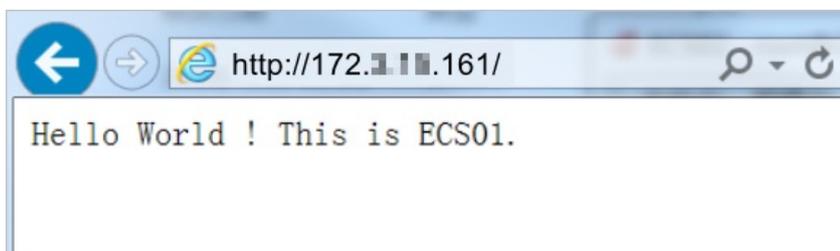
Advanced:

- **Enable Peak Bandwidth Limit** : Applications that run on backend ECS instances provide external services. You can set bandwidth caps to limit service capabilities of applications.
 - **Scheduling Algorithm**: SLB supports the following scheduling algorithms. **Round-Robin (RR)** is selected in this example.
 - **Weighted Round-Robin (WRR)**: Requests are distributed to backend servers in sequence. Backend servers that have higher weights receive more requests.
 - **Round-Robin (RR)**: Requests are distributed to backend servers in sequence.
 - **Consistent Hash (CH)**: Only high-performance SLB instances support the CH algorithm.
 - **Source IP**: specifies consistent hashing that is based on source IP addresses. Requests from the same source IP address are distributed to the same backend server.
 - **Tuple**: specifies consistent hashing that is based on four factors: source IP address, destination IP address, source port number, and destination port number. Requests that contain the same information based on the four factors are distributed to the same backend server.
4. On the **Backend Servers** wizard page, select **Default Server Group** and click **Add More** to add backend servers.
 - i. In the **Available Servers** panel, select ECS 01 and ECS 02 and click **Next**.
 - ii. A backend server that has a higher weight receives more requests. The default value is 100. We recommend that you use the default value.
 - iii. Click **Add**.
 - iv. On the **Default Server Group** tab, specify backend ports that are available to receive requests. The ports are used by backend ECS instances to receive requests. You can specify the same port for multiple backend servers that are added to the same SLB instance. The port is set to 80 in this example.
 5. Click **Next** to configure the health check feature. The default health check settings are used in this example. After you enable health checks for the SLB instance, the SLB instance periodically checks whether the backend ECS instances are healthy. When the SLB instance detects an unhealthy ECS instance, the SLB instance distributes the requests to other healthy ECS instances. When the unhealthy ECS instance recovers, the SLB instance starts to distribute requests to the ECS instance again.
 6. Click **Next**. On the **Confirm** wizard page, check the configurations and click **Submit**.
 7. Click **OK** to go back to the **Instances** page and click  to refresh the page.

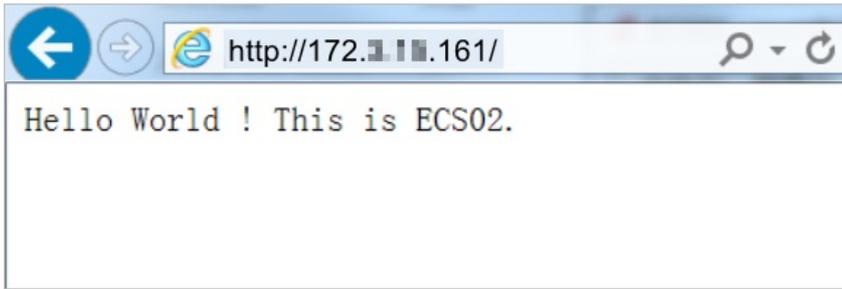
If the health check state of a backend ECS instance is **Active**, this indicates that the backend server works as expected and can process requests.

8. Enter the IP address of the SLB instance into the address bar of the browser to test load balancing services of the instance.

ECS01



ECS02



16.3.5. Release an SLB instance

If an SLB instance is no longer needed, you can release the instance to save costs. The backend ECS instances will not be deleted or affected after you delete an SLB instance.

Procedure

1. Log on to the SLB console.
2. On the **Instances** page, find the instance and click  > **Release** in the Actions column, or select the instance and click **Release** at the lower part of the page.
3. In the **Release** dialog box, select **Release Now**.

 **Note** The system performs release operations at 30-minute and hour marks. However, billing for the SLB instance is stopped at the specified release time.

4. Click **Next**.
5. Click **OK** to release the SLB instance.

 **Note** Pay-as-you-go SLB instances cannot be restored once deleted. We recommend that you exercise caution when you release SLB instances.

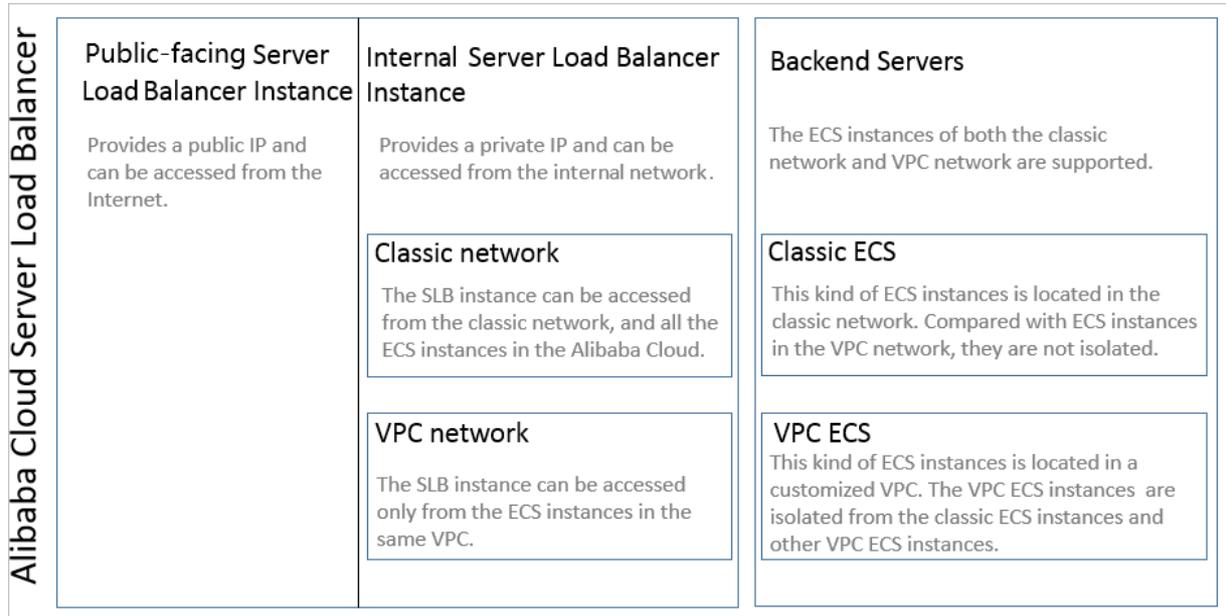
16.4. SLB instances

16.4.1. SLB instance overview

A Server Load Balancer (SLB) instance is a running entity of the SLB service. To use the SLB service, you must create an SLB instance and add listeners and backend servers to the instance.

Instance network types

Apsara Stack provides Internet-facing and internal-facing SLB instances. You can create Internet-facing SLB instances or internal-facing SLB instances based on your business requirements. If you create an Internet-facing SLB instance, a public IP address is allocated. If you create an internal-facing SLB instance, a private IP address is allocated.



• Internet-facing SLB instances

An Internet-facing SLB instance distributes client requests over the Internet to backend servers based on configured forwarding rules. An internal-facing SLB instance in a virtual private cloud (VPC) can process requests sent over the Internet only when the instance is associated with an elastic IP address (EIP). The following table provides more detailed information.

| Type | Feature |
|---|--|
| <p>Internet-facing SLB instances</p> <p>After you create an Internet-facing SLB instance, the system allocates a public IP address to the instance. You can bind a domain name to the public IP address to provide external services.</p> | <ul style="list-style-type: none"> ◦ The public IP address is allocated to the SLB instance and cannot be unbound from the instance. ◦ Internet-facing SLB instances support only pay-by-data-transfer and pay-by-bandwidth billing methods. |
| <p>Internal-facing SLB instances that are associated with EIPs</p> <p>An internal-facing SLB instance that is associated with an EIP can process requests sent over the Internet.</p> | <ul style="list-style-type: none"> ◦ A public IP address is allocated to the EIP. You can associate the EIP with the SLB instance and disassociate the EIP from the SLB instance based on your requirements. ◦ An EIP that is associated with an EIP bandwidth plan supports the 95th percentile bandwidth billing method. |

• Internal-facing SLB instances

Internal-facing SLB instances can be used only inside Apsara Stack and can forward only requests from clients that can access SLB instances over the internal network.

You can select one of the following network types for an internal-facing SLB instance:

- Classic network

If you choose classic network for an internal-facing SLB instance, the IP address of the SLB instance is allocated and maintained by Apsara Stack. This instance can be accessed only by Elastic Compute Service (ECS) instances in the classic network.

- VPC

If you choose VPC for an internal-facing SLB instance, the IP address of the SLB instance is allocated from the CIDR block of the vSwitch that is attached to the VPC. This SLB instance can be accessed only by ECS instances in the VPC.

Instance types and specifications

Alibaba Cloud provides shared-performance SLB instances and guaranteed-performance SLB instances. Guaranteed-performance SLB instances provide reliable performance metrics.

- Shared-performance SLB instances

All shared-performance SLB instances share SLB resources. This indicates that the instance performance cannot be guaranteed.

- Guaranteed-performance SLB instances

The following content describes three key metrics of guaranteed-performance SLB instances:

- Max Connection

The maximum number of concurrent connections that an SLB instance supports. When the number of concurrent connections reaches the specified limit, new connection requests are dropped.

- Connections Per Second (CPS)

The number of new connections that are established per second. When the CPS value reaches the specified limit, new connection requests are dropped.

- Queries Per Second (QPS)

The number of HTTP or HTTPS queries (requests) that can be processed per second. This metric is specific to Layer 7 listeners. When the QPS value reaches the specified limit, new connection requests are dropped.

Apsara Stack provides four types of guaranteed-performance SLB instances.

| Type | Specification | Max connection | CPS | QPS | Purchase method |
|--------|-----------------------------|----------------|--------|--------|---|
| Type 1 | Small I (slb.s1.small) | 5,000 | 3,000 | 1,000 | Available for purchase from the official website of Apsara Stack. |
| Type 2 | Standard I (slb.s2.small) | 50,000 | 5,000 | 5,000 | Available for purchase from the official website of Apsara Stack. |
| Type 3 | Standard II (slb.s2.medium) | 100,000 | 10,000 | 10,000 | Available for purchase from the official website of Apsara Stack. |

| Type | Specification | Max connection | CPS | QPS | Purchase method |
|--------|-------------------------|----------------|--------|--------|---|
| Type 4 | Higher I (slb.s3.small) | 200,000 | 20,000 | 20,000 | Available for purchase from the official website of Apsara Stack. |

The following table describes the differences between shared-performance SLB instances and guaranteed-performance SLB instances.

| Feature | Shared-performance SLB instance | Guaranteed-performance SLB instance |
|--|---------------------------------|-------------------------------------|
| Resource allocation | Shared resources | Exclusive resources |
| Service level agreement for guaranteed availability | Not supported | 99.95% |
| IPv6 | × | √ |
| Server Name Indication (SNI) certificates | × | √ |
| Support for blacklists and whitelists | × | √ |
| Elastic network interface (ENI) mounting | × | √ |
| Assignment of secondary IP addresses to ENIs that are bound to ECS instances | × | √ |
| HTTP-to-HTTPS redirection | × | √ |
| Consistent hashing | × | √ |
| TLS security policies | × | √ |
| HTTP2 | × | √ |
| Websocket(S) | × | √ |

16.4.2. Create an SLB instance

This topic describes how to create a Server Load Balancer (SLB) instance. An SLB instance is a running entity of the SLB service. You can add multiple listeners and backend servers to an SLB instance.

Prerequisites

- Elastic Compute Service (ECS) instances are created and applications are deployed on the ECS instances.
- The ECS instances and the SLB instance belong to the same organization. In addition, the security group rules of the ECS instances allow access from port 80 (HTTP) and port 443 (HTTPS).

Procedure

1. [Log on to the SLB console](#).
2. In the left-side navigation pane, choose **Instances > Instances**.

3. On the **Instances** page, click **Create Instance**.

- **Organization**: Select an organization for the SLB instance from the drop-down list.

 **Note** Make sure that the organization of the SLB instance is the same as the organization of its backend servers.

- **Resource Set**: Select a resource set for the SLB instance from the drop-down list.
- **Region**: Select the region where you want to deploy the SLB instance.
- **Zone**: Select a zone for the SLB instance from the drop-down list.
- **Instance Name**: Enter a name for the SLB instance in the Instance Name field.
The name must be 2 to 128 characters in length, and can contain letters, digits, full-width characters, hyphens (-), colons (:), periods (.), and underscores (). Line breaks and spaces are supported. It must start with a letter and cannot start with `http://` or `https://`.
- **Instance Edition**: Select one of the following options: shared-performance and guaranteed-performance. Shared-performance SLB instances share resources with each other. The performance of shared-performance SLB instances is not guaranteed. The performance of guaranteed-performance SLB instances varies by type.
- **Instance Type**: Select the type of network traffic that you want to distribute. Valid values: Internal Network and Internet.
- **Network Type**: Select the network type of the SLB instance. Valid values: Classic Network and VPC.
- **IP Version**: Select an IP version.
- **IP Address**: Enter a service IP address for the SLB instance. Make sure that the service IP address is valid. Otherwise, the SLB instance cannot be created. If you do not set this parameter, the system automatically allocates an IP address to the SLB instance.

 **Note** The private IP address that you specify must belong to the destination CIDR block of the vSwitch.

4. Click **Submit**.

What's next

[Configure a CLB instance](#)

16.4.3. Start and stop an SLB instance

This topic describes how to start and stop an SLB instance. SLB instances can be started or stopped at any time. A stopped SLB instance does not receive or forward client traffic.

Procedure

1. [Log on to the SLB console](#).
2. In the left-side navigation pane, choose **Instances > Server Load Balancers**.
3. Find the target SLB instance. In the **Actions** column, choose  > **Start** or  > **Stop**.
4. To start or stop multiple instances at a time, select the instances and click **Start** or **Stop** at the bottom of the page.

16.4.4. Tags

16.4.4.1. Tag overview

This topic provides an overview of tags in SLB. SLB provides the tag management feature that allows you to classify SLB instances by using tags.

Each tag consists of a key and a value. Before you use tags, note the following limits:

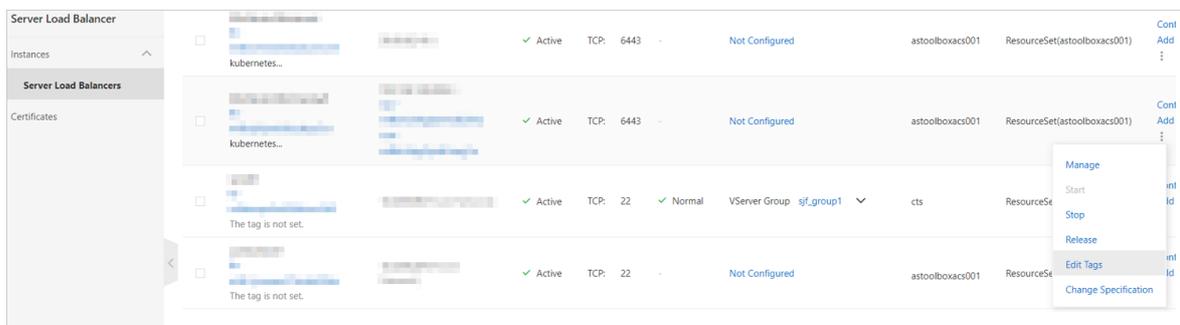
- Tags must be added to SLB instances.
- Each SLB instance can have a maximum of ten tags. You can add or remove a maximum of 5 tags at a time.
- The key of each tag added to an SLB instance must be unique. If a tag with the same key already exists, the tag is overwritten with the new value.

16.4.4.2. Add tags

This topic describes how to add tags to an SLB instance.

Procedure

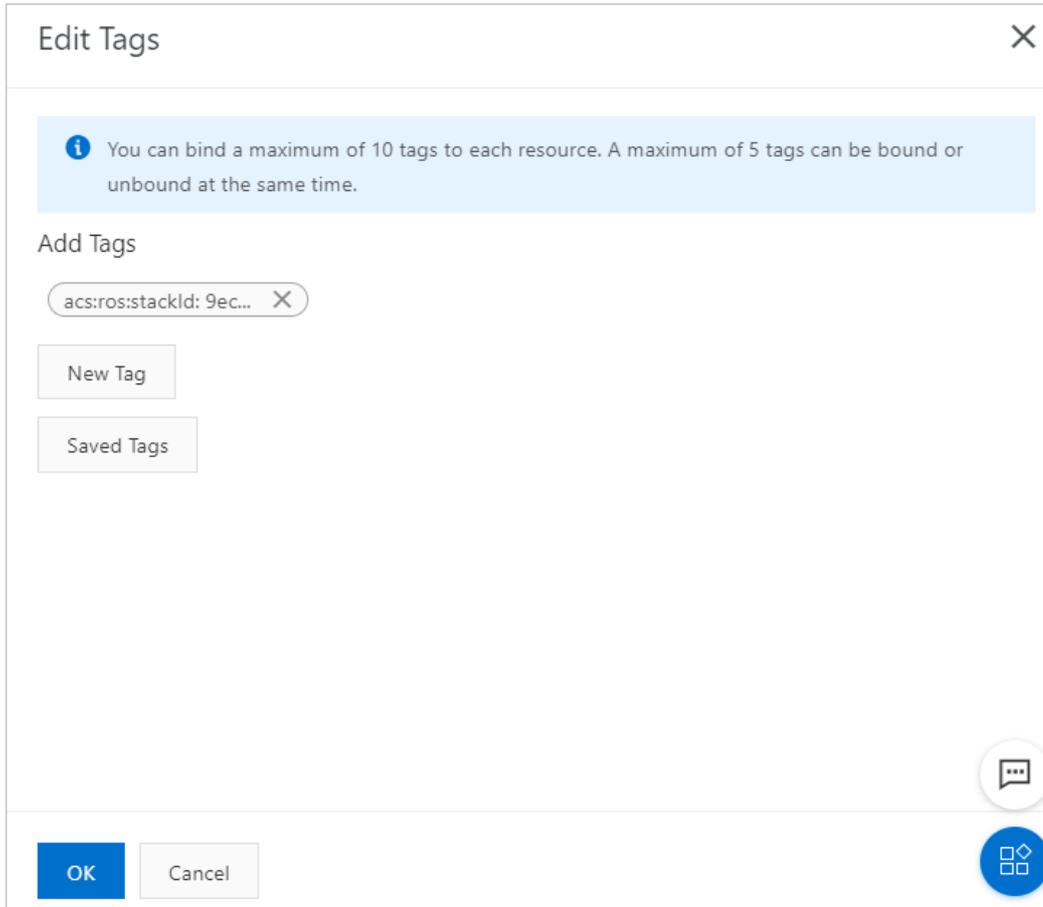
1. [Log on to the SLB console](#).
2. In the left-side navigation pane, choose **Instances > Server Load Balancers**.
3. In the **Actions** column, choose  > **Edit Tags**.



4. Edit tags in the **Edit Tags** dialog box.

To add a tag, perform the following operations:

- To add an existing tag, click **Saved Tags** and then select a tag.
- To create and add a new tag, click **New Tag** in the **Edit Tags** dialog box, enter the key and value of the new tag, and then click **OK** next to the value.



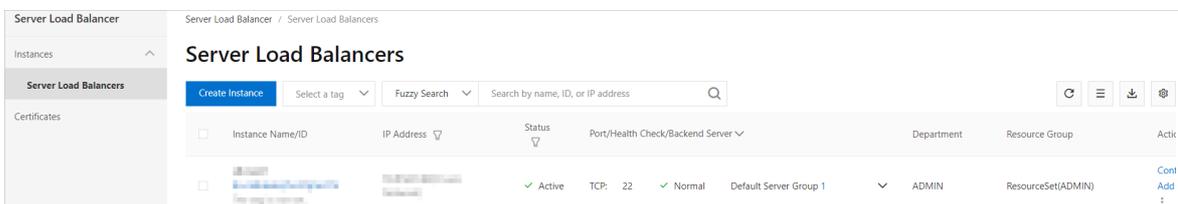
5. Click OK.

16.4.4.3. Query SLB instances by tag

This topic describes how to use tags to query SLB instances.

Procedure

1. [Log on to the SLB console](#).
2. In the left-side navigation pane, choose **Instances > Server Load Balancers**.
3. Select a data from the **Select a tag** drip-down list to filter instances.



Note To clear the search condition, move the pointer over the selected tag and click the displayed deletion icon next to it.

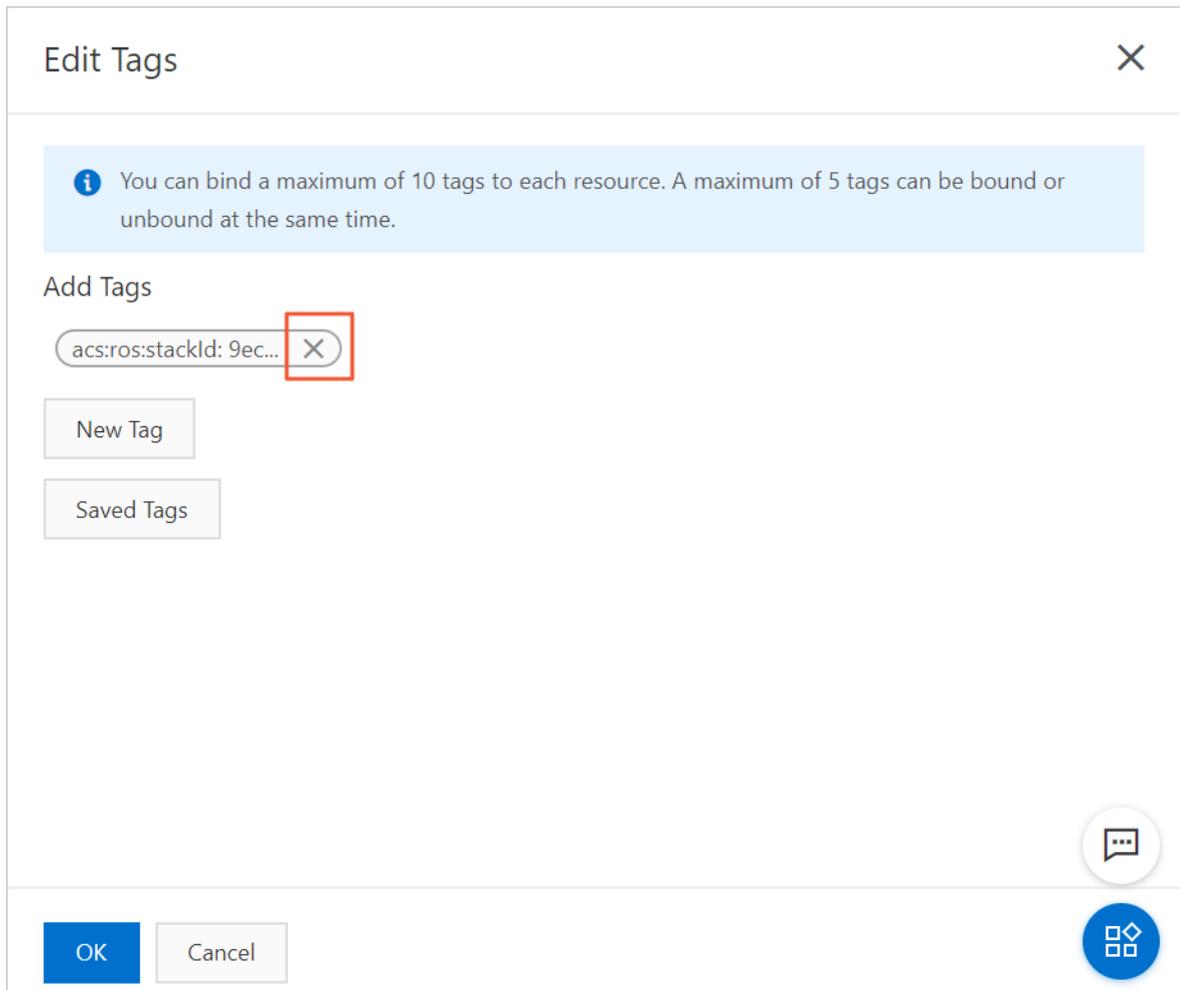
16.4.4.4. Remove a tag

This topic describes how to remove tags from a CLB instance. You can only remove tags for one CLB instance at a time.

Procedure

1. [Log on to the SLB console](#).
2. In the left-side navigation pane, choose **Instances > Server Load Balancers**.
3. In the **Actions** column, choose  **> Edit Tags**.
4. In the **Edit Tags** dialog box, click the deletion icon next to the tags to be removed, and then click **OK**.

 **Note** If a tag is removed from a CLB instance and is not added to other instances, the tag is deleted from the system.



16.4.5. Release an SLB instance

This topic describes how to release an SLB instance. You can release SLB instances immediately.

Procedure

1. [Log on to the SLB console](#).
2. In the left-side navigation pane, choose **Instances > Server Load Balancers**.
3. Find the target SLB instance and choose  **> Release**.
4. In the **Release** dialog box, select **Release Now**.
5. Click **Next**.

6. Confirm the displayed information and click **OK** to release the instance.

16.5. Listeners

16.5.1. Listener overview

This topic provides an overview of listeners. After you create a Server Load Balancer (SLB) instance, you must configure one or more listeners for it. A listener checks for connection requests and then distributes the requests to backend servers based on the forwarding rules that are defined by a specified scheduling algorithm.

SLB provides Layer 4 (TCP and UDP) and Layer 7 (HTTP and HTTPS) listeners. The following table lists the features and use cases of these listeners.

| Protocol | Feature | Use case |
|----------|---|---|
| TCP | <ul style="list-style-type: none"> A connection-oriented protocol. A logical connection must be established before data can be sent and received. Source IP address-based session persistence. Source IP addresses readable at the network layer. Fast data transmission | <ul style="list-style-type: none"> Applicable to scenarios that require high reliability and data accuracy but can tolerate low speeds, such as file transmission, sending or receiving emails, and remote logons. Web applications that do not have special requirements. <p>For more information, see Add a TCP listener.</p> |
| UDP | <ul style="list-style-type: none"> A connectionless protocol. UDP transmits data packets directly instead of making a three-way handshake with the other party before UDP sends data. UDP does not provide error recovery or data re-transmission. Fast data transmission but relatively low reliability. | <p>Applicable to scenarios where real-time transmission is more important than reliability, such as video chats and real-time financial market pushes.</p> <p>For more information, see Add a UDP listener.</p> |
| HTTP | <ul style="list-style-type: none"> An application-layer protocol that is used to package data. Cookie-based session persistence. Use X-Forward-For to obtain source IP addresses. | <p>Applicable to scenarios that require data content to be identified, such as web applications and small mobile games.</p> <p>For more information, see Add an HTTP listener.</p> |
| HTTPS | <ul style="list-style-type: none"> Encrypted data transmission that prevents unauthorized access. Centralized certificate management service. You can upload certificates to SLB. The decryption operations are directly completed on SLB. | <p>Applicable to scenarios that require encrypted transmission.</p> <p>For more information, see Add an HTTPS listener.</p> |

16.5.2. Add a TCP listener

This topic describes how to add a Transmission Control Protocol (TCP) listener to a Server Load Balancer (SLB) instance. TCP provides reliable and accurate data delivery at relatively low connection speeds. Therefore, TCP applies to file transmission, email sending or receiving, and remote logons. You can add a TCP listener to forward TCP requests.

Step 1: Open the listener configuration wizard

To open the listener configuration wizard, perform the following operations:

1. [Log on to the SLB console](#).
2. In the left-side navigation pane, choose **Instances > Instances**.
3. Use one of the following methods to open the listener configuration wizard:
 - On the **Instances** page, find the SLB instance and click **Configure Listener** in the **Actions** column.
 - On the **Instances** page, click the ID of the SLB instance. On the **Listener** tab, click **Add Listener**.

Step 2: Configure the TCP listener

To configure the TCP listener, perform the following operations:

1. Set the following parameters and click **Next**.

| Parameter | Description |
|-----------------------------------|---|
| Select Listener Protocol | Select the protocol of the listener. In this example, TCP is selected. |
| Listening Port | Specify the listening port that is used to receive requests and forward them to backend servers. Valid values: 1 to 65535. |
| Advanced | |
| Scheduling Algorithm | <p>SLB supports the following scheduling algorithms:</p> <ul style="list-style-type: none"> ◦ Weighted Round-Robin (WRR): Backend servers that have higher weights receive more requests than backend servers that have lower weights. ◦ Round-Robin (RR): Requests are distributed to backend servers in sequence. ◦ Consistent Hash (CH): <ul style="list-style-type: none"> ▪ Source IP: specifies consistent hashing that is based on source IP addresses. Requests from the same source IP address are distributed to the same backend server. ▪ Tuple: specifies consistent hashing that is based on four factors: source IP address, destination IP address, source port, and destination port. Requests that contain the same information based on the four factors are distributed to the same backend server. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note Only high-performance SLB instances support the CH algorithm.</p> </div> |
| Enable Session Persistence | <p>Specify whether to enable session persistence.</p> <p>After session persistence is enabled, SLB forwards all requests from a client to the same backend server.</p> <p>For TCP listeners, session persistence is implemented based on IP addresses. Requests from the same IP address are forwarded to the same backend server.</p> |

| Parameter | Description |
|---|---|
| Enable Peak Bandwidth Limit | <p>Specify whether to set a bandwidth limit for the listener.</p> <p>If an SLB instance is billed based on bandwidth usage, you can set different bandwidth limit values for different listeners. This limits the amount of network traffic that flows through each listener. The sum of the bandwidth limit values of all listeners that are added to an SLB instance cannot exceed the bandwidth limit of this SLB instance.</p> <p>By default, this feature is disabled and all listeners share the bandwidth of the SLB instance.</p> |
| Idle Timeout | Specify the timeout of idle TCP connections. Unit: seconds. Valid values: 10 to 900. |
| Obtain Client Source IP Address | Backend servers associated with Layer 4 listeners can obtain client IP addresses without additional configurations. |
| Automatically Enable Listener After Creation | Specify whether to immediately enable the listener after it is created. By default, this feature is enabled. |

Step 3: Add backend servers

After you configure the listener, you must add backend servers to process client requests. You can use the default server group that is configured for the SLB instance, or create a vServer group or a primary/secondary server group. For more information, see [Overview](#).

Backend servers are added to the default server group in this example.

1. On the **Backend Servers** wizard page, select **Default Server Group** and click **Add More**.
2. In the **My Servers** panel, select the ECS instances that you want to add as backend servers and click **Next**.
3. On the **Configure Ports and Weights** wizard page, specify the weights of the backend servers that you want to add. A backend server with a higher weight receives more requests.

 **Note** If the weight of a backend server is set to 0, no request is distributed to the backend server.

4. Click **Add**. On the **Default Server Group** tab, specify the ports that you want to open on the backend servers to receive requests. The backend servers are the ECS instances that you selected. Valid values: 1 to 65535.
You can specify the same port on different backend servers that are added to an SLB instance.
5. Click **Next**.

Step 4: Configure health checks

SLB performs health checks to check the availability of the ECS instances that serve as backend servers. The health check feature improves overall service availability and reduces the impact of backend server failures. Click **Modify** to configure advanced health check settings and click **Next**. For more information, see [Health check overview](#).

Step 5: Submit the configurations

To submit the configurations, perform the following operations:

1. On the **Confirm** wizard page, check the configurations. You can click **Modify** to modify the configurations.
2. After you confirm the configurations, click **Submit**.
3. In the **Configuration Successful** message, click **OK**.

After you configure the listener, you can view the listener on the **Listener** tab.

16.5.3. Add a UDP listener

This topic describes how to add a User Datagram Protocol (UDP) listener to a Server Load Balancer (SLB) instance. UDP applies to services that prioritize real-time content delivery over reliability, such as video conferencing and real-time quote services. You can add a UDP listener to forward UDP packets.

Context

Before you configure a UDP listener, take note of the following items:

- You are not allowed to specify ports 250, 4789, or 4790 for UDP listeners. They are system reserved ports.
- Fragmentation is not supported.
- You cannot view source IP addresses by using the UDP listeners of an SLB instance in the classic network.
- The following operations take effect 5 minutes after they are performed on a UDP listener:
 - Remove backend servers.
 - Set the weight of a backend server to 0 after it is detected unhealthy.

Step 2: Configure the UDP listener

To configure the listener, perform the following operations:

1. On the **Protocol and Listener** wizard page, set the following parameters and click **Next**.

| Parameter | Description |
|--------------------------|---|
| Listener Protocol | Select a protocol for the listener. In this example, UDP is selected. |
| Listening Port | Set the listening port that is used to receive requests and forward them to backend servers. Valid values: 1 to 65535. |
| Advanced | |

| Parameter | Description |
|--|--|
| Scheduling Algorithm | <p>SLB supports the following scheduling algorithms:</p> <ul style="list-style-type: none"> ◦ Weighted Round-Robin (WRR): Backend servers that have higher weights receive more requests than backend servers that have lower weights. ◦ Round-Robin (RR): Requests are distributed to backend servers in sequence. ◦ Consistent Hash (CH): <ul style="list-style-type: none"> ▪ Source IP: specifies consistent hashing that is based on source IP addresses. Requests from the same source IP address are distributed to the same backend server. ▪ Tuple: specifies consistent hashing that is based on four factors: source IP address, destination IP address, source port, and destination port. Requests that contain the same information based on the four factors are distributed to the same backend server. ▪ QUIC ID: specifies consistent hashing that is based on Quick UDP Internet Connections (QUIC) IDs. Requests that contain the same QUIC ID are distributed to the same backend server. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Notice QUIC is implemented based on draft-ietf-quic-transport-10 and is rapidly evolving. Therefore, compatibility is not guaranteed for all QUIC versions. We recommend that you perform tests before you apply the protocol to a production environment.</p> </div> |
| Enable Session Persistence | <p>Specify whether to enable session persistence.</p> <p>SLB maintains the persistence of UDP sessions by using consistent hashing that is based on source IP addresses.</p> |
| Enable Peak Bandwidth Limit | <p>Specify whether to set a bandwidth limit for the listener.</p> <p>If an SLB instance is billed based on bandwidth usage, you can specify different bandwidth limit values for different listeners. This limits the amount of network traffic that flows through each listener. The sum of the bandwidth limit values of all listeners that are added to an SLB instance cannot exceed the bandwidth limit of this SLB instance.</p> <p>By default, this feature is disabled and all listeners share the bandwidth of the SLB instance.</p> |
| Obtain Client Source IP Address | <p>Backend servers associated with UDP listeners can obtain client IP addresses without additional configurations.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note You cannot view source IP addresses by using the UDP listeners of an SLB instance in the classic network.</p> </div> |
| Automatically Enable Listener After Creation | <p>Specify whether to immediately enable the listener after it is created. By default, this feature is enabled.</p> |

Step 3: Add backend servers

After you configure the listener, you must add backend servers to process client requests. You can use the default server group that is configured for the SLB instance, or create a vServer group or a primary/secondary server group. For more information, see [Overview](#).

Backend servers are added to the default server group in this example.

1. On the **Backend Servers** wizard page, select **Default Server Group**. Then, click **Add More**.
2. In the **Available Servers** panel, select the Elastic Compute Service (ECS) instances that you want to add as backend servers and click **Next**.
3. Set weights for the backend servers that you add.

An ECS instance with a higher weight receives more requests.

 **Note** If the weight of a backend server is set to 0, no request is distributed to the backend server.

4. Click **Add**. On the **Default Server Group** tab, specify the ports that you want to open on the backend servers to receive requests. The backend servers are the ECS instances that you selected. Valid values: 1 to 65535.
You can specify the same port on different backend servers that are added to an SLB instance.
5. Click **Next**.

Step 4: Configure health checks

SLB performs health checks to check the availability of the ECS instances that serve as backend servers. The health check feature improves overall service availability and reduces the impact of backend server failures. Click **Modify** to configure advanced health check settings and click **Next**. For more information, see [Health check overview](#).

Step 5: Submit the configurations

To submit the configurations, perform the following operations:

1. On the **Confirm** wizard page, check the configurations. You can click **Modify** to modify the configurations.
2. After you confirm the configurations, click **Submit**.
3. In the **Configuration Successful** message, click **OK**.

After you configure the listener, you can view the listener on the **Listener** tab.

16.5.4. Add an HTTP listener

This topic describes how to add an HTTP listener to a Server Load Balancer (SLB) instance. HTTP is applicable to applications that must identify data from different users, such as web applications and mobile games. You can add HTTP listeners to forward HTTP requests.

Step 1: Configure an HTTP listener

1. [Log on to the SLB console](#).
2. Use one of the following methods to open the listener configuration wizard:
 - On the **Instances** page, find the SLB instance that you want to manage and click **Configure Listener** in the **Actions** column.
 - On the **Instances** page, click the ID of the SLB instance that you want to manage. On the **Listener** tab, click **Add Listener**.
3. Set the following parameters to configure the listener.

| Parameter | Description |
|--------------------------|--|
| Listener Protocol | Select a protocol for the listener. In this example, HTTP is selected. |

| Parameter | Description |
|------------------------------------|--|
| Listening Port | Specify the listening port that is used to receive requests and forward them to backend servers. Valid values: 1 to 65535. |
| Advanced | Click Modify to configure advanced settings. |
| Scheduling Algorithm | <p>Select a scheduling algorithm.</p> <ul style="list-style-type: none"> ◦ Weighted Round-Robin (WRR): Backend servers that have higher weights receive more requests than backend servers that have lower weights. ◦ Round-Robin (RR): Requests are distributed to backend servers in sequence. |
| Redirection | <p>Specify whether to redirect traffic from the HTTP listener to an HTTPS listener.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note Before you enable redirection, make sure that an HTTPS listener is created.</p> </div> |
| Enable Session Persistence | <p>Specify whether to enable session persistence.</p> <p>After session persistence is enabled, SLB forwards all requests from a client to the same backend server.</p> <p>HTTP session persistence is implemented through cookies. You can use one of the following methods to handle cookies:</p> <ul style="list-style-type: none"> ◦ Insert cookie: If you select this option, you need only to specify the timeout period of the cookie. <p>SLB inserts a cookie (SERVERID) into the first HTTP or HTTPS response that is sent to a client. The next request from the client contains the cookie. Then, the listener distributes the request to the recorded backend server.</p> <ul style="list-style-type: none"> ◦ Rewrite cookie: If you select this option, you can specify the cookie that you want to insert into an HTTP or HTTPS response. You must specify the timeout period and the lifetime of a cookie on a backend server. <p>When SLB detects a user-defined cookie, SLB rewrites the original cookie with the user-defined cookie. The next request from the client contains the user-defined cookie. Then, the listener forwards the request to the recorded backend server.</p> |
| Enable Peak Bandwidth Limit | <p>Specify whether to set a bandwidth limit for the listener. Unit: Mbit/s. Valid values: 0 to 5120.</p> <p>If an SLB instance is billed based on bandwidth usage, you can set different bandwidth limit values for different listeners. This limits the amount of network traffic that flows through each listener. The sum of the bandwidth limit values of all listeners that are added to an SLB instance cannot exceed the bandwidth limit of this SLB instance. By default, this feature is disabled and all listeners share the bandwidth of the SLB instance.</p> |

| Parameter | Description |
|---|---|
| Idle Timeout | Specify the timeout period of idle connections. Unit: seconds. Valid values: 1 to 60. If no request is received within the specified timeout period, SLB closes the connection. SLB recreates the connection when a new connection request is received. |
| Request Timeout | Specify the request timeout period. Unit: seconds. Valid values: 1 to 180. If no response is received from the backend server within the request timeout period, SLB returns an HTTP 504 error to the client. |
| Enable Gzip Compression | Specify whether to enable Gzip compression to compress specific types of files. Gzip supports the following file types: text/xml, text/plain, text/css, application/javascript, application/x-javascript, application/rss+xml, application/atom+xml, and application/xml. |
| Add HTTP Header Fields | You can add the following HTTP headers: <ul style="list-style-type: none"> Use the X-Forwarded-For header to retrieve the real IP addresses of clients. Use the SLB-ID header to retrieve the ID of the SLB instance. Use the SLB-IP header to retrieve the public IP address of the SLB instance. Use the X-Forwarded-Proto header to retrieve the listener protocol used by the SLB instance. |
| Obtain Client Source IP Address | Specify whether to retrieve the real IP address of the client. By default, this feature is enabled. |
| Automatically Enable Listener After Creation | Specify whether to immediately enable the listener after it is created. By default, this feature is enabled. |

- Click **Next**.

Step 2: Add backend servers

After you configure the listener, you must add backend servers to process client requests. You can use the default server group that is configured for the SLB instance. You can also create a vServer group or a primary/secondary server group. For more information, see [Overview](#).

The default server group is selected in this example.

- On the **Backend Servers** wizard page, select **Default Server Group**. Then, click **Add More**.
- In the **My Servers** panel, select the Elastic Compute Service (ECS) instances that you want to add as backend servers and click **Next**.
- On the **Configure Ports and Weights** wizard page, specify the weights of the backend servers that you want to add. A backend server with a higher weight receives more requests.

 **Note** If the weight of a backend server is set to 0, no request is distributed to the backend server.

- Click **Add**. On the **Default Server Group** tab, specify the ports that you want to open on the backend servers to receive requests. The backend servers are the ECS instances that you selected. Valid values: 1 to 65535.

You can specify the same port on different backend servers that are added to an SLB instance.

5. Click **Next**.

Step 4: Configure health checks

SLB performs health checks to check the availability of the ECS instances that serve as backend servers. The health check feature improves overall service availability and reduces the impact of backend server failures. Click **Modify** to configure advanced health check settings and click **Next**. For more information, see [Health check overview](#).

Step 5: Submit the configurations

To submit the configurations, perform the following operations:

1. On the **Confirm** wizard page, check the configurations. You can click **Modify** to modify the configurations.
2. After you confirm the configurations, click **Submit**.
3. In the **Configuration Successful** message, click **OK**.

After you configure the listener, you can view the listener on the **Listener** tab.

16.5.5. Add an HTTPS listener

This topic describes how to add an HTTPS listener to a Server Load Balancer (SLB) instance. HTTPS is intended for applications that require encrypted data transmission. You can add an HTTPS listener to forward HTTPS requests.

Step 1: Configure an HTTPS listener

1. [Log on to the SLB console](#).
2. Use one of the following methods to open the listener configuration wizard:
 - On the **Instances** page, find the SLB instance that you want to manage and click **Configure Listener** in the **Actions** column.
 - On the **Instances** page, click the ID of the SLB instance that you want to manage. On the **Listener** tab, click **Add Listener**.
3. Set the following parameters and click **Next**.

| Parameter | Description |
|-----------------------------|--|
| Listener Protocol | Select the protocol type of the listener. In this example, HTTPS is selected. |
| Listening Port | Specify the listening port that is used to receive requests and forward them to backend servers. Valid values: 1 to 65535. |
| Advanced | Click Modify to configure advanced settings. |
| Scheduling Algorithm | Select a scheduling algorithm. <ul style="list-style-type: none">◦ Weighted Round-Robin (WRR): Backend servers that have higher weights receive more requests than backend servers that have lower weights.◦ Round-Robin (RR): Requests are distributed to backend servers in sequence. |

| Parameter | Description |
|------------------------------------|--|
| Enable Session Persistence | <p>Specify whether to enable session persistence.</p> <p>After session persistence is enabled, SLB forwards all requests from a client to the same backend server.</p> <p>SLB maintains the persistence of HTTP sessions based on cookies. SLB allows you to use the following methods to process cookies:</p> <ul style="list-style-type: none"> ◦ Insert cookie: If you select this option, you need only to specify the timeout period of the cookie. <p>SLB inserts a cookie (SERVERID) into the first HTTP or HTTPS response that is sent to a client. The next request from the client contains the cookie. Then, the listener distributes the request to the recorded backend server.</p> <ul style="list-style-type: none"> ◦ Rewrite cookie: If you select this option, you can specify the cookie that you want to insert into an HTTP or HTTPS response. You must specify the timeout period and the lifetime of a cookie on a backend server. <p>After you specify a cookie, SLB overwrites the original cookie with the specified cookie. The next time SLB receives a client request that carries the specified cookie, the listener distributes the request to the recorded backend server.</p> |
| Enable HTTP/2 | Select whether to enable HTTP/2.0 for the frontend protocol of the SLB instance. |
| Enable Peak Bandwidth Limit | <p>Specify whether to set a bandwidth limit for the listener.</p> <p>If an SLB instance is billed based on bandwidth usage, you can set different bandwidth limit values for different listeners. This limits the amount of traffic that flows through each listener. The sum of the bandwidth limit values of all listeners that are added to an SLB instance cannot exceed the bandwidth limit of this SLB instance. By default, this feature is disabled and all listeners share the bandwidth of the SLB instance.</p> |
| Idle Timeout | <p>Specify the timeout period of idle connections. Unit: seconds. Valid values: 1 to 60.</p> <p>If no request is received within the specified timeout period, SLB closes the connection. SLB recreates the connection when a new connection request is received.</p> |
| Request Timeout | <p>Specify the request timeout period. Unit: seconds. Valid values: 1 to 180.</p> <p>If no response is received from the backend server within the request timeout period, SLB returns an HTTP 504 error to the client.</p> |
| Enable Gzip Compression | <p>Specify whether to enable Gzip compression to compress specific types of files.</p> <p>Gzip supports the following file types: text/xml, text/plain, text/css, application/javascript, application/x-javascript, application/rss+xml, application/atom+xml, and application/xml.</p> |
| Add HTTP Header Fields | <p>You can add the following HTTP headers:</p> <ul style="list-style-type: none"> ◦ Use the X-Forwarded-For header to retrieve the real IP addresses of clients. ◦ Use the SLB-ID header to retrieve the ID of the SLB instance. ◦ Use the SLB-IP header to retrieve the public IP address of the SLB instance. ◦ Use the X-Forwarded-Proto header to retrieve the listener protocol used by the SLB instance. |

| Parameter | Description |
|---|--|
| Obtain Client Source IP Address | Specify whether to retrieve the real IP address of the client. By default, this feature is enabled. |
| Automatically Enable Listener After Creation | Specify whether to immediately enable the listener after it is created. By default, this feature is enabled. |

Step 2: Configure an SSL certificate

When you add an HTTPS listener, you must upload a server certificate or certification authority (CA) certificate, as shown in the following table.

| Item | Description | Required for one-way authentication | Required for mutual authentication |
|---------------------------|---|--|--|
| Server certificate | The certificate that is used to identify the server. Your browser uses the server certificate to verify whether the certificate sent by the server is signed and issued by a trusted CA. | Yes You must upload the server certificate to the certificate management system of SLB. | Yes You must upload the server certificate to the certificate management system of SLB. |
| Client certificate | The certificate that is used to identify the client. The server identifies the client by checking the certificate sent by the client. You can sign a client certificate with a self-signed CA certificate. | No | Yes You must install the client certificate on the client. |
| CA certificate | The server uses a CA certificate to verify the signature on the client certificate. If the signature is invalid, the connection request is denied. | No | Yes You must upload the CA certificate to the certificate management system of SLB. |

Before you upload a certificate, take note of the following items:

- SLB supports the following public key algorithms: RSA 1024, RSA 2048, RSA 4096, ECDSA P-256, ECDSA P-384, and ECDSA P-521.
- The certificate that you want to upload must be in the PEM format.
- After you upload a certificate to SLB, SLB can manage the certificate. You do not need to bind the certificate to backend servers.
- It may take a few minutes to upload, load, and verify the certificate. Therefore, an HTTPS listener is not enabled immediately after it is created. It requires about 1 to 3 minutes to enable an HTTPS listener.
- The ECDHE cipher suite used by HTTPS listeners supports forward secrecy. It does not support the security enhancement parameters that are required by the DHE cipher suite. Therefore, you cannot upload certificates (PEM files) that contain the `BEGIN DH PARAMETERS` field. For more information, see [Certificate requirements](#).
- HTTPS listeners do not support Server Name Indication (SNI). You can choose TCP listeners and configure SNI on backend servers.
- By default, the timeout period of session tickets for HTTPS listeners is 300 seconds.
- The actual amount of data transfer on an HTTPS listener is larger than the billed amount because a portion of data is used for handshaking.
- If a large number of connections are established, a large amount of data is used for handshaking.

1. On the **SSL Certificates** wizard page, select the server certificate that you uploaded. You can also click

Create Server Certificate to upload a server certificate.

2. To enable mutual authentication or configure a TLS security policy, click **Modify** next to **Advanced**.
3. Enable mutual authentication, and select an uploaded CA certificate. You can also upload a CA certificate.

Step 3: Add backend servers

After you configure the listener, you must add backend servers to process client requests. You can use the default server group that is configured for the SLB instance. You can also create a vServer group or a primary/secondary server group. For more information, see [Overview](#).

The default server group is selected in this example.

1. On the **Backend Servers** wizard page, select **Default Server Group**. Then, click **Add More**.
2. In the **My Servers** panel, select the Elastic Compute Service (ECS) instances that you want to add as backend servers and click **Next**.
3. On the **Configure Ports and Weights** wizard page, specify the weights of the backend servers that you want to add. A backend server with a higher weight receives more requests.

 **Note** If the weight of a backend server is set to 0, no request is distributed to the backend server.

4. Click **Add**. On the **Default Server Group** tab, specify the ports that you want to open on the backend servers to receive requests. The backend servers are the ECS instances that you selected. Valid values: 1 to 65535.

You can specify the same port on different backend servers that are added to an SLB instance.

5. Click **Next**.

Step 4: Configure health checks

SLB performs health checks to check the availability of the ECS instances that serve as backend servers. The health check feature improves overall service availability and reduces the impact of backend server failures. Click **Modify** to configure advanced health check settings and click **Next**. For more information, see [Health check overview](#).

Step 5: Submit the configurations

1. On the **Confirm** wizard page, check the configurations. You can click **Modify** to modify the configurations.
2. After you confirm the configurations, click **Submit**.
3. In the **Configuration Successful** message, click **OK**.

After you configure the listener, you can view the listener on the **Listener** tab.

16.5.6. Configure forwarding rules

This topic describes how to configure forwarding rules for a Server Load Balancer (SLB) instance. You can configure domain name-based or URL-based forwarding rules for an SLB instance that uses Layer 7 listeners. Layer 7 listeners distribute requests destined for different domain names or URLs to different Elastic Compute Service (ECS) instances.

Context

You can add multiple forwarding rules to a listener. Each forwarding rule is associated with a unique server group. Each server group contains one or more ECS instances. For example, you can configure a listener to forward read requests to one server group and write requests to another server group. This allows you to optimize load balancing among your server resources.

SLB forwards requests based on the following rules:

- If a request matches a domain name-based or URL-based forwarding rule of a listener, the request is forwarded to the corresponding server group based on the forwarding rule.
- If a request does not match the domain name-based or URL-based forwarding rules of a listener but the listener

is associated with a server group, the request is forwarded to the server group.

- If none of the preceding conditions are met, requests are forwarded to the ECS instances in the default server group of the SLB instance.

Procedure

1. Log on to the SLB console.
2. Click the ID of the SLB instance that you want to manage. On the Listener tab, find the listener that you want to manage.

You can configure domain name-based or URL-based forwarding rules only for HTTP and HTTPS listeners.

3. Click **Configure Forwarding Rule** in the **Actions** column.

4. Configure forwarding rules based on the following information:

- Configure a domain name-based forwarding rule

- When you configure a domain name-based forwarding rule, leave the URL field empty. You do not need to enter a forward slash (/) in this field. The domain name can contain only letters, digits, hyphens (-), and periods (.).
- Domain-based forwarding rules support both exact matching and wildcard matching. For example, www.aliyun.com is an exact domain name, whereas *.aliyun.com and *.market.aliyun.com are wildcard domain names. When a request matches multiple domain name-based forwarding rules, an exact match prevails over wildcard matches, as described in the following table. Domain name matching rule

| Type | Request URL | Domain name matching rule (√ indicates that the domain name is matched whereas x indicates that the domain name is not matched.) | | |
|----------------|------------------------|--|--------------|---------------------|
| | | www.aliyun.com | *.aliyun.com | *.market.aliyun.com |
| Exact match | www.aliyun.com | √ | x | x |
| Wildcard match | market.aliyun.com | x | x | x |
| | info.market.aliyun.com | x | x | √ |

- Configure a URL-based forwarding rule

- When you configure a URL-based forwarding rule, leave the Domain Name field empty.
- The URL can contain only letters, digits, and hyphens (-)./%?#&
- The URL must start with a forward slash (/).

Note If you enter only a forward slash (/) in the URL field, the URL-based forwarding rule is invalid.

- URL-based forwarding rules support string matching and adopt sequential matching. Examples: /admin , /bbs_ , and /ino_test .

- Configure both domain name-based and URL-based forwarding rules

You can configure both domain name-based and URL-based forwarding rules to forward traffic destined for different URLs of the same domain name. We recommend that you configure a default forwarding rule with the URL field left empty in case errors are returned when the URLs of requests are not matched.

For example, the domain name of a website is www.aaa.com . You are required to forward requests destined for www.aaa.com/index.html to Server Group 1 and forward requests destined for other URLs of the domain name to Server Group 2. To meet the preceding requirements, you must configure two forwarding rules, as shown in the following figure. Otherwise, a 404 error code is returned when a request destined for the www.aaa.com domain name does not match all forwarding rules.

5. Click **Save**.

16.5.7. Enable access control

This topic describes how to enable access control for a listener. You can enable access control for each listener of a Classic Load Balancer (CLB). You can set whitelists for different listeners to regulate network access control.

Procedure

1. [Log on to the SLB console](#).
2. Click the ID of the CLB instance.
3. Click the **Listener** tab, find the listener that you want to manage, and then choose  > **Set Access Control** in the **Actions** column.
4. Set the following parameters and click **OK**.

| Parameter | Description |
|------------------------------|--|
| Enable Access Control | Enable access control. |
| Access Control Method | <p>Whitelist: After you set a whitelist for a listener, the listener forwards only requests from IP addresses or CIDR blocks that are added to the whitelist.</p> <p>However, your business may be adversely affected if the whitelist is not set properly. After you set a whitelist for a CLB listener, only requests from IP addresses or CIDR blocks that are added to the whitelist are distributed by the listener. After you enable the whitelist, if no IP address is added to the whitelist, the listener does not forward requests.</p> |
| Access Control List | <p>Select a network ACL.</p> <p>IPv6 instances can be associated only with IPv6 network ACLs, and IPv4 instances can be associated only with IPv4 network ACLs.</p> <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e6f2ff;"> <p> Note Separate multiple IP entries with commas (,). You can add up to 300 IP entries to each network ACL. IP entries must be unique within each network ACL.</p> </div> |

16.5.8. Disable access control

This topic describes how to disable access control for a listener.

Procedure

1. [Log on to the SLB console](#).
2. Find an SLB instance and click its instance ID.
3. Click the **Listener** tab next to the **Instance Details** tab.
4. Find the listener for which you want to disable access control, and choose  > **Set Access Control** in the **Actions** column.
5. In the **Access Control Settings** panel, disable access control and then click **OK**.

16.6. Backend servers

16.6.1. Backend server overview

Before you use the Server Load Balancer (SLB) service, you must add Elastic Compute Service (ECS) instances as backend servers to an SLB instance to process client requests.

Backend server overview

You can set virtual IP addresses for an SLB instance. This way, the added ECS instances in the same region can be virtualized into an application service pool that provides high performance and availability. You can manage backend servers by using vServer groups. A listener of an SLB instance can be associated with a specific vServer group so that different listeners can forward requests to their associated backend servers that use different ports.

 **Note** If you associate a vServer group with a listener, the listener distributes requests to backend servers in the associated vServer group instead of those in the default server group.

Limits

You can increase or decrease the number of backend ECS instances at any time and switch ECS instances to receive client requests. When you perform the operations, make sure that the health check feature is enabled and at least one ECS instance is running as expected to maintain service stability.

When you add backend ECS instances, take note of the following items:

- You can add ECS instances of different operating systems to an SLB instance. However, the applications deployed on the ECS instances must be the same and have consistent data. We recommend that you use ECS instances of the same operating system to facilitate management and maintenance.
- Up to 50 listeners can be added to a single SLB instance. Each listener corresponds to an application deployed on backend ECS instances. Listening ports of an SLB instance correspond to application service ports that are opened on backend ECS instances.
- You can specify a weight for each ECS instance in the application service pool. An ECS instance with a higher weight receives more requests.
- If session persistence is enabled, requests may not be evenly distributed to backend ECS instances. To solve this problem, we recommend that you disable session persistence and check whether the problem persists.

If requests are not evenly distributed, troubleshoot the issue in the following way:

- i. Collect statistics on the access logs of the web service on backend ECS instances for a specified period.
 - ii. Check whether the numbers of access logs of backend ECS instances match SLB configurations. If session persistence is enabled, you must differentiate the access logs for the same IP address. If different weights are configured for backend ECS instances, you must check whether the percentage of access logs is normal based on the percentage of weights.
- When an ECS instance is undergoing hot migration, persistent connections to SLB may be interrupted. You can solve this problem by reestablishing the connections.

Default server groups

A default server group contains ECS instances that are used to receive requests. If a listener is not associated with a vServer group or a primary/secondary server group, the listener forwards requests to ECS instances in the default server group.

Before you use the SLB service, you must add at least one default backend server to receive client requests forwarded by SLB. For more information, see [Add a default backend server](#).

vServer groups

You can use a vServer group if you want to distribute different requests to different backend servers or configure forwarding rules based on domain names and URLs. For more information, see [Create a vServer group](#).

Primary/secondary server groups

A primary/secondary server group contains only two ECS instances. One ECS instance acts as the primary server and the other acts as the secondary server. Health checks are not performed on the secondary server. If the primary server is detected unhealthy, traffic is redirected to the secondary server. After the primary server recovers and is considered healthy, traffic is switched back to the primary server. For more information, see [Create a primary/secondary server group](#).

 **Note** You can add primary/secondary server groups only for TCP and UDP listeners.

Related information

- [Add a default backend server](#)
- [Create a vServer group](#)
- [Create a primary/secondary server group](#)

16.6.2. Default server groups

16.6.2.1. Add a default backend server

This topic describes how to add a default backend server. Before you use the Classic Load Balancer (CLB) service, you must add at least one default backend server to receive client requests.

Prerequisites

Before you add an Elastic Compute Service (ECS) instance to the default server group, make sure that the following requirements are met:

- A CLB instance is created. For more information, see [Create an SLB instance](#).
- ECS instances are created and applications are deployed on the ECS instances to receive requests.

Procedure

1. [Log on to the SLB console](#).
2. Find the CLB instance that you want to manage and click its ID.
3. Click the **Default Server Group** tab.
4. Click **Add**.
5. In the **My Servers** panel, for **Select Servers**, select one or more ECS instances that you want to add to the default server group.
6. Click **Next**.
7. For **Configure Ports and Weights**, specify the weight of each ECS instance.

An ECS instance that has a higher weight receives more requests.

You can change the weight of a server and then move the pointer over  to change the weights of other servers:

- If you click **Replicate to Below**, the weights of all servers below the current server are set to the weight of the current server.
- If you click **Replicate to Above**, the weights of all servers above the current server are set to the weight of the current server.
- If you click **Replicate to All**, the weights of all servers in the default server group are set to the weight of

the current server.

- o If you click **Reset**, when you clear the weight of the current server, the weights of all servers in the default server group are cleared.

Notice

- o Valid values of weights: 0 to 100. If you set the weight of a server to 0, the server does not receive requests.
- o If two servers have the same weight, only one server receives requests.

8. Click **Add**.

9. Click **OK**.

16.6.2.2. Add IDC servers to the default server group

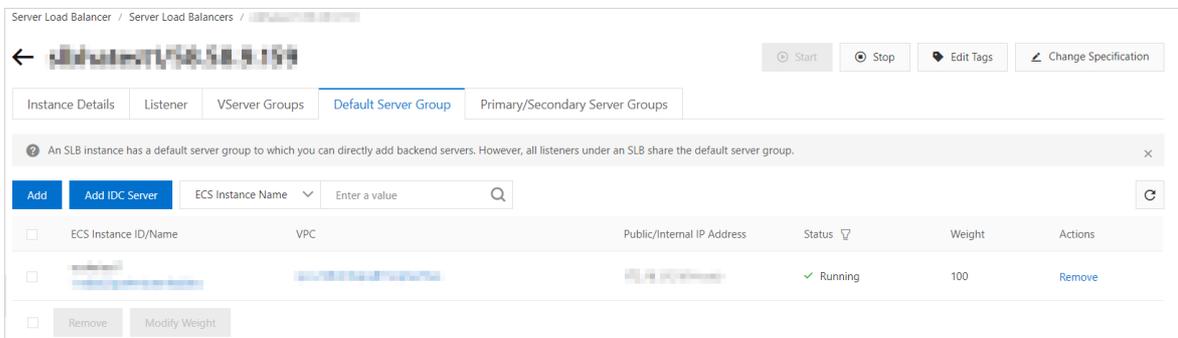
This topic describes how to add servers in on-premises Internet Data Centers (IDCs) as default backend servers to the default server group of an SLB instance. Before you use the SLB service, you must add at least one default backend server to receive client requests forwarded by SLB.

Prerequisites

Applications are deployed on the IDC servers, and the IDC servers are ready to receive distributed requests.

Procedure

1. [Log on to the SLB console](#).
2. Find the target SLB instance and click its instance ID.
3. Click the **Default Server Group** tab.
4. Click **Add IDC Server**.



5. In the **My Servers** dialog box, click **Add**.

6. Select a VPC from the VPC Connected to IDC drop-down list, enter a name for the IDC server, and specify the IP address of the IDC server.

7. Click **Next**.

8. In the **Configure Ports and Weights** step, specify the weight of each added IDC server.
An IDC server with a higher weight receives more requests.

You can change the weight of a server and then move the pointer over  to change the weights of multiple servers:

- Click **Replicate to Below**: The weights of all servers below the current server are set to the weight of the current server.
- Click **Replicate to Above**: The weights of all servers above the current server are set to the weight of the current server.
- Click **Replicate to All**: The weights of all servers in the default server group are set to the weight of the current server.
- Click **Reset**: The weight fields of all servers in the default server group are cleared.

 **Notice** If the weight of a backend server is set to 0, this backend server no longer receives new requests.

9. Click **Add**.

10. Click **OK**.

16.6.2.3. Change the weight of a backend server

This topic describes how to change the weight of a backend server to adjust the proportion of requests sent to the backend server.

Procedure

1. [Log on to the SLB console](#).
2. Find the target SLB instance and click its instance ID.
3. Click the **Default Server Group** tab.
4. Move the pointer over the weight value of the target backend server and click the  icon.
5. Change the weight and then click **OK**.

A backend server (ECS instance or IDC server) with a higher weight receives more requests.

 **Notice** The weight value ranges from 0 to 100. If the weight of a backend server is set to 0, no requests are sent to the backend server.

16.6.2.4. Remove a backend server

This topic describes how to remove a backend server that is no longer needed.

Procedure

1. [Log on to the SLB console](#).
2. Find the target SLB instance and click its instance ID.
3. Click the **Default Server Group** tab.
4. Find the target backend server and click **Remove** in the **Actions** column.
5. In the dialog box that appears, click **OK**.

16.6.3. VServer groups

16.6.3.1. Add ECS instances to a VServer group

This topic describes how to create a VServer group and then add ECS instances as backend servers to the VServer group. If you associate a VServer group with a listener, the listener distributes requests only to the backend servers in the VServer group instead of other backend servers.

Prerequisites

Before you create a VServer group, make sure that the following conditions are met:

- An SLB instance is created. For more information, see [Create an SLB instance](#).
- You have created ECS instances and deployed applications on these ECS instances to process requests.

Context

Note the following items before you create a VServer group:

- An ECS instance can be added to multiple VServer groups.
- A VServer group can be associated with multiple listeners of an SLB instance.
- The settings of the VServer group include the settings of ECS instances and application ports.

Procedure

1. [Log on to the SLB console](#).
2. Find the target SLB instance and click its instance ID.
3. Click the **VServer Groups** tab.
4. On the **VServer Groups** tab, click **Create VServer Group**.

5. On the **Create VServer Group** page, configure the VServer group.
 - i. In the **VServer Group Name** field, enter a name for the VServer group.
 - ii. Click **Add**. In the **My Servers** dialog box, select ECS instances that you want to add.
 - iii. Click **Next**.
 - iv. Specify a port and a weight for each ECS instance and then click **Add**.

Set the ports and weights based on the following information:

- **Port**: The backend port opened on an ECS instance to receive requests.

You can set the same port number for multiple backend servers of the same SLB instance. In addition, you can click **Add Port** to add multiple ports for a backend server.

- **Weight**: An ECS instance with a higher weight receives more requests.

 **Notice** If the weight of an ECS instance is set to 0, the ECS instance no longer receives new requests.

You can change the weight of a server and then move the pointer over  to change the weights of multiple servers:

- Click **Replicate to Below**: The weights of all servers below the current server are set to the weight of the current server.
- Click **Replicate to Above**: The weights of all servers above the current server are set to the weight of the current server.
- Click **Replicate to All**: The weights of all servers in the VServer group are set to the weight of the current server.
- Click **Reset**: The weight fields of all servers in the VServer group are cleared.

 **Notice** If the weight of a backend server is set to 0, this backend server no longer receives new requests.

- v. Click **Add**.

6. Click **Create**.

16.6.3.2. Add IDC servers to a VServer group

This topic describes how to create a VServer group and then add IDC servers to the VServer group. You can add ECS instances and IDC servers as backend servers to a VServer group. If you associate a VServer group with a listener, the listener distributes requests only to the backend servers in the VServer group instead of other backend servers.

Prerequisites

Before you create a VServer group, make sure that applications are deployed on the IDC servers and the IDC servers are ready to receive distributed requests.

Context

Note the following items before you create a VServer group:

- An IDC server can be added to multiple VServer groups.
- A VServer group can be associated with multiple listeners of an SLB instance.
- The settings of the VServer group include the settings of IDC servers and application ports.

Procedure

1. **Log on to the SLB console.**
2. Find the target SLB instance and click its instance ID.
3. Click the **VServer Groups** tab.
4. On the **VServer Groups** tab, click **Create VServer Group**.
5. On the **Create VServer Group** page, configure the VServer group.
 - i. In the **VServer Group Name** field, enter a name for the VServer group.
 - ii. Click **Add IDC Server**.
 - iii. In the **My Servers** dialog box, click **Add**.
 - iv. Select a VPC from the VPC Connected to IDC drop-down list, enter a name for the IDC server, and specify the IP address of the IDC server.

The IP address of the IDC server must be accessible to the VPC.
 - v. Click **Next**.
 - vi. Specify a port and weight for each IDC server, and then click **Add**.

Set the ports and weights based on the following information:

- **Port** : The backend port opened on an IDC server to receive requests. Multiple ports can be added to an IDC server.

You can set the same port number for multiple backend servers of the same SLB instance.

- **Weight** : An IDC server with a higher weight receives more requests.

 **Notice** If the weight of an IDC server is set to 0, the IDC server no longer receives new requests.

You can change the weight of a server and then move the pointer over  to change the weights of multiple servers:

- Click **Replicate to Below**: The weights of all servers below the current server are set to the weight of the current server.
- Click **Replicate to Above**: The weights of all servers above the current server are set to the weight of the current server.
- Click **Replicate to All**: The weights of all servers in the VServer group are set to the weight of the current server.
- Click **Reset**: The weight fields of all servers in the VServer group are cleared.

 **Notice** If the weight of a backend server is set to 0, this backend server no longer receives new requests.

- vii. Click **Add**.
6. Click **Create**.

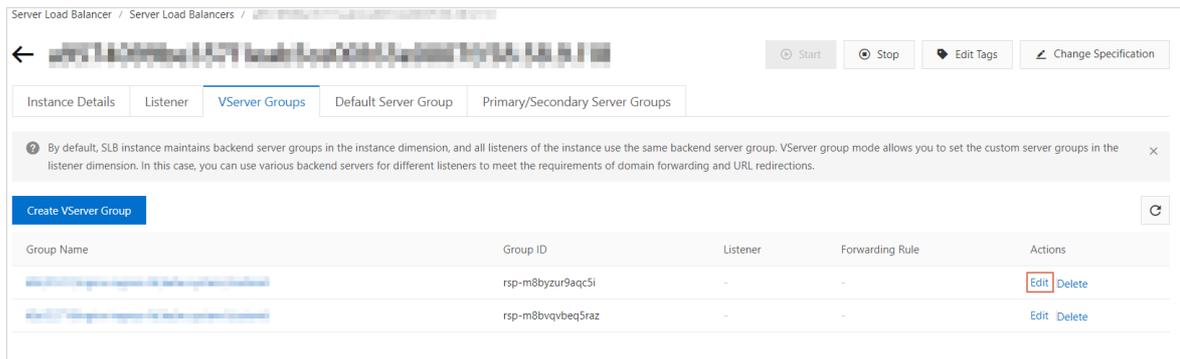
16.6.3.3. Modify a VServer group

This topic describes how to modify the settings of ECS instances or IDC servers in a VServer group.

Procedure

1. **Log on to the SLB console.**
2. Find the target SLB instance and click its instance ID.
3. Click the **VServer Groups** tab.

- Find the target VServer group and then click **Edit** in the Actions column.



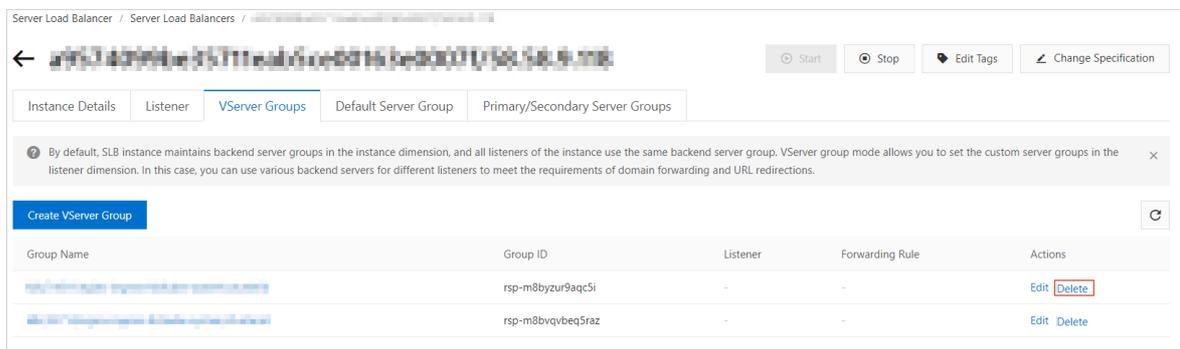
- Modify the ports and weights of ECS instances or IDC servers, and then click **Save**.

16.6.3.4. Delete a VServer group

This topic describes how to delete a VServer group that is no longer needed for traffic distribution.

Procedure

- Log on to the SLB console.
- Find the target SLB instance and click its instance ID.
- Click the **VServer Groups** tab.
- Find the target VServer group, and then click **Delete** in the Actions column.



- In the dialog box that appears, click **OK**.

16.6.4. Active/standby server groups

16.6.4.1. Add ECS instances to a primary/secondary server group

This topic describes how to create a primary/secondary server group and then add ECS instances to the primary/secondary server group. You can use a primary/secondary server group to implement failover between a primary server and a secondary server. By default, the primary server handles all distributed requests. When the primary server fails, traffic is redirected to the secondary server.

Prerequisites

Before you create a primary/secondary server group, make sure that the following conditions are met:

- An SLB instance is created. For more information, see [Create an SLB instance](#).
- You have created ECS instances and deployed applications on these ECS instances to process requests.

Procedure

1. Log on to the SLB console.
2. Find the target SLB instance and click its instance ID.
3. Click the **Primary/Secondary Server Groups** tab.
4. On the **Primary/Secondary Server Groups** tab, click **Create Primary/Secondary Server Group**.
5. On the **Create Primary/Secondary Server Group** page, configure the primary/secondary server group.
 - i. In the **Primary/Secondary Server Group Name** field, enter a name for the primary/secondary server group.

← Create Primary/Secondary Server Group

Note: The network type of the SLB instance is Classic Network, and the instance type is Private Network. You can add ECS instances in a classic or VPC network to the primary/secondary server group.

* Primary/Secondary Server Group Name

Enter a server group name

Added Servers

Add Add IDC Server Search by server name, ID, or IP

| ECS Instance ID/Name | Region | VPC | Public/Private IP | Status | Port | Reset | Type | Actions |
|----------------------|--------|-----|-------------------|--------|------|-------|------|---------|
| No data available. | | | | | | | | |

Create Cancel

- ii. Click **Add**. In the **My Servers** dialog box, select ECS instances that you want to add in the **Select Servers** step.

My Servers

1 Select Servers 2 Configure Ports and Weights

ECS Instance Name Search by name, ID, or IP VPC Select

Show Available Instances Only Advanced Mode

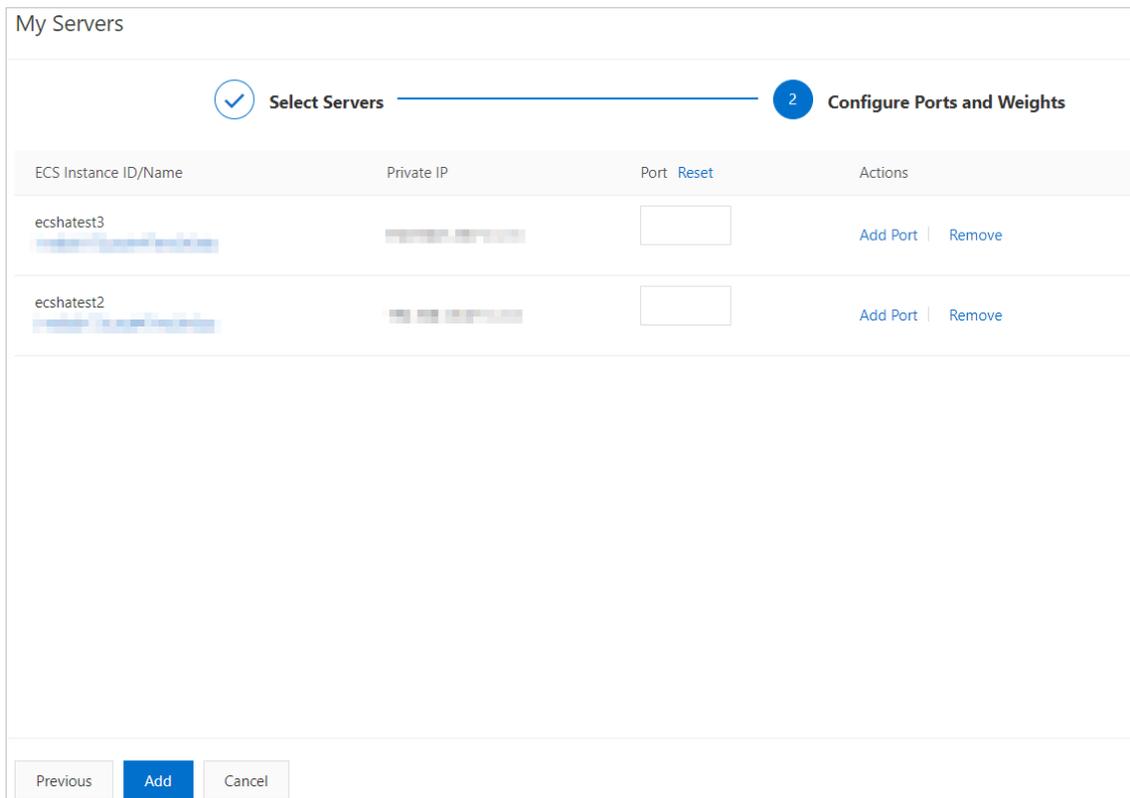
| <input type="checkbox"/> | ECS Instance ID/Name | Private IP | Public IP/VPC Property | Status | Number of SLB Associations |
|-------------------------------------|----------------------|-------------|------------------------|---------|----------------------------|
| <input type="checkbox"/> | ecshatest1 | 192.168.1.1 | 192.168.1.1 | Running | 1 |
| <input checked="" type="checkbox"/> | ecshatest2 | 192.168.1.2 | 192.168.1.2 | Running | 1 |
| <input checked="" type="checkbox"/> | ecshatest3 | 192.168.1.3 | 192.168.1.3 | Running | 1 |

You have selected 2 servers. Next Cancel

You can add up to two ECS instances to a primary/secondary server group.

- iii. Click **Next**.

iv. Configure the backend ports opened on ECS instances to receive requests, and then click **Add**.



You can set multiple ports for an ECS instance.

- v. Set an ECS instance as the primary server.
- vi. Click **Create**.

16.6.4.2. Add IDC servers to a primary/secondary server group

This topic describes how to create a primary/secondary server group and then add IDC servers to the primary/secondary server group. You can use a primary/secondary server group to implement failover between a primary server and a secondary server. By default, the primary server handles all distributed requests. When the primary server fails, traffic is redirected to the secondary server.

Prerequisites

The IDC servers are created, configured to deploy applications, and ready to receive distributed requests.

Procedure

1. [Log on to the SLB console](#).
2. Find the target SLB instance and click its instance ID.
3. Click the **Primary/Secondary Server Groups** tab.
4. On the **Primary/Secondary Server Groups** tab, click **Create Primary/Secondary Server Group**.
5. On the **Create Primary/Secondary Server Group** page, configure the primary/secondary server group.

- i. In the **Primary/Secondary Server Group Name** field, enter a name for the primary/secondary server group, and then click **Add IDC Server**.

← **Create Primary/Secondary Server Group**

Note: The network type of the SLB instance is Classic Network, and the instance type is Private Network. You can add ECS instances in a classic or VPC network to the primary/secondary server group.

* Primary/Secondary Server Group Name

Added Servers

Add **Add IDC Server**

| ECS Instance ID/Name | Region | VPC | Public/Private IP | Status | Port | Reset | Type | Actions |
|----------------------|--------|-----|-------------------|--------|------|-------|------|---------|
| No data available. | | | | | | | | |

Create **Cancel**

- ii. In the **My Servers** dialog box, click **Add**.

My Servers

1 **Select Servers** | 2 **Configure Ports and Weights**

| VPC Connected to IDC | IDC Server Name | IDC Server IP | Actions |
|---|--|--|---------------|
| <input type="text" value="vpc-xxxxxx"/> | <input type="text" value="DocuIDCServer"/> | <input type="text" value="192.168.1.1"/> | Remove |

+ Add

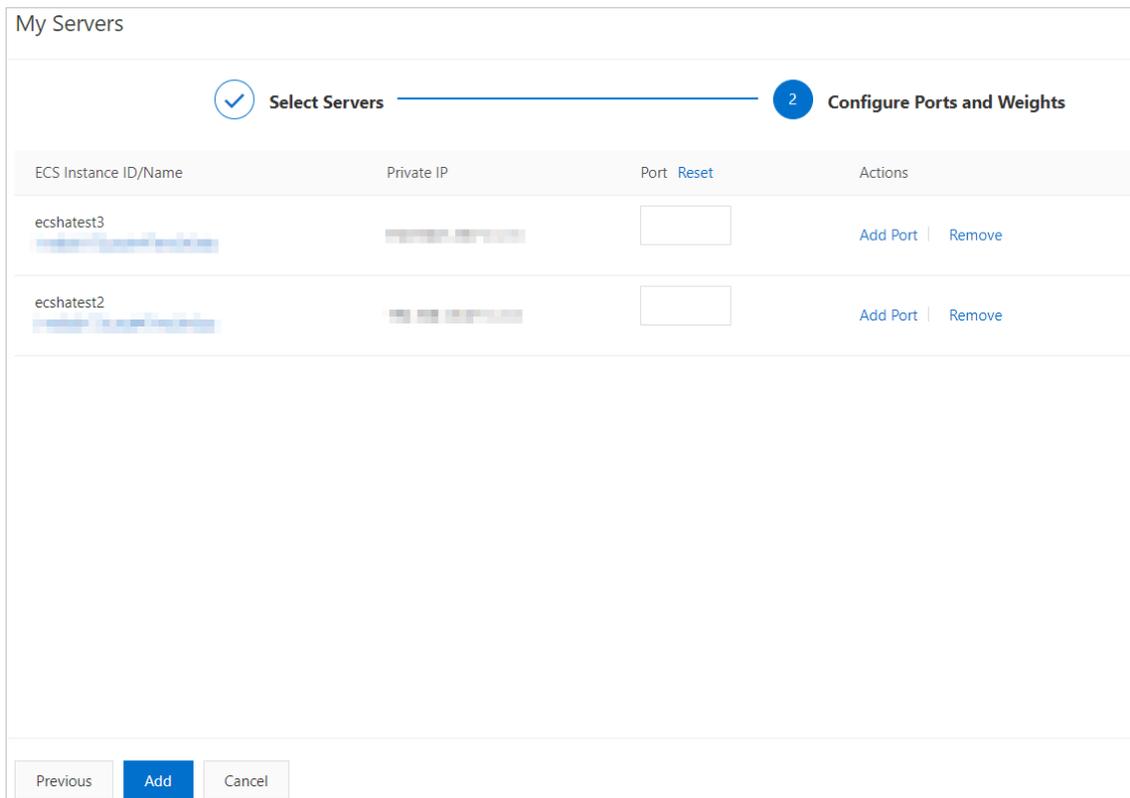
You have selected 1 servers. **Next** **Cancel**

- iii. Select a VPC from the **VPC Connected to IDC** drop-down list, enter a name for the IDC server, and specify the IP address of the IDC server.

The IP address of the IDC server must be accessible to the VPC.

- iv. Click **Next**.

- v. Configure the backend ports opened on ECS instances to receive requests, and then click **Add**.



You can set multiple ports for an IDC server.

- vi. Set a backend server as the primary server.
- vii. Click **Create**.

16.6.4.3. Delete a primary/secondary server group

This topic describes how to delete a primary/secondary server group of a Server Load Balancer (SLB) instance. If a primary/secondary server group is no longer needed to forward traffic, you can delete the primary/secondary server group.

Procedure

1. [Log on to the SLB console](#).
2. Find the SLB instance and click its instance ID.
3. Click the **Primary/Secondary Server Groups** tab.
4. On the **Primary/Secondary Server Groups** tab, find the primary/secondary server group that you want to delete and click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

16.7. Health check

16.7.1. Health check overview

This topic describes the health check feature of Server Load Balancer (SLB). SLB checks the availability of Elastic Compute Service (ECS) instances that act as backend servers by performing health checks. The health check feature improves the overall availability of your frontend business and mitigates the impacts of exceptions that occur on backend ECS instances.

After you enable the health check feature, SLB stops distributing requests to ECS instances that are declared unhealthy and distributes new requests to healthy ECS instances. When the unhealthy ECS instances have recovered, SLB starts forwarding requests to these ECS instances again.

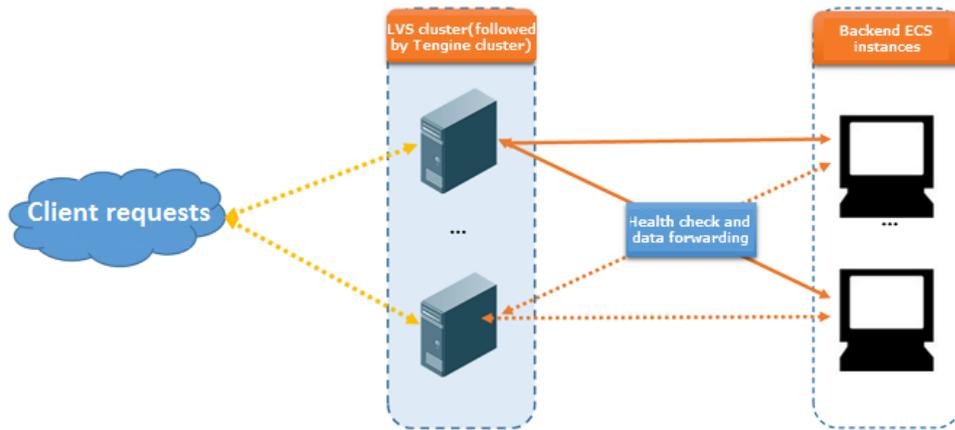
If your business is highly sensitive to traffic loads, frequent health checks may impact the availability of normal business. To reduce the impacts of health checks on your business, you can reduce the health check frequency, increase the health check interval, or change Layer 7 health checks to Layer 4 health checks. We recommend that you do not disable the health check feature to ensure business continuity.

Health check process

SLB is deployed in clusters. Node servers in the LVS or Tengine cluster forward data and perform health checks.

The node servers in the LVS cluster forward data and perform health checks independently and in parallel based on configured load balancing policies. If an LVS node server detects that a backend ECS instance is unhealthy, this node server no longer sends new client requests to this ECS instance. This operation is synchronized among all node servers in the LVS cluster.

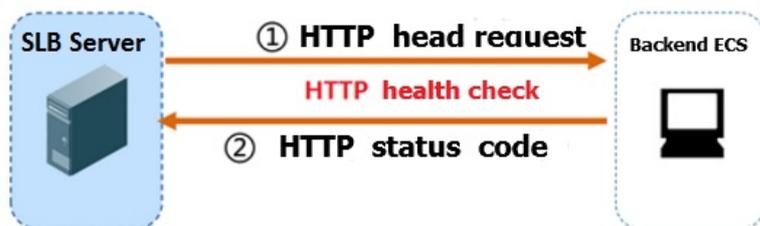
SLB uses the CIDR block of 100.64.0.0/10 for health checks. Make sure that backend ECS instances do not block this CIDR block. You do not need to configure a security group rule to allow access from this CIDR block. However, if you have configured security rules such as iptables, you must allow access from this CIDR block. 100.64.0.0/10 is reserved by Alibaba Cloud. Other users cannot use any IP addresses within this CIDR block, and therefore no relevant security risks exist.



Health checks of HTTP or HTTPS listeners

For Layer 7 (HTTP or HTTPS) listeners, SLB checks the status of backend ECS instances by sending HTTP HEAD requests. The following figure shows the process.

For HTTPS listeners, certificates are managed in SLB. To improve system performance, HTTPS is not used for data exchange (including health check data and business interaction data) between SLB and backend ECS instances.



The following section describes the health check process of a Layer 7 listener:

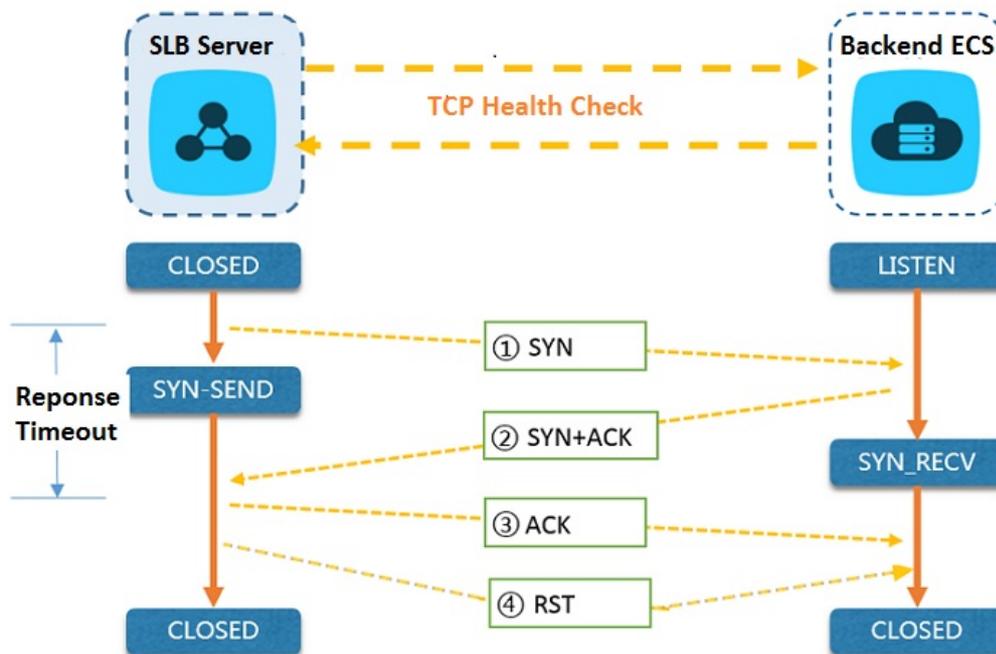
1. A Tengine node server sends an HTTP HEAD request that contains the configured domain name to the internal IP address, health check port, and health check path of a backend ECS instance based on health check settings.
2. After the backend ECS instance receives the request, the ECS instance returns an HTTP status code based on

the running status.

3. If the Tengine node server does not receive a response from the backend ECS instance within the specified response timeout period, the backend server is declared unhealthy.
4. If the Tengine node server receives a response from the backend ECS instance within the specified response timeout period, the node server compares the response with the configured status code. If the response contains the status code that indicates a healthy server, the backend server is declared healthy. Otherwise, the backend server is declared unhealthy.

Health checks of TCP listeners

For TCP listeners, SLB checks the status of backend servers by establishing TCP connections to improve health check efficiency. The following figure shows the process.



The following section describes the health check process of a TCP listener:

1. An LVS node server sends a TCP SYN packet to the internal IP address and health check port of a backend ECS instance.
2. After the backend ECS instance receives the request, the ECS instance returns an SYN-ACK packet if the corresponding port is listening normally.
3. If the LVS node server does not receive a packet from the backend ECS instance within the specified response timeout period, the backend ECS instance is declared unhealthy. Then, the node server sends an RST packet to the backend ECS instance to terminate the TCP connection.
4. If the LVS node server receives a packet from the backend ECS instance within the specified response timeout period, the node server determines that the service runs properly and the health check succeeds. Then, the node server sends an RST packet to the backend ECS instance to terminate the TCP connection.

Note A TCP three-way handshake is conducted to establish a TCP connection. After the LVS node server receives the SYN+ACK packet from the backend ECS instance, the node server sends an ACK packet, and then immediately sends an RST packet to terminate the TCP connection.

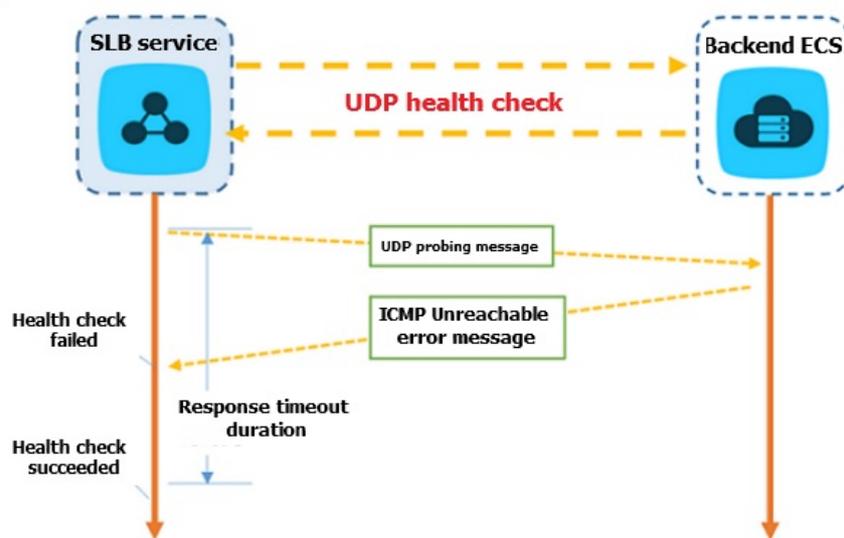
This process may cause backend ECS instances to think that an error such as an abnormal exit occurred in the TCP connection. Then, these instances may report a corresponding error message, such as `Connection reset by peer`, in logs such as Java connection pool logs.

Solution:

- You can implement HTTP health checks.
- If you have enabled the feature of obtaining actual client IP addresses on backend ECS instances, you can ignore connection errors caused by the access of the SLB CIDR block.

Health checks of UDP listeners

For UDP listeners, SLB checks the status of backend ECS instances by sending UDP packets. The following figure shows the process.



The following section describes the health check process of a UDP listener:

- An LVS node server sends a UDP packet to the internal IP address and health check port of an ECS instance based on health check configurations.
- If the corresponding port of the ECS instance is not listening normally, the system returns an ICMP error message, such as `port XX unreachable`. Otherwise, no message is returned.
- If the LVS node server receives the ICMP error message within the response timeout period, the backend ECS instance is declared unhealthy.
- If the LVS node server does not receive any messages from the backend ECS instance within the response timeout period, the ECS instance is declared healthy.

Note For UDP health checks, the health check result may not reflect the real status of a backend ECS instance in the following situation:

If the backend ECS instance uses a Linux operating system, the speed at which ICMP messages in high concurrency scenarios are sent is limited due to the ICMP attack prevention feature of Linux. In this case, even if a service exception occurs, SLB may declare the backend ECS instance healthy because the error message `port XX unreachable` is not returned. Consequently, the health check result deviates from the actual service status.

Solution:

You can specify a request and a response for UDP health checks. The ECS instance is considered healthy only when the specified response is returned. However, the client must be configured accordingly to return responses.

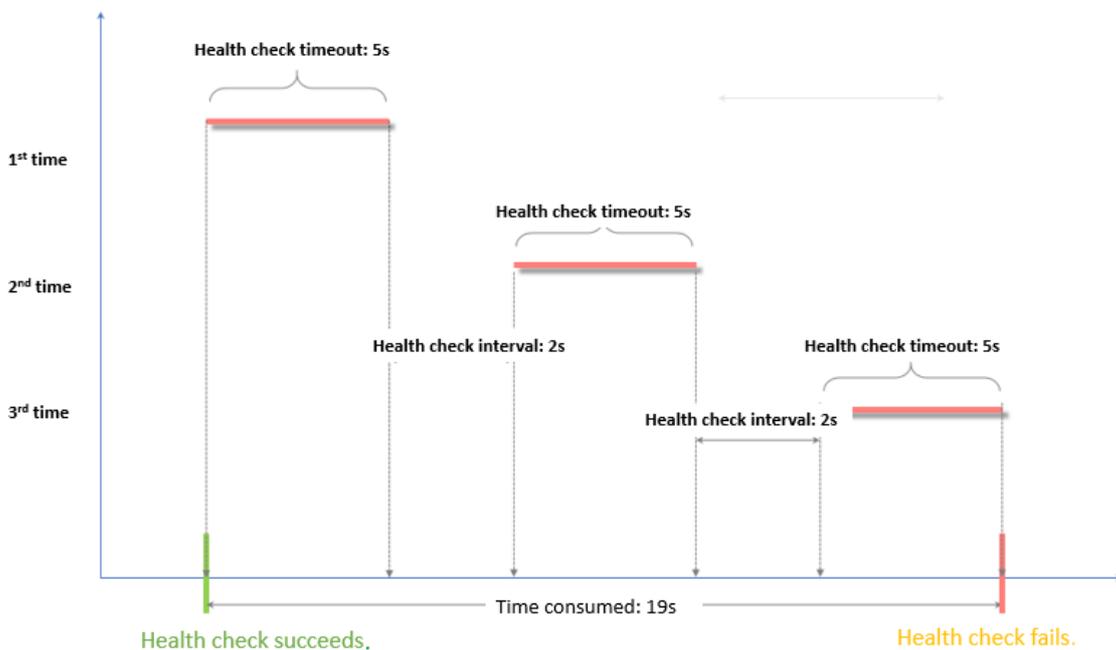
Health check time window

The health check feature effectively improves the availability of your services. However, to avoid impacts on system availability caused by frequent switching after failed health checks, the health check status switches only when health checks successively succeed or fail for a specified number of times within a certain time window. The health check time window is determined by the following factors:

- Health check interval: how often health checks are performed
- Response timeout: the length of time to wait for a response
- Health check threshold: the number of consecutive successes or failures of health checks

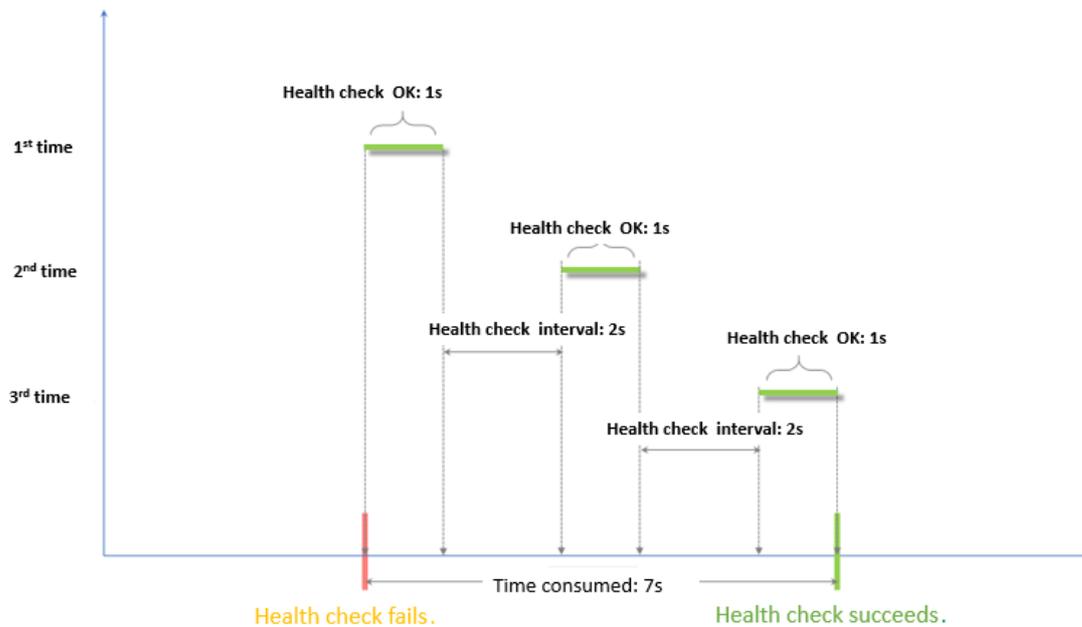
The health check time window is calculated based on the following formula:

- Time window for health check failures = Response timeout × Unhealthy threshold + Health check interval × (Unhealthy threshold - 1)



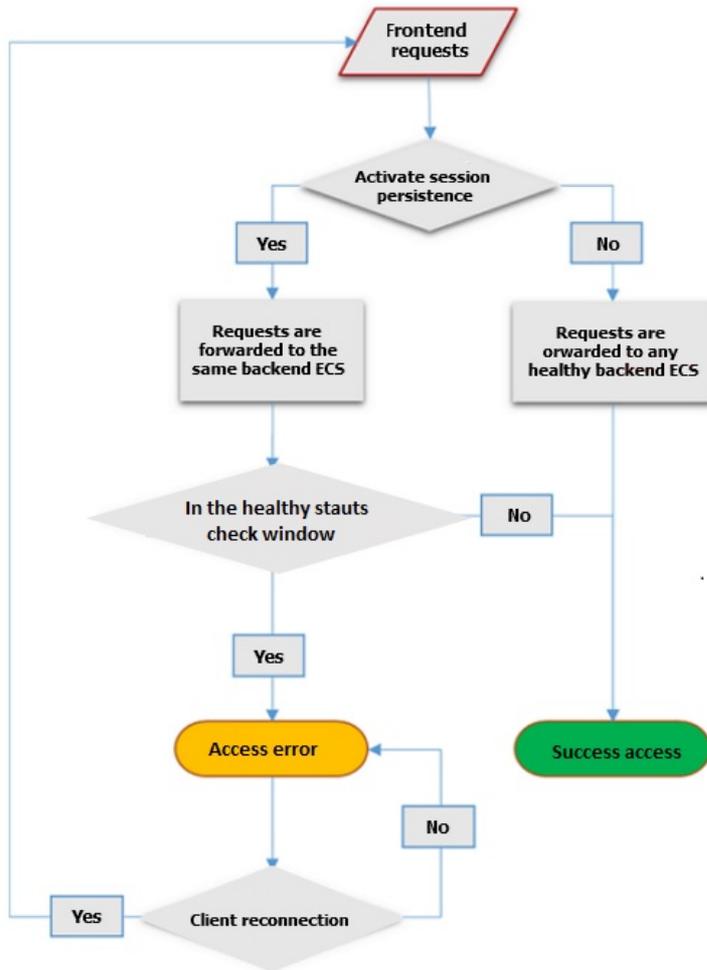
- Time window for health check successes = Response time of a successful health check × Healthy threshold + Health check interval × (Healthy threshold - 1)

Note The response time of a successful health check is the duration from the time when the health check request is sent to the time when the response is received. When TCP health checks are used, the response time is short and almost negligible because only whether the specific port is alive is checked. For HTTP health checks, the response time depends on the performance and load of the application server and is typically within a few seconds.



The health check result has the following impacts on request forwarding:

- If the health check of the backend ECS instance fails, new requests are distributed to other backend ECS instances. This does not affect client access.
- If the health check of the backend ECS instance succeeds, new requests are distributed to this instance. The client access is normal.
- If an exception occurs on the backend ECS instance and a request arrives during a time window for health check failures, the request is still sent to the backend ECS instance. This is because the number of failed health checks has not reached the unhealthy threshold (3 times by default). In this case, the client access fails.



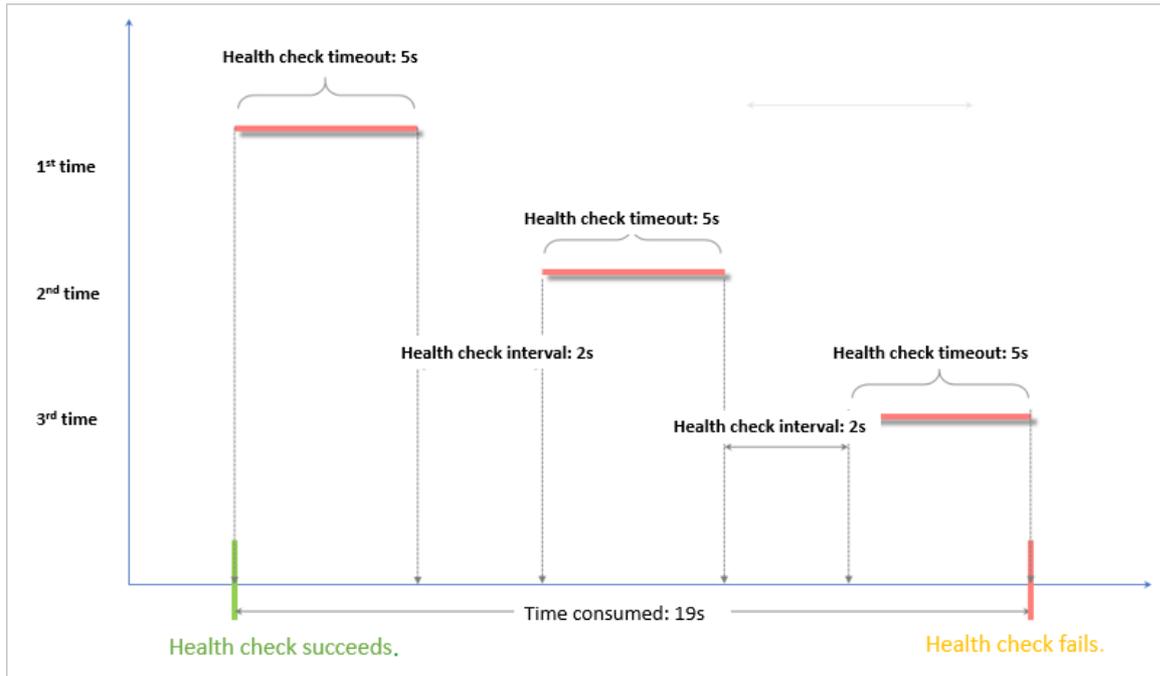
Examples of health check response timeout and health check interval

The following health check settings are used in these examples:

- Response Timeout Period: 5 Seconds
- Health Check Interval: 2 Seconds
- Healthy Threshold: 3 Times
- Unhealthy Threshold: 3 Times

Time window for health check failures = Response timeout × Unhealthy threshold + Health check interval × (Unhealthy threshold - 1). That is, $5 \times 3 + 2 \times (3 - 1) = 19$ seconds. If the response time of a health check exceeds 19 seconds, the health check fails.

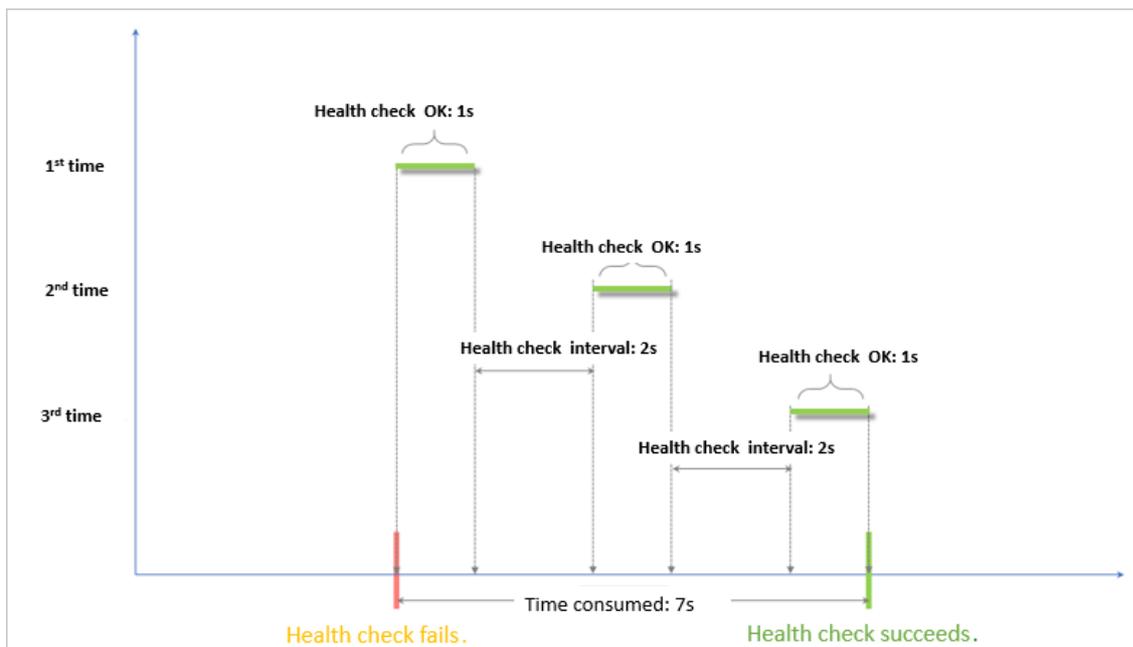
The following figure shows the time window from a healthy status to an unhealthy status.



Time window for health check successes = Response time of a successful health check × Healthy threshold + Health check interval × (Healthy threshold - 1). That is, $(1 \times 3) + 2 \times (3 - 1) = 7$ seconds. If the response time of a successful health check is less than seven seconds, the health check succeeds.

Note The response time of a successful health check is the duration from the time when the health check request is sent to the time when the response is received. When TCP health checks are used, the response time is short and almost negligible because only whether the specific port is alive is checked. For HTTP health checks, the response time depends on the performance and load of the application server and is typically within a few seconds.

The following figure shows the time window from an unhealthy status to a healthy status (assume that it takes 1 second for the server to respond to a health check request).



Domain name setting in HTTP health checks

When HTTP health checks are used, you can set a domain name for health checks. The setting is optional. Some application servers verify the host field in requests. In this case, the request header must contain the host field. If a domain name is configured in health check setting, SLB adds the domain name to the host field when SLB forwards a request to an application server. If no domain name is configured, the health check request is denied by the application server because it does not contain a host field and the health check may fail. If your application server verifies the host field in requests, you must configure a domain name to make sure that the health check feature works.

16.7.2. Configure health checks

This topic describes how to configure health checks. You can configure health checks when you create a listener or for an existing listener. The default health check settings can meet your requirements in most cases.

Procedure

1. [Log on to the SLB console](#).
2. Find an SLB instance and click the instance ID.
3. On the page that appears, click the **Listener** tab.
4. Click **Add Listener**, or find an existing listener and click **Modify Listener** in the **Actions** column.
5. Click **Next** to go to the **Health Check** step and configure the health check.

We recommend that you use the default settings when you configure health checks.

Health check parameters

| Parameter | Description |
|---|--|
| Health Check Protocol | <p>Select the protocol that the SLB instance uses when it performs health checks. For TCP listeners, both TCP health checks and HTTP health checks are supported.</p> <ul style="list-style-type: none"> ◦ A TCP health check implements detection at the network layer by sending SYN packets to check whether a port is open. ◦ An HTTP health check verifies the health of a backend server by sending HEAD or GET requests to simulate browser access. |
| Health Check Method (for the HTTP and HTTPS health checks only) | <p>Health checks of Layer 7 (HTTP or HTTPS) listeners support both the HEAD and GET methods. The HEAD method is used by default.</p> <p>If your backend application server does not support the HEAD method or if the HEAD method is disabled, the health check may fail. To solve this issue, you can use the GET method instead.</p> <p>If the GET method is used and the response size exceeds 8 KB, the response is truncated. However, the health check result is not affected.</p> |

| Parameter | Description |
|---|---|
| Health Check Path and Health Check Domain Name (Optional) (for the HTTP health checks only) | <p>By default, SLB sends HTTP HEAD requests to the default homepage configured on the application server through the internal IP address of the backend ECS instance to perform health checks.</p> <p>If you do not use the default homepage of the application server for health checks, you must specify the path for health checks.</p> <p>Some application servers verify the host field in requests. In this case, the request header must contain the host field. If a domain name is configured in health check settings, SLB adds this domain name to the host field when SLB forwards a health check request to one of the preceding application servers. If no domain name is configured, SLB does not include the host field in requests and the requests are rejected by the application server, which may cause health checks to fail. If your application server verifies the host field in requests, you must configure a domain name in health check settings to ensure that the health check feature functions properly.</p> |
| Normal Status Code (for the HTTP health checks only) | <p>Select the HTTP status code that indicates successful health checks.</p> <p>Default values: http_2xx and http_3xx.</p> |
| Health Check Port | <p>The detection port used by the health check feature to access backend servers.</p> <p>By default, the backend port configured for the listener is used.</p> <div style="background-color: #e6f2ff; padding: 5px; border: 1px solid #d9e1f2;"> <p> Note If a VServer group or a primary/secondary server group is configured for the listener, and the ECS instances in the group use different ports, leave this parameter empty. SLB uses the backend port of each ECS instance to perform health checks.</p> </div> |
| Response Timeout | <p>The length of time to wait for a health check response. If the backend ECS instance does not send an expected response within the specified period of time, the health check fails.</p> <p>Valid values: 1 to 300. Unit: seconds. Default value for UDP listeners: 10. Default value for HTTP, HTTPS, and TCP listeners: 5.</p> |
| Health Check Interval | <p>The interval between two consecutive health checks.</p> <p>All nodes in the LVS cluster perform health checks independently and in parallel on backend ECS instances at the specified interval. The health check statistics of a single ECS instance cannot reflect the health check interval because the nodes perform health checks at different times.</p> <p>Valid values: 1 to 50. Unit: seconds. Default value for UDP listeners: 5. Default value for HTTP, HTTPS, and TCP listeners: 2.</p> |

| Parameter | Description |
|---------------------|--|
| Unhealthy Threshold | The number of consecutive failed health checks that must occur on a backend ECS instance before this ECS instance is declared unhealthy. Valid values: 2 to 10. Default value: 3. |

6. Click **Next**.

16.7.3. Disable the health check feature

This topic describes how to disable the health check feature. If you disable the health check feature, requests may be distributed to unhealthy ECS instances and cause impacts on your business. We recommend that you enable the health check feature.

Context

 **Note** You can only disable the health check feature for HTTP and HTTPS listeners. The health check feature for UDP and TCP listeners cannot be disabled.

Procedure

1. [Log on to the SLB console](#).
2. On the **Server Load Balancer** page, find the target SLB instance and click its instance ID.
3. On the **Listeners** tab, find the target listener and click **Configure** in the **Actions** column.
4. On the **Configure Listener** page, click **Next** until the **Health Check** step appears.
5. Turn off **Enable Health Check**.
6. Click **Next**.
7. Click **Submit**, and then click **OK**.

16.8. Certificate management

16.8.1. Certificate overview

This topic provides an overview of the certificates that can be deployed on SLB instances. To use an HTTPS listener, you must upload the required third-party server certificate and digital identification issued by a certificate authority (CA) to SLB. You do not need to configure certificates on backend servers after uploading the certificates to SLB.

To upload a third-party certificate, you must have the files that contain the public key and private key of the certificate.

HTTPS server certificates and client CA certificates are supported.

You can create a maximum of 100 certificates per account.

16.8.2. Certificate requirements

Server Load Balancer (SLB) supports only certificates in the PEM format. Before you upload a certificate, make sure that the certificate content, certificate chain, and private key meet the corresponding format requirements.

Certificates issued by a root CA

If the certificate was issued by a root certification authority (CA), the received certificate is the only one that needs to be uploaded to SLB. In this case, the website that is configured with this certificate is regarded as a trusted website and does not require additional certificates.

The certificate must meet the following format requirements:

- The certificate must start with -----BEGIN CERTIFICATE----- and end with -----END CERTIFICATE-----.
- Each line (except the last line) must contain 64 characters. The last line can contain 64 or fewer characters.
- The certificate content cannot contain spaces.

Certificates issued by an intermediate CA

If the certificate was issued by an intermediate CA, the received certificate file contains multiple certificates. You must upload both the server certificate and the required intermediate certificates to SLB.

The format of the certificate chain must meet the following requirements:

- The server certificate must be put first and the content of the one or more required intermediate certificates must be put underneath without blank lines between the certificates.
- The certificate content cannot contain spaces.
- Blank lines are not allowed between the certificates. Each line must contain 64 characters. For more information, see [RFC1421](#).
- Certificates must meet the corresponding format requirements. In most cases, the intermediate CA provides instructions about the certificate format when certificates are issued. The certificates must meet the format requirements.

The following section provides a sample certificate chain:

```
-----BEGIN CERTIFICATE-----
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
-----END CERTIFICATE-----
```

Public keys of certificates

SLB supports the following public key algorithms:

- RSA 1024
- RSA 2048
- RSA 4096
- ECDSA P-256
- ECDSA P-384
- ECDSA P-521

RSA private keys

When you upload a server certificate, you must upload the private key of the certificate.

An RSA private key must meet the following format requirements:

- The private key must start with -----BEGIN RSA PRIVATE KEY----- and end with -----END RSA PRIVATE KEY-----, and these parts must also be uploaded.
- Blank lines are not allowed in the content. Each line (except the last line) must contain 64 characters. The last line can contain 64 or fewer characters. For more information, see [RFC1421](#).

You may use an encrypted private key. For example, the private key starts with `-----BEGIN PRIVATE KEY-----` and ends with `-----END PRIVATE KEY-----`, or starts with `-----BEGIN ENCRYPTED PRIVATE KEY-----` and ends with `-----END ENCRYPTED PRIVATE KEY-----`. The private key may also contain `Proc-Type: 4,ENCRYPTED`. In this case, you must first run the following command to convert the private key:

```
openssl rsa -in old_server_key.pem -out new_server_key.pem
```

16.8.3. Upload certificates

This topic describes how to create and upload certificates to Server Load Balancer (SLB). Before you create an HTTPS listener, you must upload the required server certificate and CA certificate to SLB. You do not need to configure certificates on backend servers after you upload the certificates to SLB.

Prerequisites

- A server certificate is purchased.
- A CA certificate and a client certificate are generated.

Context

Note that you can create up to 100 certificates with each account.

Procedure

1. In the left-side navigation pane, click **Certificates**.
2. On the Certificates page, click **Create Certificate**.
3. In the **Create Certificate** panel, set the required parameters and click **Create**.

| Parameter | Description |
|-------------------------------|---|
| Certificate Name | Enter a name for the certificate. The name must be 1 to 80 characters in length, and can contain only letters, digits, hyphens (-), forward slashes (/), periods (.), underscores (_), and asterisks (*). |
| Organization | The organization to which the certificate belongs. |
| Resource Group | The resource set to which the certificate belongs. |
| Certificate Standard | Select the type of certificate. You can select International Standard Certificate or National Standard Certificate . |
| Public Key Certificate | The content of the server certificate. Paste the content into the editor. Click Example to view the valid certificate formats. For more information, see Certificate requirements . |
| Private Key | The private key of the server certificate. Paste the private key into the editor. Click Example to view the valid certificate formats. For more information, see Certificate requirements . <div style="border: 1px solid #ccc; background-color: #e0f2f1; padding: 5px; margin-top: 10px;"> Notice A private key is required only when you upload a server certificate.</div> |
| Region | The region where you want to deploy the certificate. |

4. Click **Create**.

16.8.4. Generate a CA certificate

When you configure an HTTPS listener, you can use a self-signed CA certificate. This topic describes how to generate a CA certificate and use the CA certificate to sign a client certificate.

Generate a CA certificate by using Open SSL

1. Run the following commands to create a `ca` folder in the `/root` directory and then create four subfolders under the `ca` folder.

```
sudo mkdir ca
cd ca
sudo mkdir newcerts private conf server
```

- `newcerts` is used to store the digital certificate signed by the CA certificate.
 - `private` is used to store the private Key of the CA certificate.
 - `conf` is used to store the configuration files used for simplifying parameters.
 - `server` is used to store the server certificate.
2. Create an `openssl.conf` file that contains the following information in the `conf` directory.

```
[ ca ]
default_ca = foo
[ foo ]
dir = /root/ca
database = /root/ca/index.txt
new_certs_dir = /root/ca/newcerts
certificate = /root/ca/private/ca.crt
serial = /root/ca/serial
private_key = /root/ca/private/ca.key
RANDFILE = /root/ca/private/.rand
default_days = 365
default_cr_days = 30
default_md = md5
unique_subject = no
policy = policy_any
[ policy_any ]
countryName = match
stateOrProvinceName = match
organizationName = match
organizationalUnitName = match
localityName = optional
commonName = supplied
emailAddress = optional
```

3. Run the following command to generate a private Key.

```
cd /root/ca
sudo openssl genrsa -out private/ca.key
```

The following figure is an example of the key generation.

```
root@izbplhfivcqx1jwap31iz:~/ca/conf# cd /root/ca
root@izbplhfivcqx1jwap31iz:~/ca# sudo openssl genrsa -out private/ca.key
Generating RSA private key, 2048 bit long modulus
.....+++
.....+++
e is 65537 (0x10001)
```

4. Run the following command and input the required information according to the prompts. Press Enter to generate a *csr* file.

```
sudo openssl req -new -key private/ca.key -out private/ca.csr
```

```
root@iZbp1hfivcqx1jwv3p3liZ:~/ca# sudo openssl req -new -key private/ca.key -out private/ca.csr
You are about to be asked to enter information that will be incorporated into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]:CN
State or Province Name (full name) [Some-State]:ZheJiang
Locality Name (eg, city) []:HangZhou
Organization Name (eg, company) [Internet Widgits Pty Ltd]:Alibaba
Organizational Unit Name (eg, section) []:Test
Common Name (e.g. server FQDN or YOUR name) []:mydomain
Email Address []:a@alibaba.com

Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []:
An optional company name []:
root@iZbp1hfivcqx1jwv3p3liZ:~/ca#
```

 Note

Common Name is the domain name of the SLB instance.

5. Run the following command to generate a *crt* file:

```
sudo openssl x509 -req -days 365 -in private/ca.csr -signkey private/ca.key -out private/ca.crt
```

6. Run the following command to set the start sequence number for the private Key, which can be any four characters.

```
sudo echo FACE > serial
```

7. Run the following command to create a CA Key library:

```
sudo touch index.txt
```

8. Run the following command to create a certificate revocation list for removing the client certificate:

```
sudo openssl ca -gencrl -out /root/ca/private/ca.crl -crl days 7 -config "/root/ca/conf/openssl.conf"
```

The output is:

```
Using configuration from /root/ca/conf/openssl.conf
```

Sign the client certificate

1. Run the following command to generate a *users* folder under the *ca* directory to store the client Key.

```
sudo mkdir users
```

2. Run the following command to create a Key for the client certificate:

```
sudo openssl genrsa -des3 -out /root/ca/users/client.key 1024
```

Note

Enter a pass phrase when creating the Key. It is the password to protect the private Key from unauthorized access. Enter the same password twice.

3. Run the following command to create a *csr* file for the client Key.

```
sudo openssl req -new -key /root/ca/users/client.key -out /root/ca/users/client.csr
```

Enter the pass phrase set in the previous step and other required information when prompted.

Note

A challenge password is the password of the client certificate. Note that it is not the password of the client Key.

4. Run the following command to sign the client Key.

```
sudo openssl ca -in /root/ca/users/client.csr -cert /root/ca/private/ca.crt -keyfile /root/ca/private/ca.key -out /root/ca/users/client.crt -config "/root/ca/conf/openssl.conf"
```

Enter *y* twice when prompted to confirm the operation.

```
root@iZbp1hfivcqx1jwbp3liZ:~/ca# sudo openssl ca -in /root/ca/users/client.csr
-cert /root/ca/private/ca.crt -keyfile /root/ca/private/ca.key -out /root/ca/us
ers/client.crt -config "/root/ca/conf/openssl.conf"
Using configuration from /root/ca/conf/openssl.conf
Check that the request matches the signature
Signature ok
The Subject's Distinguished Name is as follows
countryName       :PRINTABLE:'CN'
stateOrProvinceName :ASN.1 12:'ZheJiang'
localityName      :ASN.1 12:'HangZhou'
organizationName  :ASN.1 12:'Alibaba'
organizationalUnitName:ASN.1 12:'Test'
commonName        :ASN.1 12:'mydomain'
emailAddress       :IA5STRING:'a@alibaba.com'
Certificate is to be certified until Jun  4 15:28:55 2018 GMT (365 days)
Sign the certificate? [y/n]:y

1 out of 1 certificate requests certified, commit? [y/n]:y
Write out database with 1 new entries
Data Base Updated
root@iZbp1hfivcqx1jwbp3liZ:~/ca#
```

5. Run the following command to convert the certificate to a *PKCS12* file.

```
sudo openssl pkcs12 -export -clcerts -in /root/ca/users/client.crt -inkey /root/ca/users/client.key -out /root/ca/users/client.p12
```

Follow the prompts to enter the pass phrase of client Key. Then enter the password used for exporting the client certificate. This password is used to protect the client certificate, which is required when you install the client certificate.

6. Run the following commands to view the generated client certificate:

```
cd users
ls
```

16.8.5. Convert the certificate format

Server Load Balancer (SLB) supports PEM certificates only. Certificates in other formats must be converted to the PEM format before they can be uploaded to SLB. We recommend that you use Open SSL for conversion.

Convert DER to PEM

DER: This format is usually used on a Java platform. The certificate file suffix is generally *.der*, *.cer*, or *.crt*.

- Run the following command to convert the certificate format:

```
openssl x509 -inform der -in certificate.cer -out certificate.pem
```

- Run the following command to convert the private key:

```
openssl rsa -inform DER -outform PEM -in privatekey.der -out privatekey.pem
```

Convert P7B to PEM

P7B: This format is usually used in a Windows server and Tomcat.

Run the following command to convert the certificate format:

```
openssl pkcs7 -print_certs -in incertificate.p7b -out outcertificate.cer
```

Convert PFX to PEM

PFX: This format is usually used in a Windows server.

- Run the following command to extract the certificate:

```
openssl pkcs12 -in certname.pfx -nokeys -out cert.pem
```

- Run the following command to extract the private key:

```
openssl pkcs12 -in certname.pfx -nocerts -out key.pem -nodes
```

16.8.6. Replace a certificate

This topic describes how to replace a certificate with a new certificate. We recommend that you replace certificates before they expire to avoid impacts on your service.

Procedure

1. Create and upload a new certificate.
For more information, see [Certificate overview](#).
2. Configure the certificate for the target HTTPS listener.
For more information, see [Add an HTTPS listener](#).
3. On the **Certificates** page, find the certificate to be replaced and click **Delete** in the Actions column.
4. In the dialog box that appears, click **OK**.

17.Virtual Private Cloud (VPC)

17.1. What is a VPC?

A virtual private cloud (VPC) is a private network dedicated for your use. You have full control over your VPC. For example, you can specify the CIDR block and configure route tables and gateways. In a VPC, you can deploy Apsara Stack resources, such as Elastic Compute Service (ECS) instances, ApsaraDB RDS instances, and Server Load Balancer (SLB) instances.

Furthermore, you can connect your VPC to other VPCs or on-premises networks to create a custom network environment. This way, you can migrate applications to the cloud and extend data centers.



Components

Each Virtual Private Cloud consists of at least one private network segment, one router, and at least one switch.

- Private CIDR blocks

When you create a VPC and a vSwitch, you must specify the private IP address range for the VPC in CIDR notation.

You can use the standard private CIDR blocks listed in the following table and their subsets as CIDR blocks for your VPCs. For more information, see [网络规划](#). For more information, see The network planning section in *User Guide*.

| CIDR blocks | Number of available private IP addresses (system reserved ones excluded) |
|----------------|--|
| 192.168.0.0/16 | 65,532 |
| 172.16.0.0/12 | 1,048,572 |
| 10.0.0.0/8 | 16,777,212 |

- VRouter

A VRouter is a hub that connects all vSwitches in a VPC and serves as a gateway between the VPC and other networks. After a VPC is created, the system creates a VRouter for the VPC. Each vRouter is associated with a route table.

For more information, see [Overview](#).

For more information, see the Route table overview topic in *User Guide*.

- vSwitches

A vSwitch is a basic network component that connects different cloud resources in a VPC. After you create a VPC, you can create a vSwitch to divide your VPC into multiple subnets. vSwitches deployed in a VPC can communicate with each other over the private network. You can deploy your applications in vSwitches that belong to different zones to improve service availability.

For more information, see [Switch](#).

For more information, see the [Create a vSwitch](#) topic in *User Guide*.

17.2. Log on to the VPC console

This topic describes how to log on to the Virtual Private Cloud (VPC) console of Apsara Uni-manager by using the Google Chrome browser.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

 **Note** When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top navigation bar, choose **Products > Networking > Virtual Private Cloud**.

17.3. Quick start

17.3.1. Plan and design a VPC

Before you create virtual private clouds (VPCs) and VSwitches, you need to plan the quantity and Classless Inter-domain Routing (CIDR) blocks of VPCs and VSwitches.

- [How many VPCs are required?](#)
- [How many VSwitches are required?](#)
- [How do I specify CIDR blocks?](#)
- [How do I specify CIDR blocks if I want to connect a VPC to other VPCs or on-premises data centers?](#)

How many VPCs are required?

- One VPC

We recommend that you create one VPC if you do not need to deploy systems in multiple regions or separate VPCs.

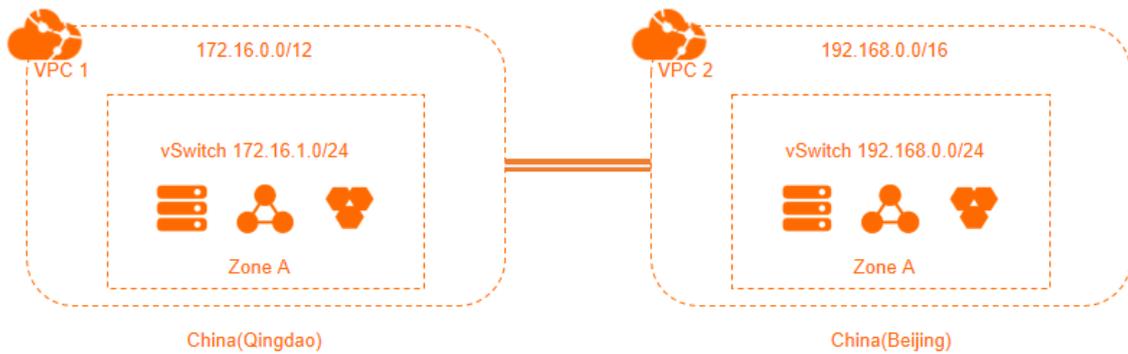


• Multiple VPCs

We recommend that you create multiple VPCs if you need to:

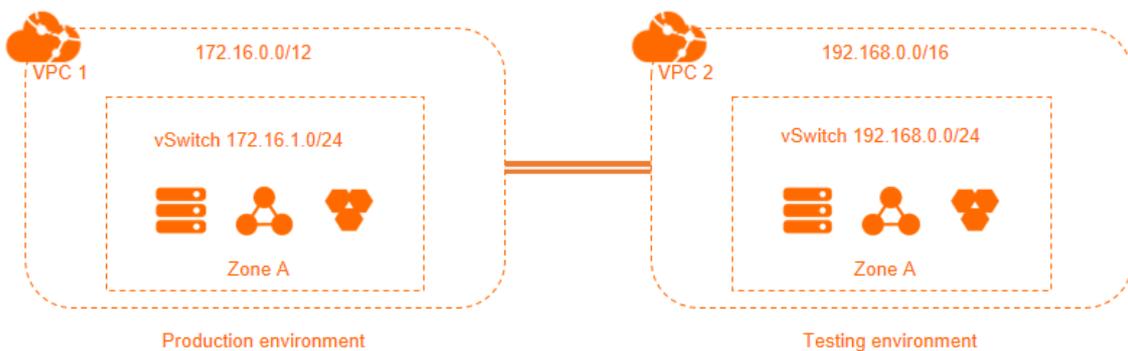
- Deploy application systems across regions.

A VPC cannot be deployed across regions. If you want to deploy your application systems in different regions, you must create multiple VPCs. You can use Express Connect and VPN Gateway to connect VPCs.



- Separate IT systems

To separate IT systems, you must create multiple VPCs. The following figure shows an example of isolating a production environment from a test environment by deploying them in separate VPCs.



How many VSwitches are required?

We recommend that you create at least two VSwitches for each VPC and deploy these VSwitches in different zones to achieve zone-disaster recovery.

After you deploy your applications in different zones within a region, you must measure the network latency between these applications. This is because the cross-zone network latency may be higher than expected due to complex data processing or cross-zone calls. An ideal approach is to optimize and adjust your systems to strike a balance between availability and latency.

In addition, the sizes and designs of your IT systems must also be taken into consideration when you create VSwitches. If you allow traffic from the Internet to be routed to and from the frontend systems, you can deploy the front-end systems in different VSwitches and the backend systems in other VSwitches to create a robust disaster recovery strategy.

How do I specify CIDR blocks?

When you create VPCs and VSwitches, you must specify their private IP address ranges in the form of CIDR blocks.

- VPC CIDR blocks

You can use 192.168.0.0/16, 172.16.0.0/12, 10.0.0.0/8, or their subsets as the CIDR blocks of your VPCs. To specify CIDR blocks for VPCs, follow these rules:

- If you have only one VPC and this VPC does not need to communicate with any on-premises data center, you can use one of the preceding CIDR blocks or one of their subsets as the CIDR block of the VPC.
- If you have multiple VPCs, or you need to build a hybrid cloud to integrate VPCs and on-premises data centers, we recommend that you use the subsets of the preceding CIDR blocks for your VPCs. In this case, the mask cannot be longer than 16 bits.

- VSwitch CIDR blocks

The CIDR block of a VSwitch must be a subset of the CIDR block of the VPC this VSwitch resides in. For example, if the CIDR block of a VPC is 192.168.0.0/16, the CIDR block of a VSwitch in the VPC must be a segment from 192.168.0.0/17 to 192.168.0.0/29.

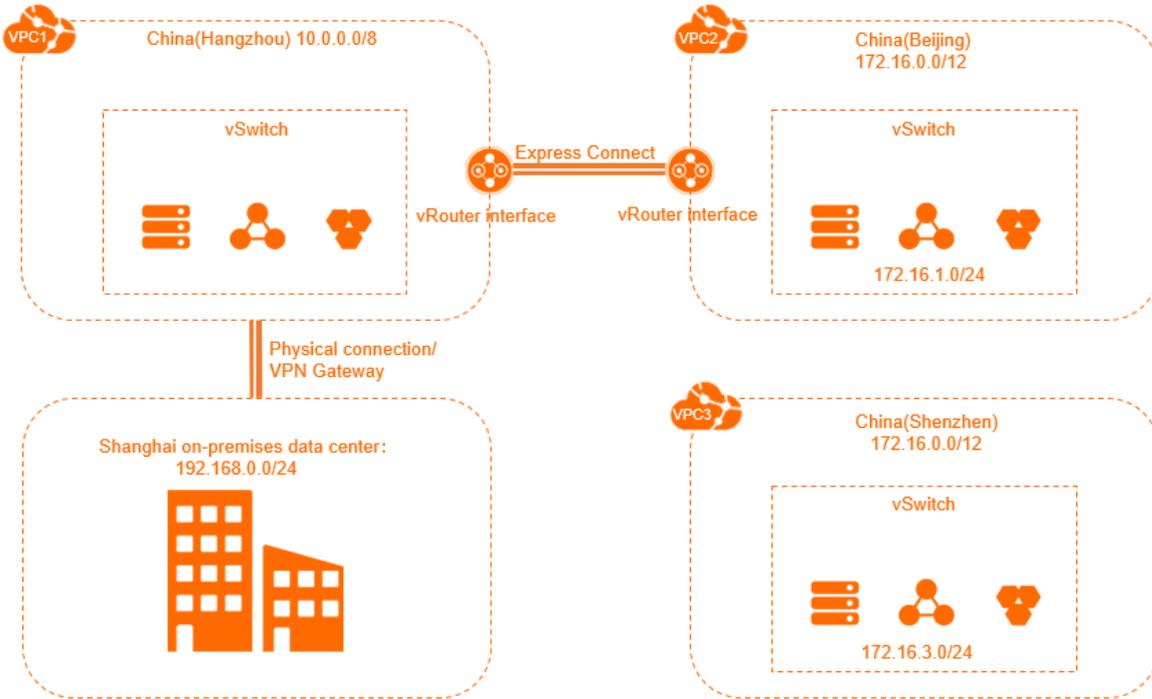
To specify CIDR blocks for VSwitches, follow these rules:

- The CIDR block size for a VSwitch is between a 16-bit mask and a 29-bit mask. It means that 8 to 65,536 IP addresses can be provided. This range is set because a 16-bit host address space provides addressing for 65,534 ECS instances, which can meet your needs in most cases, while a mask smaller than 29 bits can only allow very few usable host addresses.
- The first and the last three IP addresses in each VSwitch CIDR block are reserved by the system. For example, if the CIDR block of a VSwitch is 192.168.1.0/24, the IP addresses 192.168.1.0, 192.168.1.253, 192.168.1.254, and 192.168.1.255 are reserved.
- You must check the number of ECS instances in the VSwitch before you specify the CIDR block of a VSwitch.

How do I specify CIDR blocks if I want to connect a VPC to other VPCs or on-premises data centers?

Before you connect your VPC to another VPC or an on-premises data center, you must make sure that the CIDR block of your VPC does not conflict with that of the peer network.

For example, assume you have three VPCs: VPC1 in China (Hangzhou), VPC2 in China (Beijing), and VPC3 in China (Shenzhen), as shown in the following figure. Express Connect circuit is used for VPC1 and VPC2 to communicate with each other. VPC3 does not communicate with other VPCs, but may need to communicate with VPC2 in the future. Additionally, you have an on-premises data center in Shanghai, and you need to connect it to VPC1 by using an Express Connect circuit.



In this example, the CIDR block of VPC2 is different from the CIDR block of VPC1, but is the same with the CIDR block of VPC3. However, considering that VPC2 and VPC3 may need to communicate with each other later in the private network, the VSwitches in these VPCs are assigned with different CIDR blocks. This example demonstrates that VPCs communicating with each other can have identical CIDR blocks, but their VSwitches must have different CIDR blocks.

When you specify CIDR blocks for multiple VPCs that need to communicate with each other, follow these rules:

- The preferred practice is to specify different CIDR blocks for different VPCs. You can use the subsets of the standard CIDR blocks to increase the number of available CIDR blocks.
- If you cannot assign different CIDR blocks for VPCs, try to specify different CIDR blocks for the VSwitches in these VPCs.
- If you cannot assign different CIDR blocks for all VSwitches in these VPCs, make sure that different CIDR blocks are configured for the VSwitches communicating with each other.

17.3.2. Create an IPv4 VPC

This topic describes how to create a virtual private cloud (VPC) with an IPv4 CIDR block and create an Elastic Compute Service (ECS) instance in the VPC.

Prerequisites

To deploy cloud resources in a VPC, you must first prepare network subnetting. For more information, see [Plan and design a VPC](#).

Step 1: Create a VPC

Perform the following steps to create a VPC:

1. Log on to the VPC console.
2. On the VPCs page, click **Create VPC**.
3. On the **Create VPC** page, set the following parameters and click **Submit**.

| Parameter | Description |
|------------------------|---|
| Organization | Select the organization to which the VPC belongs. |
| Resource Set | Select the resource set to which the VPC belongs. |
| Region | Select the region where you want to deploy the VPC. |
| Sharing Scope | <p>Select the participants that can use the VPC to create resources.</p> <ul style="list-style-type: none"> ◦ Current Resource Set: Only the administrator of the current resource set can create resources for the shared VPC. ◦ Current Organization and Subordinate Organization: Only the administrators of the current organization and its subordinate organization can create resources for the shared VPC. ◦ Current Organization: Only the administrator of the current organization can create resources for the shared VPC. <p>Current Resource Set is selected in this example.</p> |
| VPC Name | <p>Enter a name for the VPC.</p> <p>The name must be 2 to 128 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code>.</p> <p>VPCtest is used in this example.</p> |
| IPv4 CIDR Block | <p>Select an IPv4 CIDR block for the VPC. The following settings are supported:</p> <ul style="list-style-type: none"> ◦ Recommended CIDR Block: You can use one of the following standard IPv4 CIDR blocks: 192.168.0.0/16 and 172.16.0.0/16. ◦ Custom CIDR Block: You can use 192.168.0.0/16, 172.16.0.0/16, or a subset of the preceding CIDR blocks. The subnet mask must be 8 to 28 bits in length. For example, enter 192.168.0.0/16. <p>In this example, 192.168.0.0/16 is used as the IPv4 CIDR block of the VPC.</p> <div style="background-color: #e0f2f7; padding: 5px; border: 1px solid #ccc;"> <p> Note After you create a VPC, you cannot change its IPv4 CIDR block.</p> </div> |
| IPv6 CIDR Block | <p>Specify whether to assign an IPv6 CIDR block.</p> <ul style="list-style-type: none"> ◦ Do Not Assign: The system does not assign an IPv6 CIDR block to the VPC. ◦ Assign: The system automatically assigns an IPv6 CIDR block to the VPC. <p>Do Not Assign is selected in this example.</p> |
| Description | <p>Enter a description for the VPC.</p> <p>The description must be 2 to 256 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code>.</p> |

4. Click **Back to Console**. On the VPCs page, you can view the VPCs that are created.

Step 2: Create a vSwitch

Perform the following steps to create a vSwitch in a VPC:

1. In the left-side navigation pane, click **vSwitches**.
2. On the **vSwitches** page, click **Create vSwitch**.
3. On the **vSwitch** page, set the following parameters and click **Submit**.

| Parameter | Description |
|---|---|
| Organization | Select the organization to which the vSwitch belongs. |
| Resource Set | Select the resource set to which the vSwitch belongs. |
| Region | Select the region where the vSwitch is deployed. |
| Zone | <p>Select the zone to which the vSwitch belongs.</p> <p>In a VPC, each vSwitch can be deployed in only one zone. You cannot deploy a vSwitch across zones. However, you can deploy cloud resources in vSwitches that belong to different zones to achieve cross-zone disaster recovery.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note An instance can be deployed in only one vSwitch.</p> </div> |
| Sharing Scope | <p>Select the participants that can use the vSwitch to create resources.</p> <ul style="list-style-type: none"> ◦ Current Resource Set: Only the administrator of the current resource set can create resources in the vSwitch. ◦ Current Organization and Subordinate Organization: Only the administrators of the current organization and its subordinate organizations can create resources in the vSwitch. ◦ Current Organization: Only the administrator of the current organization can create resources in the vSwitch. <p>Current Resource Set is selected in this example.</p> |
| VPC | <p>The VPC in which you want to create the vSwitch.</p> <p>VPCtest is selected in this example.</p> |
| Dedicated for Out-of-cloud Physical Machines | <p>Specify whether the vSwitch is dedicated for bare-metal servers.</p> <p>For more information about bare-metal servers, see the Bare-metal servers in VPCs topic in <i>Bare-metal Server Management Service User Guide</i>.</p> <p>No is selected in this example.</p> |
| vSwitch Name | <p>Enter a name for the vSwitch.</p> <p>The name must be 2 to 128 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code>.</p> |
| IPv4 CIDR Block | <p>Specify an IPv4 CIDR block for the vSwitch.</p> <p>The default IPv4 CIDR block is used in this example.</p> |

| Parameter | Description |
|-----------------|---|
| IPv6 CIDR Block | Specify an IPv6 CIDR block for the vSwitch. Do Not Assign is selected in this example. |
| Description | Enter a description for the vSwitch. The description must be 2 to 256 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . |

Step 3: Create a security group

Perform the following steps to create a security group:

1. In the top navigation bar, choose **Products > Elastic Computing > Elastic Compute Service**.
2. Choose **Networks and Security > Security Groups**.
3. On the **Security Groups** page, click **Create Security Group**.
4. On the **Create Security Group** page, set the following parameters and click **Submit**.

| Parameter | Description |
|---------------------|---|
| Organization | Select the organization to which the security group belongs. |
| Resource Set | Select the resource set to which the security group belongs. |
| Region | Select the region to which the security group belongs. Make sure that the security group and the VPC belong to the same region. |
| Zone | Select the zone to which the security group belongs. |
| Sharing Scope | Select the participants that can use the security group to create resources. <ul style="list-style-type: none"> ◦ Current Resource Set: Only the administrator of the current resource set can create resources for the security group. ◦ Current Organization and Subordinate Organization: Only the administrators of the current organization and its subordinate organization can create resources for the security group. ◦ Current Organization: Only the administrator of the current organization can create resources for the security group. Current Resource Set is selected in this example. |
| VPC | The VPC to which the security group belongs. |
| Security Group Name | Enter a name for the security group. The name must be 2 to 128 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . |

| Parameter | Description |
|-------------|--|
| Description | Enter a description for the security group. The description must be 2 to 256 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . |

Step 4: Create an ECS instance

Perform the following steps to create an ECS instance in the VPC:

1. In the top navigation bar, choose **Products > Networking > Virtual Private Cloud**.
2. In the left-side navigation pane, click **vSwitches**.
3. In the top navigation bar, select the region where the vSwitch is deployed.
4. On the **vSwitches** page, find the vSwitch that you want to manage, and choose **Create > ECS Instance** in the **Actions** column.
5. On the **Create ECS Instance** page, set the parameters and click **Submit**.

For more information about how to configure ECS instances, see [Create an instance](#) in *Quick start of ECS user guide*.

17.3.3. Create an IPv6 VPC

This topic describes how to create a virtual private cloud (VPC) that supports IPv6 CIDR blocks and then create an Elastic Compute Service (ECS) instance that is assigned an IPv6 address in the VPC to access IPv6 services.

Step 1: Create a VPC and a vSwitch

Before you deploy cloud resources in a VPC, you must create a VPC and a vSwitch.

Perform the following steps to create a VPC and a vSwitch:

1. [Log on to the VPC console](#) Log on to the VPC console.
2. On the **VPCs** page, click **Create VPC**.
3. On the **Create VPC** page, set the following parameters to configure the VPC and click **OK**.

| Parameter | Description |
|--------------|---|
| Organization | Select the organization to which the VPC belongs. |
| Resource Set | Select the resource set to which the VPC belongs. |
| Region | Select the region where you want to deploy the VPC. |

| Parameter | Description |
|------------------------|--|
| Sharing Scope | <p>Specify the scope of entities that are allowed to use the VPC.</p> <ul style="list-style-type: none"> ◦ Current Resource Set: If you select this option, the administrator of the current resource set can create resources in the VPC. ◦ Current Organization and Subordinate Organizations: If you select this option, administrators that belong to the current organization and subordinate organizations can create resources in the VPC. ◦ Current Organization: If you select this option, administrators that belong to the current organization can create resources in the VPC. <p>In this example, Current Resource Set is selected.</p> |
| VPC Name | <p>Enter a name for the VPC.</p> <p>The name must be 2 to 128 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter or Chinese character and cannot start with <code>http://</code> or <code>https://</code>.</p> <p>In this example, VPCTest is entered.</p> |
| IPv4 CIDR Block | <p>Specify the IPv4 CIDR block of the VPC. You can specify an IPv4 CIDR block in one of the following ways:</p> <ul style="list-style-type: none"> ◦ Recommended CIDR Block: You can use one of the following standard IPv4 CIDR blocks: 192.168.0.0/16, 172.16.0.0/12, and 10.0.0.0/8. ◦ Custom CIDR Block: Enter 192.168.0.0/16, 172.16.0.0/12, 10.0.0.0/8, or a subset of these CIDR blocks. The subnet mask must be 8 to 28 bits in length. For example, you can enter 192.168.0.0/16. <p>In this example, Recommended CIDR Block is selected and 192.168.0.0/16 is selected as the IPv4 CIDR block of the VPC.</p> <div style="background-color: #e6f2ff; padding: 5px; border: 1px solid #d9e1f2;"> <p> Note After you create a VPC, you cannot change its IPv4 CIDR block.</p> </div> |
| IPv6 CIDR Block | <p>Specify whether to assign an IPv6 CIDR block.</p> <ul style="list-style-type: none"> ◦ Do Not Assign: If you select this option, the system does not assign an IPv6 CIDR block to the VPC. ◦ Assign: If you select this option, the system automatically assigns an IPv6 CIDR block to the VPC. <p>In this example, Assign is selected.</p> |
| Description | <p>Enter a description for the VPC.</p> <p>The description must be 2 to 256 characters in length and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter or Chinese character. It cannot start with <code>http://</code> or <code>https://</code>.</p> |

4. Click **Back to Console**. In the left-side navigation pane, click **VSwitches**.
5. On the **VSwitches** page, click **Create VSwitch**.
6. On the **vSwitch** page, set the following parameter to configure the vSwitch and click **Submit**.

| Parameter | Description |
|---------------------|---|
| Organization | Select the organization to which the vSwitch belongs. |
| Resource Set | Select the resource set to which the vSwitch belongs. |

| Parameter | Description |
|---|---|
| Region | Select the region where you want to deploy the vSwitch. |
| Zone | <p>Select the zone where you want to deploy the vSwitch.</p> <p>In a VPC, each vSwitch can be deployed in only one zone. You cannot deploy a vSwitch across zones. However, you can deploy cloud resources in vSwitches that belong to different zones to achieve zone-disaster recovery.</p> <div style="background-color: #e6f2ff; padding: 5px; border: 1px solid #d9e1f2;"> <p> Note Each cloud resource can be added to only one vSwitch.</p> </div> |
| Sharing Scope | <p>Specify the scope of entities that are allowed to use the vSwitch.</p> <ul style="list-style-type: none"> ◦ Current Resource Set: If you select this option, the administrator of the current resource set can create resources in the vSwitch. ◦ Current Organization and Subordinate Organizations: If you select this option, administrators that belong to the current organization and subordinate organizations can create resources in the vSwitch. ◦ Current Organization: If you select this option, administrators that belong to the current organization can create resources in the vSwitch. <p>In this example, Current Resource Set is selected.</p> |
| VPC | <p>Select the VPC where you want to deploy the vSwitch.</p> <p>In this example, VPCtest is selected.</p> |
| Dedicated for Out-of-cloud Physical Machines | <p>Specify whether the vSwitch to be created is dedicated to bare-metal servers.</p> <p>For more information about bare-metal servers, see the Bare-metal servers in VPCs topic in <i>Bare-metal Server Management Service User Guide</i>.</p> <p>In this example, No is selected.</p> |
| vSwitch Name | <p>Enter a name for the vSwitch.</p> <p>The name must be 2 to 128 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter or Chinese character. It cannot start with <code>http://</code> or <code>https://</code>.</p> |
| IPv4 CIDR Block | <p>Enter an IPv4 CIDR block for the vSwitch.</p> <p>In this example, the default IPv4 CIDR block is used.</p> |
| IPv6 CIDR Block | <p>Enter an IPv6 CIDR block for the vSwitch.</p> <p>In this example, the default IPv6 CIDR block is used.</p> |
| Description | <p>Enter a description for the vSwitch.</p> <p>The description must be 2 to 256 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter or Chinese character. It cannot start with <code>http://</code> or <code>https://</code>.</p> |

Step 2: Create a security group

Perform the following steps to create a security group:

1. In the top navigation bar, choose **Products > Elastic Computing > Elastic Compute Service**.

2. Choose **Networks and Security > Security Groups**.
3. On the **Security Groups** page, click **Create Security Group**.
4. On the **Create Security Group** page, set the following parameters to configure the security group and click **Submit**.

| Parameter | Description |
|----------------------------|---|
| Organization | Select the organization to which the security group belongs. |
| Resource Set | Select the resource set to which the security group belongs. |
| Region | Select the region to which the security group belongs. Make sure that the security group and the VPC belong to the same region. |
| Zone | Select the zone to which the security group belongs. |
| VPC | Select the VPC to which the security group belongs. |
| Security Group Name | Enter a name for the security group. The name must be 2 to 128 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter or Chinese character. It cannot start with <code>http://</code> or <code>https://</code> . |
| Description | Enter a description for the security group. The description must be 2 to 256 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter or Chinese character. It cannot start with <code>http://</code> or <code>https://</code> . |

Step 3: Create and configure an ECS instance

After you create a VPC and a vSwitch, you must create an ECS instance and assign an IPv6 address to the ECS instance. You must associate this IPv6 address with the network interface controller (NIC) of the ECS instance.

Perform the following steps to create and configure an ECS instance:

1. In the top navigation bar, choose **Products > Networking > Virtual Private Cloud**.
2. In the left-side navigation pane, click **vSwitches**.
3. Select the region where the vSwitch is created.
4. On the **vSwitches** page, find the vSwitch that you want to manage and choose **Create > ECS Instance** in the **Actions** column.
5. On the **Create ECS Instance** page, configure the ECS instance and click **Submit**.

In this example, **Assign** is selected. Therefore, an IPv6 IP address is assigned to the ECS instance. For more information about other parameters that you are required to specify when you create an ECS instance, see [Create an instance](#) **Create an ECS instance in Quick Start** of *Elastic Compute Service User Guide*.

6. Return to the **Instances** page and click the instance ID to view the IPv6 address that is assigned to the ECS instance.
7. Configure a static IPv6 address.
 - o If the image of your ECS instance supports DHCPv6, you do not need to manually configure a static IPv6 address. DHCPv6 enables automatic configuration of IPv6 addresses. Therefore, if your ECS instance image supports DHCPv6, the ECS instance can use the assigned IPv6 address to communicate within the private network.

The following images support DHCPv6:

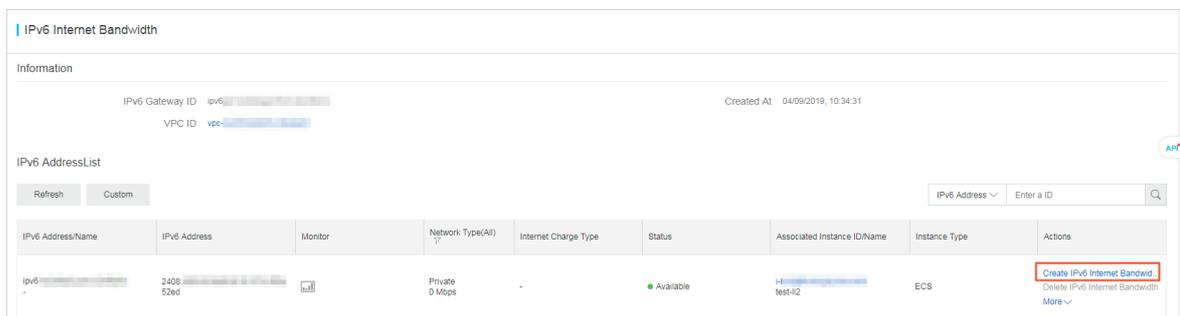
- Linux images:
 - CentOS 7.6 IPV6 64Bit
 - CentOS 6.10 64Bit
 - SUSE Linux Enterprise Server 12 SP4 64Bit
 - Windows Server images
- If the image of your ECS instance does not support DHCPv6, you must manually configure an IPv6 address for the ECS instance. We recommend that you refer to the related documentation for each image for configuration guidance.

Step 4: Purchase an IPv6 Internet bandwidth plan

By default, IPv6 addresses are only used for communication within private networks. If you want to allow an instance that is assigned an IPv6 address to access the Internet or receive requests from IPv6 clients over the Internet, you must purchase an Internet bandwidth plan for the IPv6 address.

Perform the following steps to purchase an Internet bandwidth plan for the IPv6 address:

1. In the top navigation bar, choose **Products > Networking > IPv6 Gateway**.
2. Select the region where the IPv6 gateway is created.
3. On the **IPv6 Gateway** page, find the IPv6 gateway that you want to manage and click **Manage** in the **Actions** column.
4. In the left-side navigation pane, click **IPv6 Internet Bandwidth**.
5. On the **IPv6 Internet Bandwidth** page, find the IPv6 address that you want to manage and click **Enable IPv6 Internet Bandwidth** in the **Actions** column.



6. Select a bandwidth plan and click **OK**.

The maximum IPv6 Internet bandwidth for an IPv6 gateway of the Free, Enterprise, or Enhanced Edition is 2 Gbit/s.

Step 5: Configure security group rules

IPv4 and IPv6 addresses are independent of each other. If the current security group rules do not apply to your IPv6 services, you must configure security group rules for the ECS instances to regulate communication with IPv6 addresses.

For more information about how to configure security rules, see [Add a security group rule](#) the **Add security group rules** chapter in *Security Groups of Elastic Compute Service User Guide*.

Step 6: Test the network connectivity

Log on to an ECS instance and ping an IPv6 service to test the network connectivity.

```
[root@iZhp3aehva ~]# ping6 ipv6.baidu.com
PING ipv6.baidu.com(2400:da00:2::29 (2400:da00:2::29)) 56 data bytes
64 bytes from 2400:da00:2::29 (2400:da00:2::29): icmp_seq=1 ttl=45 time=77.1 ms
64 bytes from 2400:da00:2::29 (2400:da00:2::29): icmp_seq=2 ttl=45 time=77.1 ms
64 bytes from 2400:da00:2::29 (2400:da00:2::29): icmp_seq=3 ttl=45 time=77.0 ms
^C
--- ipv6.baidu.com ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2002ms
rtt min/avg/max/mdev = 77.070/77.101/77.127/0.227 ms
[root@iZhp3aehva ~]#
```

17.4. VPCs and VSwitches

17.4.1. Overview

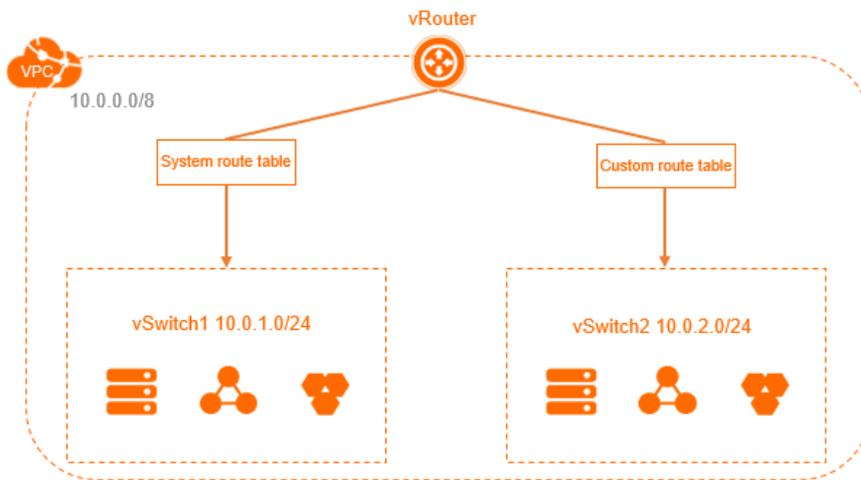
Before you can use cloud resources in a virtual private cloud (VPC), you must create a VPC and a vSwitch. You can create more than one vSwitch to divide a VPC into multiple subnets. By default, the subnets in a VPC can communicate with each other.

VPCs and vSwitches

A VPC is a private network dedicated for your use.

Note You cannot directly deploy cloud resources in a VPC. You must deploy cloud resources in vSwitches of a VPC.

vSwitches are basic components of VPCs and are used to connect different cloud resources. You can deploy a VPC only in one region and cannot deploy a VPC across regions. However, a VPC contains all zones of the region to which the VPC belongs. You can create one or more vSwitches in a zone to divide a VPC into multiple subnets.



CIDR blocks and IP addresses

VPCs support both IPv4 and IPv6. By default, VPCs use IPv4. You can enable IPv6 based on your business requirements.

VPCs support the dual-stack mode. In dual-stack mode, resources in a VPC can communicate through both IPv4 and IPv6 addresses. IPv4 and IPv6 addresses are independent of each other. Therefore, you must configure routes and security groups for both IPv4 and IPv6 addresses.

The following table lists the differences between IPv4 and IPv6 addresses.

| IPv4 VPC | IPv6 VPC |
|---|---|
| An IPv4 address is 32 bits in length and contains four groups. Each group consists of at most three decimal digits. | An IPv6 address is 128 bits in length and contains eight groups. Each group consists of four hexadecimal digits. |
| By default, IPv4 is enabled. | You can enable IPv6. |
| The subnet mask of a VPC CIDR block can range from /8 to /24. | The subnet mask of a VPC CIDR block is /61. |
| The subnet mask of a vSwitch CIDR block can range from /16 to /29. | The subnet mask of a vSwitch CIDR block is /64. |
| You can specify an IPv4 CIDR block. | You cannot specify an IPv6 CIDR block. The system automatically assigns an IPv6 CIDR block to your VPC from the IPv6 address pool. |
| Supported by all instance types. | Not supported by specific instance types. For more information, see Instance types under What is ECS? in the <i>Elastic Compute Service User Guide</i> . |
| IPv4 elastic IP addresses (EIPs) are supported. | IPv6 EIPs are not supported. |
| VPN gateways and NAT gateways are supported. | VPN gateways and NAT gateways are not supported. |

By default, IPv4 and IPv6 addresses provided for VPCs support only communication over private networks. Cloud resources deployed in different vSwitches that belong to the same VPC can communicate with each other over private networks. To connect a VPC to another VPC or a data center, you can use Express Connect, or VPN Gateway.

To enable cloud resources in a VPC to communicate with the Internet, perform the following operations:

- Through IPv4 addresses

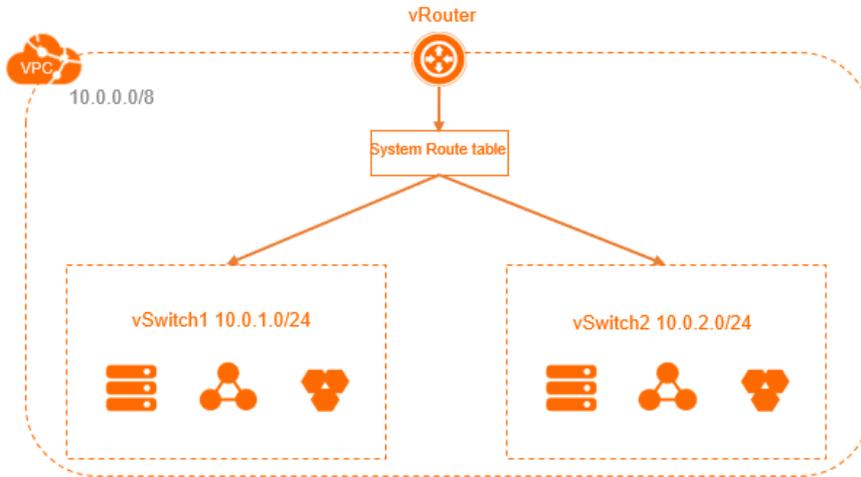
You can configure NAT gateways or associate EIPs with Elastic Compute Service (ECS) instances in a VPC. This way, these ECS instances can communicate with the Internet through IPv4 addresses.

- Through IPv6 addresses

To enable cloud resources in a VPC to communicate with the Internet through IPv6 addresses, you must purchase an IPv6 Internet bandwidth plan. You can configure egress-only rules for IPv6 addresses. This allows cloud resources in the VPC to access the Internet through IPv6 addresses. However IPv6 clients cannot access the cloud resources over the Internet.

Routes

The system automatically creates a system route table and adds system route entries to control the traffic of the VPC. A VPC has only one system route table. You cannot create or delete a system route table.



If multiple route entries match the destination CIDR block, the route entry with the largest prefix prevails and determines the next hop. This ensures that the traffic is routed to the most precise destination. You can also add a custom route entry to route traffic to a specified destination.

17.4.2. VPC management

17.4.2.1. Create a VPC

A virtual private cloud (VPC) is a private network dedicated for your use. You have full control over your VPC. For example, you can specify CIDR blocks, and configure route tables and gateways for your VPC. You can create a VPC, and then use Alibaba Cloud resources, such as Elastic Compute Service (ECS), ApsaraDB RDS, and Server Load Balancer (SLB) instances in the VPC. This topic describes how to create a VPC.

Prerequisites

Before you create a VPC, you must plan your networks. For more information, see [网络规划](#).

Procedure

1. [Log on to the VPC console](#).
2. In the top navigation bar, select the region where you want to deploy the VPC.

Note The VPC must be deployed in the same region as that of the cloud resources that you want to deploy in this VPC.

3. On the **VPC** page, click **Create VPC**.
4. On the **Create VPC** page, configure the VPC and click **OK**. The following table describes the parameters.

| Parameter | Description |
|---------------------|---|
| Organization | Select the organization to which the VPC belongs. |
| Resource Set | Select the resource set to which the VPC belongs. |
| Region | Select the region where you want to deploy the VPC. |

| Parameter | Description |
|-----------------|--|
| Sharing Scope | <p>Select the participants that can use the VPC to create resources.</p> <ul style="list-style-type: none"> ◦ Current Resource Set : Only the administrator of the current resource set can create resources for the shared VPC. ◦ Current Organization and Subordinate Organization: Only the administrators of the current organization and its subordinate organization can create resources for the shared VPC. ◦ Current Organization: Only the administrator of the current organization can use create resources for the shared VPC. |
| VPC Name | <p>Enter a name for the VPC.</p> <p>The name must be 2 to 128 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code>.</p> |
| IPv4 CIDR Block | <p>Select an IPv4 CIDR block for the VPC. The following settings are supported:</p> <ul style="list-style-type: none"> ◦ Recommended CIDR Block: You can use one of the following standard IPv4 CIDR blocks: 192.168.0.0/16 and 172.16.0.0/16. ◦ Custom CIDR Block: You can use 192.168.0.0/16, 172.16.0.0/16, or a subset of these CIDR blocks. The subnet mask must be 8 to 28 bits in length. For example, enter 192.168.0.0/24. <div style="background-color: #e6f2ff; padding: 5px;"> <p> Note After you create a VPC, you cannot change its IPv4 CIDR block.</p> </div> |
| IPv6 CIDR Block | <p>Specify whether to assign an IPv6 CIDR block.</p> <ul style="list-style-type: none"> ◦ Do Not Assign: The system does not assign an IPv6 CIDR block to the VPC. ◦ Assign: The system automatically assigns an IPv6 CIDR block to the VPC. <p>If you set this parameter to Assign, the system automatically creates a free IPv6 gateway for this VPC, and assigns an IPv6 CIDR block with the subnet mask /56, such as 2XX1:db8::/56. By default, IPv6 addresses can be used for communication within only private networks. If you want to use the IPv6 address to access the Internet or to be accessed by IPv6 clients over the Internet, you must purchase an Internet bandwidth plan for the IPv6 address. For more information, see Enable Internet connectivity for an IPv6 address the Activate IPv6 Internet bandwidth section of the Manage IPv6 Internet bandwidth topic of the <i>IPv6 gateway user guide</i>.</p> |
| Description | <p>Enter a description for the VPC.</p> <p>The description must be 2 to 256 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code>.</p> |

17.4.2.2. Add a secondary IPv4 CIDR block

This topic describes how to expand a virtual private cloud (VPC) by adding a secondary IPv4 CIDR block to the VPC.

Prerequisites

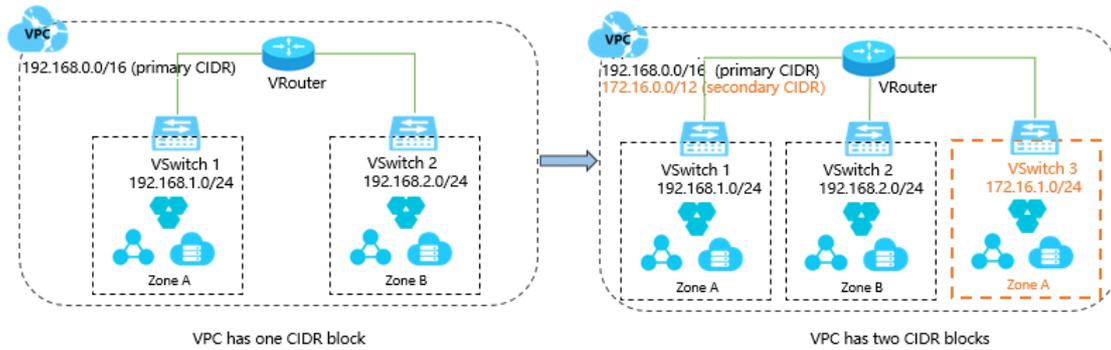
A VPC is created. For more information, see [Work with VPCs](#).

Context

When you create a VPC, the IPv4 CIDR block that you specified is the primary CIDR block. After the VPC is created, the primary IPv4 CIDR block of the VPC cannot be modified. However, you can add a secondary IPv4 CIDR block to expand the VPC. After you add the secondary IPv4 CIDR block, both the primary and secondary IPv4 CIDR blocks

take effect. You can create a vSwitch with the primary or secondary IPv4 CIDR block. However, each vSwitch belongs to only one VPC CIDR block.

The system automatically adds a vSwitch route to the VPC route table when you create a vSwitch with the primary or secondary IPv4 CIDR block. The destination CIDR block of a vSwitch route is the CIDR block with which the vSwitch is created. The CIDR block range cannot be the same as or larger than those of other routes in the route table of the VPC.



Note You can add only one secondary IPv4 CIDR block to a VPC and cannot increase the quota.

Procedure

1. Log on to the VPC console.
2. In the top navigation bar, select the region where the VPC is deployed.
3. On the VPC page, find the VPC and click **Manage** in the **Actions** column.
4. On the **CIDRs** tab, click **Add IPv4 CIDR**.
5. In the **Add Secondary CIDR** panel, set the following parameters and click **OK**.

| Parameter | Description |
|----------------|--|
| VPC | The VPC to which you want to add the secondary IPv4 CIDR block. |
| Secondary CIDR | <p>Select a method to configure the secondary IPv4 CIDR block:</p> <ul style="list-style-type: none"> ◦ Default CIDR Block: You can specify one of the following standard IPv4 CIDR blocks as the secondary IPv4 CIDR block: 192.168.0.0/16 and 172.16.0.0/12. ◦ Custom CIDR Block: You can specify one of the following standard IPv4 CIDR blocks and their subnets as the secondary IPv4 CIDR block: 192.168.0.0/16 and 172.16.0.0/12. <p>When you add a secondary IPv4 CIDR block, take note of the following limits:</p> <ul style="list-style-type: none"> ◦ The CIDR block cannot start with 0. The subnet mask must be 8 to 24 bits in length. ◦ The secondary IPv4 CIDR block cannot overlap with the primary IPv4 CIDR block or an existing secondary IPv4 CIDR block. <p>For example, if the primary IPv4 CIDR block of a VPC is 192.168.0.0/16, you cannot specify one of the following CIDR blocks as the secondary IPv4 CIDR block:</p> <ul style="list-style-type: none"> ▪ A CIDR block that is larger than 192.168.0.0/16, such as 192.168.0.0/8. ▪ 192.168.0.0/16. ▪ A CIDR block that is smaller than 192.168.0.0/16, such as 192.168.0.0/24. |

What's next

[Work with vSwitches](#)

17.4.2.3. Delete a secondary IPv4 CIDR block

This topic describes how to delete a secondary IPv4 CIDR block of a virtual private cloud (VPC). However, you cannot delete the primary IPv4 CIDR block of a VPC.

Prerequisites

Before you delete a secondary IPv4 CIDR block, make sure that you have deleted the vSwitch that is created within the secondary IPv4 CIDR block. For more information, see [Delete a VSwitch](#).

Procedure

1. [Log on to the VPC console](#).
2. In the top navigation bar, select the region where the VPC is deployed.
3. On the **VPCs** page, find the VPC that you want to manage and click **Manage** in the **Actions** column.
4. On the **CIDRs** tab, find the secondary IPv4 CIDR block that you want to delete and click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

17.4.2.4. Modify the name and description of a VPC

This topic describes how to modify the name and description of a virtual private cloud (VPC).

Procedure

1. [Log on to the VPC console](#).
2. In the top navigation bar, select the region where your VPC is deployed.
3. On the **VPCs** page, find the target VPC network and click **Manage** in the **Actions** column.
4. In the **VPC Details** section, click **Edit** next to **Name**. In the dialog box that appears, enter a new name for the VPC and click **OK**.
The name must be 2 to 128 characters in length and can contain letters, Chinese characters, digits, underscores (_), and hyphens (-). It must start with a letter or a Chinese character.
5. Click **Edit** next to **Description**. In the dialog box that appears, enter a new description, and click **OK**.
The description must be 2 to 256 characters in length and cannot start with `http://` or `https://`.

17.4.2.5. Delete a VPC

This topic describes how to delete a virtual private cloud (VPC). After you delete a VPC, the vRouters and route tables associated with the VPC are also deleted.

Prerequisites

Before you delete a VPC, make sure that the following requirements are met:

- No vSwitch exists in the VPC. If a vSwitch exists in the VPC, delete the vSwitch first. For more information, see [Delete a VSwitch](#).
- No IPv6 gateway exists in the VPC. If an IPv6 gateway exists in the VPC, delete the IPv6 gateway first.

Procedure

1. [Log on to the VPC console](#).
2. In the top navigation bar, select the region where the VPC is deployed.

3. On the VPCs page, find the VPC that you want to delete, and click **Delete** in the **Actions** column.
4. In the **Delete VPC** message, click **OK**.

17.4.3. VSwitch management

17.4.3.1. Create a vSwitch

A vSwitch is a basic network component that connects different cloud resources in a virtual private cloud (VPC).

Context

After you create a VPC, you can create vSwitches to divide the VPC into one or more subnets. vSwitches within the same VPC can communicate with each other. Cloud resources must be deployed in vSwitches. You can deploy applications in vSwitches that belong to different zones to improve service availability.

 **Note** vSwitches do not support multicasting or broadcasting.

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **vSwitches**.
3. Select the region of the VPC in which you want to create a vSwitch.
4. On the **vSwitches** page, click **Create vSwitch**.
5. On the **vSwitch** page, set the following parameters and click **Submit**.

| Parameter | Description |
|----------------------|---|
| Organization | Select the organization to which the vSwitch belongs. |
| Resource Set | Select the resource set to which the vSwitch belongs. |
| Region | Select the region where you want to deploy the vSwitch. |
| Zone | <p>Select the zone to which the vSwitch belongs.</p> <p>In a VPC, each vSwitch can be deployed in only one zone. You cannot deploy a vSwitch across zones. However, you can deploy cloud resources in vSwitches that belong to different zones to achieve cross-zone disaster recovery.</p> <p> Note A cloud instance can be deployed in only one vSwitch.</p> |
| Sharing Scope | <p>Select the participants that can use the vSwitch to create resources.</p> <ul style="list-style-type: none"> ◦ Current Resource Set: Only the administrator of the current resource set can create resources in the vSwitch. ◦ Current Organization and Subordinate Organization: Only the administrators of the current organization and its subordinate organizations can create resources in the vSwitch. ◦ Current Organization: Only the administrator of the current organization can create resources in the vSwitch. |
| VPC | Select the VPC for which you want to create the vSwitch. |

| Parameter | Description |
|---|--|
| Dedicated for Out-of-cloud Physical Machines | Specify whether the vSwitch is dedicated for bare-metal servers. For more information about bare-metal servers, see the Bare-metal servers in VPCs topic in <i>Bare-metal Server Management Service User Guide</i> . |
| vSwitch Name | Enter a name for the vSwitch. The name must be 2 to 128 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . |
| IPv4 CIDR Block | Enter an IPv4 CIDR block for the vSwitch. <ul style="list-style-type: none"> You must specify the IP address range of the vSwitch in CIDR notation. The subnet mask must be 16 to 29 bits in length. It means that 8 to 65,536 IP addresses can be provided. The CIDR block of a vSwitch must be a proper subset of the CIDR block of the VPC to which the vSwitch belongs. The first IP address and last three IP addresses of a vSwitch are reserved. For example, if the CIDR block of a vSwitch is 192.168.1.0/24, the IP addresses 192.168.1.0, 192.168.1.253, 192.168.1.254, and 192.168.1.255 are reserved. The CIDR block of a vSwitch cannot be the same as the destination CIDR block in a route entry of the VPC to which the vSwitch belongs. However, the CIDR block of the vSwitch can be a proper subset of the destination CIDR block of the route entry. After a vSwitch is created, you cannot modify its CIDR block. |
| IPv6 CIDR Block | Enter an IPv6 CIDR block for the vSwitch. <ul style="list-style-type: none"> You must check whether IPv6 is enabled for the specified VPC. If IPv6 is disabled, you cannot assign an IPv6 CIDR block to the vSwitch. If IPv6 is enabled, you can enter a decimal number ranging from 0 to 255 to define the last 8 bits of the IPv6 CIDR block of the vSwitch. For example, if the IPv6 CIDR block of the VPC is 2xx1:db8::/64, specify 255 to define the last 8 bits of the IPv6 CIDR block. In this case, the IPv6 CIDR block of the vSwitch is 2XX1:db8:ff::/64. ff is the hexadecimal value of 255. |
| Description | Enter a description for the vSwitch. The description must be 2 to 256 characters in length, and can contain digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . |

17.4.3.2. Create cloud resources in a vSwitch

You cannot directly deploy cloud resources in a virtual private cloud (VPC). You can deploy cloud resources only in a vSwitch that belongs to a VPC. This topic describes how to create cloud resources in a vSwitch.

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **vSwitches**.
3. Select the region of the VPC to which the vSwitch belongs.
4. On the **vSwitches** page, find the vSwitch, click **Create** in the **Actions** column, and select the cloud resource that you want to create.

You can create Elastic Compute Service (ECS) instances, ApsaraDB RDS instances, and Server Load Balancer

(SLB) instances in a vSwitch.

5. On the page that appears, set the parameters.

17.4.3.3. Modify a vSwitch

This topic describes how to modify the name and description of a vSwitch.

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **vSwitches**.
3. Select the region of the VPC to which the vSwitch belongs.
4. On the **vSwitches** page, find the vSwitch that you want to manage and click **Manage** in the **Actions** column.
5. In the **vSwitch Basic Information** section, click **Edit** next to **Name** to modify the name of the vSwitch.
The name must be 2 to 128 characters in length, and can contain letters, digits, underscores (_), and hyphens (-). The name must start with a letter.
6. Click **Edit** next to **Description** to modify the description of the vSwitch.
The description must be 2 to 256 characters in length and cannot start with `http://` or `https://`.

17.4.3.4. Delete a vSwitch

This topic describes how to delete a vSwitch. After you delete a vSwitch, you cannot deploy cloud resources in it.

Prerequisites

Before you delete a vSwitch, make sure that the following requirements are met:

- All instances deployed in the vSwitch are deleted, such as Elastic Compute Service (ECS) instances, Server Load Balancer (SLB) instances, and ApsaraDB RDS instances.
- All resources associated with the vSwitch are deleted, such as high-availability virtual IP addresses (HAVIPs) and Source Network Address Translation (SNAT) entries.

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **vSwitches**.
3. Select the region of the VPC to which the vSwitch belongs.
4. On the **vSwitch** page, find the vSwitch that you want to delete and click **Delete** in the **Actions** column.
5. In the **Delete vSwitch** message, click **OK**.

17.5. Route tables

17.5.1. Overview

After you create a virtual private cloud (VPC), the system creates a system route table for the VPC and adds system routes to the route table. The system routes are used to route traffic within the VPC. You cannot create or delete system routes. However, you can create custom routes to route traffic from specific CIDR blocks to a specified destination.

Route tables

After you create a VPC, the system creates a system route table to manage routes of the VPC. By default, vSwitches in the VPC use this route table. You cannot create or delete the system route table of a VPC.

Each entry in a route table is a *route entry*. A route entry specifies the destination of traffic and consists of the destination CIDR block, next hop type, and next hop. Route entries include system route entries and custom route entries.

Regions that support custom route tables

The following table describes the regions that support custom route tables by default.

| Area | Supported region |
|-----------------------|--|
| Asia Pacific | China (Qingdao), China (Zhangjiakou), China (Hohhot), China (Ulanqab), China (Shanghai), China (Heyuan), China (Guangzhou), China (Chengdu), China (Hong Kong), Japan (Tokyo), Singapore (Singapore), Australia (Sydney), Malaysia (Kuala Lumpur), and Indonesia (Jakarta) |
| Europe & Americas | US (Silicon Valley), US (Virginia), Germany (Frankfurt), and UK (London) |
| Middle East and India | India (Mumbai) and UAE (Dubai) |

The custom route table feature is in public preview in the following regions. You can apply for participating in the public review.

| Area | Supported region |
|--------------|---|
| Asia Pacific | China (Beijing), China (Shanghai), and China (Shenzhen) |

System routes

After you create a VPC, the system automatically adds the following system routes to the route table:

- A route entry with a destination CIDR block of 100.64.0.0/10. This route is used for communication among cloud resources within the VPC.
- Route entries whose destination CIDR blocks are the same as the CIDR blocks of the vSwitches in the VPC. These routes are used for communication among cloud resources within the vSwitches.

For example, if you create a VPC whose CIDR block is 192.168.0.0/16 and two vSwitches whose CIDR blocks are 192.168.1.0/24 and 192.168.0.0/24, three system routes are automatically added to the route table of the VPC.

The following table describes the system routes.

| Destination CIDR block | Next hop | Route entry type |
|------------------------|----------|------------------|
| 100.64.0.0/10 | - | System route |
| 192.168.1.0/24 | - | System route |
| 192.168.0.0/24 | - | System route |

Custom routes

You can add custom routes to replace system routes or route traffic to a specified destination. You can specify the following types of next hops when you create a custom route:

- Elastic Compute Service (ECS) instance: Traffic that is destined for the destination CIDR block is routed to a specified ECS instance in the VPC.

You can select this type if you want to access the Internet or other applications through the applications deployed on the ECS instance.

- VPN gateway: Traffic destined for the destination CIDR block is routed to a specified VPN gateway.

You can select this type if you want to connect a VPC to another VPC or an on-premises network through the VPN gateway.

- **NAT gateway:** Traffic destined for the destination CIDR block is routed to a specified NAT gateway.
You can select this type if you want to connect a VPC to the Internet through the NAT gateway.
- **Router interface (to VPC):** Traffic that is destined for the destination CIDR block is routed to a specified VPC.
You can select this type if you want to connect two VPCs through Express Connect circuits.
- **Router interface (to VBR):** Traffic that is destined for the destination CIDR block is routed to a specified virtual border router (VBR).
You can select this type if you want to connect a VPC to an on-premises network through Express Connect circuits.
- **Secondary ENI:** Traffic that is destined for the destination CIDR block is routed to a specified secondary elastic network interface (ENI).

IPv6 routes

If IPv6 is enabled for your VPC, the following route entries are automatically added to the system route table of the VPC:

- A custom route entry whose destination CIDR block is `::/0` and whose next hop is the IPv6 gateway. Cloud resources deployed in the VPC use this route to access the Internet through IPv6 addresses.
- A system route entry of which the destination CIDR block is the IPv6 CIDR block of a vSwitch. This route is used for communication within the vSwitch.

Routing rules

If multiple route entries match the destination CIDR block, the route entry with the largest prefix prevails and determines the next hop. This ensures that the traffic is routed to the most precise destination.

For example, the following table describes the route table of a VPC.

| Destination CIDR block | Next hop type | Next hop | Route entry type |
|------------------------|---------------|------------|------------------|
| 100.64.0.0/10 | - | - | System route |
| 192.168.0.0/24 | - | - | System route |
| 0.0.0.0/0 | Instance | i-12345678 | Custom route |
| 10.0.0.0/24 | Instance | i-87654321 | Custom route |

The route entries that are destined for `100.64.0.0/10` and `192.168.0.0/24` are system route entries. The route entries that are destined for `0.0.0.0/0` and `10.0.0.0/24` are custom route entries. Traffic that is destined for `0.0.0.0/0` is routed to the ECS instance `i-12345678`, whereas traffic that is destined for `10.0.0.0/24` is routed to the ECS instance `i-87654321`. Based on the preceding rule, traffic that is destined for `10.0.0.1` is routed to the ECS instance `i-87654321`, whereas traffic that is destined for `10.0.1.1` is routed to the ECS instance `i-12345678`.

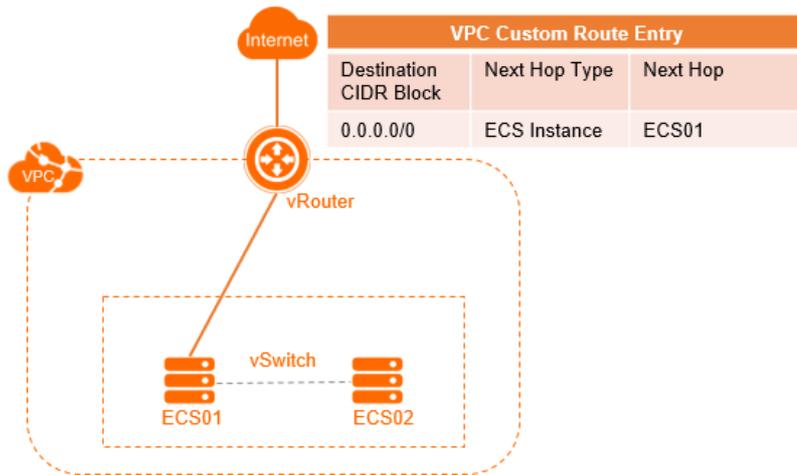
Routing examples

You can add custom route entries to a route table to control inbound and outbound traffic that are transmitted over the VPC.

- **Routes within a VPC**

The following figure shows a NAT gateway that is deployed on an ECS instance (ECS 01) in a VPC. To enable the cloud resources in the VPC to access the Internet through the ECS instance, you must add the following route entry to the route table.

| Destination CIDR block | Next hop type | Next hop |
|------------------------|---------------|----------|
| 0.0.0.0/0 | ECS instance | ECS01 |



• Connect two VPCs through Express Connect

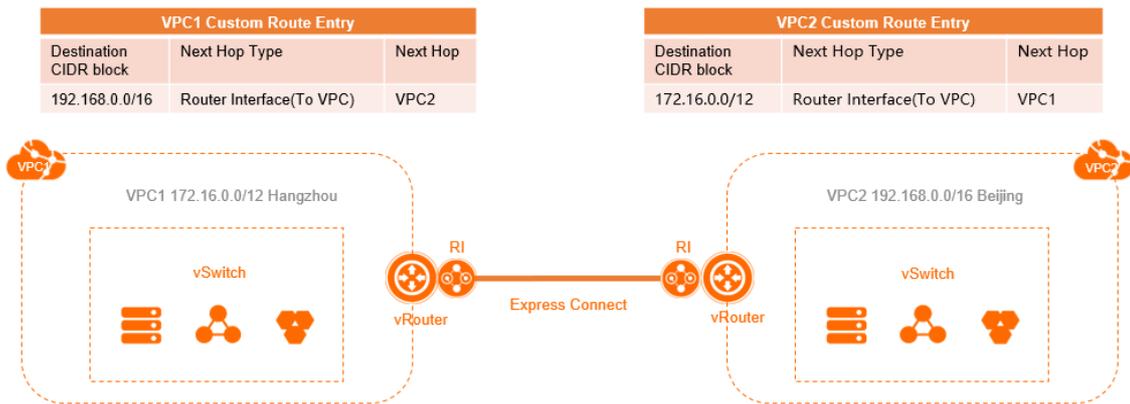
The following figure shows that VPC 1 (172.16.0.0/12) is connected to VPC 2 (192.168.0.0/16) through Express Connect. After you create router interfaces, you must add the following route entries to the VPCs:

- Add the following route entry to VPC 1

| Destination CIDR block | Next hop type | Next hop |
|------------------------|---------------------------|----------|
| 192.168.0.0/16 | Router interface (to VPC) | VPC2 |

- Add the following route entry to VPC 2

| Destination CIDR block | Next hop type | Next hop |
|------------------------|---------------------------|----------|
| 172.16.0.0/12 | Router interface (to VPC) | VPC1 |



• Connect two VPCs through a VPN gateway

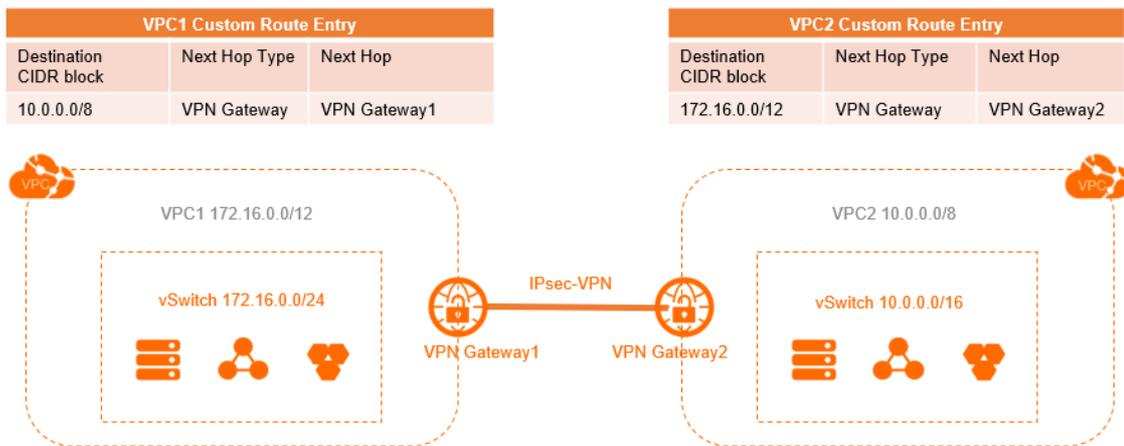
The following figure shows that VPC 1 (172.16.0.0/12) is connected to VPC 2 (10.0.0.0/8) through a VPN gateway. After you configure the VPN gateway, you must add the following route entries to the VPCs.

- o Add the following route entry to VPC 1

| Destination CIDR block | Next hop type | Next hop |
|------------------------|---------------|---------------|
| 10.0.0.0/8 | VPN gateway | VPN gateway 1 |

- o Add the following route entry to VPC 2

| Destination CIDR block | Next hop type | Next hop |
|------------------------|---------------|---------------|
| 172.16.0.0/12 | VPN gateway | VPN gateway 2 |



- Connect a VPC to a data center through Express Connect

The following figure shows that a VPC is connected to an on-premises network through Express Connect. After you configure a connection over an Express Connect circuit and a virtual border router (VBR), you must add the following route entries:

- o Add the following route entry to the VPC

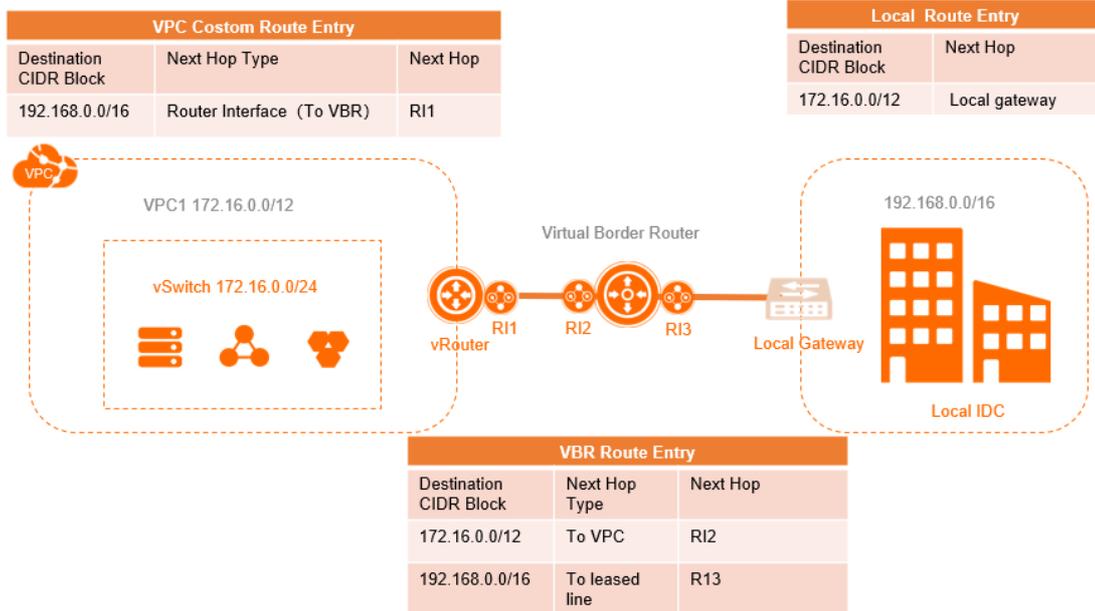
| Destination CIDR block | Next hop type | Next hop |
|------------------------|-------------------------------------|-----------------------|
| 192.168.0.0/16 | Router interfaces (general routing) | Router interface RI 1 |

- o Add the following route entry to the VBR

| Destination CIDR block | Next hop type | Next hop |
|------------------------|-------------------------|-----------------------|
| 192.168.0.0/16 | Express Connect circuit | Router interface RI 3 |
| 172.16.0.0/12 | VPC | Router interface RI 2 |

- o Add the following route entry to the on-premises network

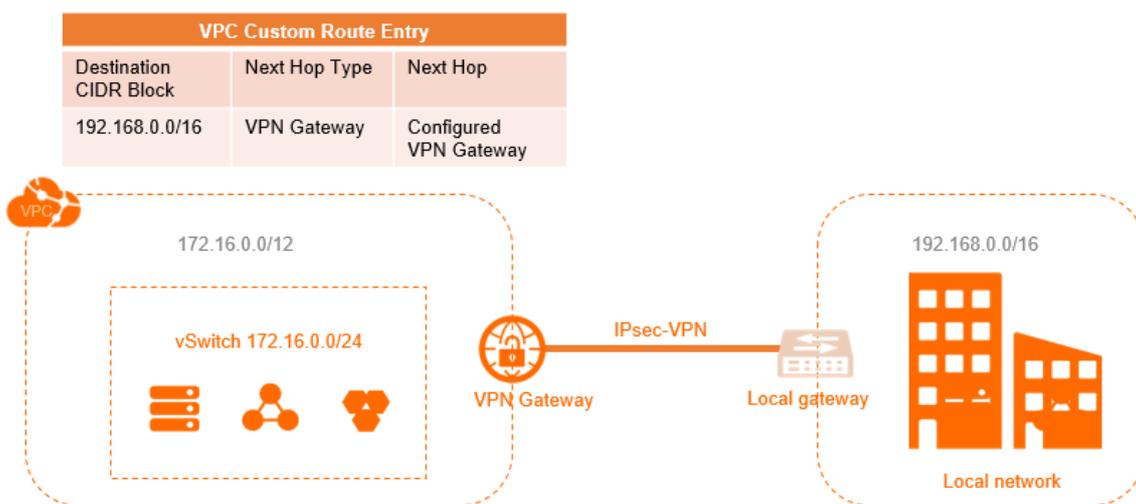
| Destination CIDR block | Next hop type | Next hop |
|------------------------|---------------|----------------------------|
| 172.16.0.0/12 | - | On-premises gateway device |



- Connect a VPC to a data center through a VPN gateway

The following figure shows that a VPC (172.16.0.0/12) is connected to a data center (192.168.0.0/16) through a VPN gateway. After you configure the VPN gateway, you must add the following route entry to the VPC.

| Destination CIDR block | Next hop type | Next hop |
|------------------------|---------------|----------------------------------|
| 192.168.0.0/16 | VPN gateway | The VPN gateway that you created |



17.5.2. Add a custom route entry

This topic describes how to add a custom route entry. After you create a virtual private cloud (VPC), the system creates a system route table and adds system route entries to the route table. The system route entries are used to route traffic within the VPC. You cannot create or delete system route entries. However, you can create custom route entries to route traffic from source CIDR blocks to specific destinations.

Context

Each item in the route table is a route entry. A route entry, which consists of the destination CIDR block, type of next hop, and next hop, specifies the destination for network traffic. Route entries include system route entries and custom route entries.

Procedure

1. [Log on to the VPC console.](#)
2. In the left-side navigation pane, click **Route Tables**.
3. In the top navigation bar, select the region to which the route table belongs.
4. On the **Route Tables** page, find the route table and click **Manage** in the **Actions** column.
5. In the **Route Table Details** section, click **Add Route Entry**.
6. In the **Add Route Entry** panel, set the following parameters and click **OK**:

| Parameter | Description |
|-------------------------------|---|
| Destination CIDR block | Enter a destination CIDR block for the route entry. |

| Parameter | Description |
|--|---|
| Next Hop Type | <p>Select the type of next hop. Valid values:</p> <ul style="list-style-type: none"> ◦ ECS Instance: Traffic destined for the specified CIDR block is routed to the specified Elastic Compute Service (ECS) instance. <p>Select this type if you want to route traffic to an ECS instance for centralized traffic forwarding and management. For example, you can configure an ECS instance as the Internet-facing gateway to route traffic from other ECS instances to the Internet.</p> <ul style="list-style-type: none"> ◦ HaVip Address: Traffic destined for the specified CIDR block is routed to the high-availability virtual IP address (HAVIP) that you select. ◦ VPN Gateway: Traffic destined for the specified CIDR block is routed to the specified VPN gateway. ◦ NAT Gateway: Traffic destined for the specified CIDR block is routed to the specified NAT gateway. ◦ Secondary ENI: Traffic destined for the specified CIDR block is routed to the secondary elastic network interface (ENI) that you select. ◦ Router Interface (To VPC): Traffic destined for the specified CIDR block is routed to the specified VPC. <p>Select this type if you want to connect VPCs through Express Connect circuits.</p> <ul style="list-style-type: none"> ◦ Router Interface (To VBR): Traffic destined for the specified CIDR block is routed to the router interface that is associated with a virtual border router (VBR). <p>Select this type if you want to connect the VPC to a data center through an Express Connect circuit.</p> <p>If you specify Router Interface (To VBR), you must select a routing mode:</p> <ul style="list-style-type: none"> ▪ General Routing: Select an associated router interface. ▪ Active/Standby Routing: Select two instances as the next hop. The active route has a weight of 100 and the standby route has a weight of 0. The standby route takes effect when the active route fails the health check. ▪ Load Balancing: Select two to four router interfaces as the next hop. The peer router of each router interface must be a VBR. Valid values of the instance weight: 1 to 255. The value must be an integer and the default value is 100. The weight of each instance must be the same. This way, traffic can be evenly distributed to the next-hop instances. |
| ECS Instance/VPN Gateway/NAT Gateway/Secondary ENI/HAVIP/VPC | Select an instance as the next hop. |

17.5.3. Export route entries

This topic describes how to export route entries from a route table for backup.

Procedure

1. [Log on to the VPC console.](#)
2. In the left-side navigation pane, click **Route Tables**.
3. In the top navigation bar, select the region to which the route table belongs.

4. On the **Route Tables** page, find the route table and click **Manage** in the **Actions** column.
5. In the **Route Table Details** section, click the **Route Entry List** tab, and then click **Export**.

The route entries are exported to a .csv file in your local computer.

17.5.4. Modify a route table

This topic describes how to modify the name and description of a route table.

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **Route Tables**.
3. In the top navigation bar, select the region to which the route table belongs.
4. On the **Route Tables** page, find the route table and click its ID.
5. In the **Route Table Details** section, click **Edit** next to **Name**. In the dialog box that appears, enter a new name for the route table and click **OK**.

The name must be 2 to 128 characters in length and can contain letters, Chinese characters, digits, underscores (_), and hyphens (-). It must start with a letter or a Chinese character.

6. Click **Edit** next to **Description**. In the dialog box that appears, enter a new description of the route table, and click **OK**.

The description must be 2 to 256 characters in length and cannot start with `http://` or `https://`.

17.5.5. Delete a custom route entry

This topic describes how to delete a custom route entry. A route table consists of one or more route entries that determine which way to forward traffic. Note that system route entries cannot be deleted.

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **Route Tables**.
3. In the top navigation bar, select the region to which the route table belongs.
4. On the **Route Tables** page, find the route table and click its ID.
5. On the **Route Entry List** tab, find the target route entry and then click **Delete** in the **Actions** column.
6. In the **Delete Route Entry** dialog box, click **OK**.

17.6. HAVIPs

17.6.1. Overview

High-availability virtual IP addresses (HAVIPs) are private IP addresses that can be created and released as independent resources. You can use HAVIPs with high-availability (HA) software such as Keepalived to provide active/standby services. This improves the availability of your services.

Overview

Each Elastic Compute Service (ECS) instance is assigned a private IP address as the primary IP address. You can associate HAVIPs with an ECS instance to increase the number of private IP addresses available for the ECS instance. Both the primary IP address and HAVIPs of an ECS instance can be used to access networks. In addition, you can use HAVIPs with HA software such as Keepalived to provide active/standby services. This improves the availability of your services. You can associate an HAVIP with ECS instances by using the following methods:

- Directly associate an HAVIP with ECS instances.

Each HAVIP can be associated with two ECS instances. After an HAVIP is associated with ECS instances, the ECS instances can send Address Resolution Protocol (ARP) messages to advertise the HAVIP. After the ECS instances advertise the HAVIP, one of the ECS instances serves as the primary instance, and the other ECS instance serves as the secondary instance. If the primary ECS instance is down, the secondary ECS instance takes over to provide services.

- Attach secondary elastic network interfaces (ENIs) to ECS instances. Then, associate the HAVIP with the secondary ENIs.

Each HAVIP can be associated with ENIs of two ECS instances. After the HAVIP is associated with the ENIs, the ECS instances can send ARP messages to advertise the HAVIP. After the ECS instances advertise the HAVIP, one of the ECS instances serves as the primary instance, and the other ECS instance serves as the secondary instance. If the primary ECS instance is down, the secondary ECS instance takes over to provide services.

 **Note** Before you associate an HAVIP with secondary ENIs, make sure that the secondary ENIs are attached to two ECS instances.

HAVIPs have the following features:

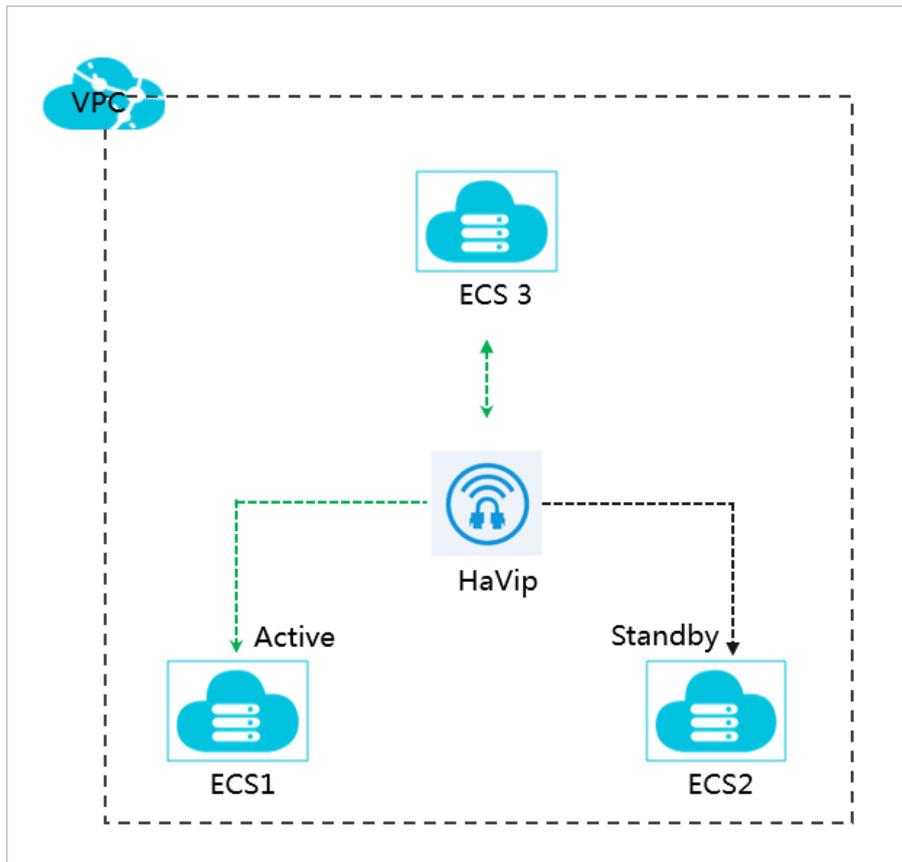
- HAVIPs are floating private IP addresses and are not statically assigned to specified ECS instances. HAVIPs can be associated with or disassociated from ECS instances through ARP announcements.
- An HAVIP can be associated with only ECS instances or ENIs that belong to the same vSwitch.
- You can associate each HAVIP with two ECS instances or two secondary ENIs. However, you cannot associate an HAVIP with an ECS instance and a secondary ENI at the same time.

Scenarios

HAVIPs support flexible configurations in the following scenarios:

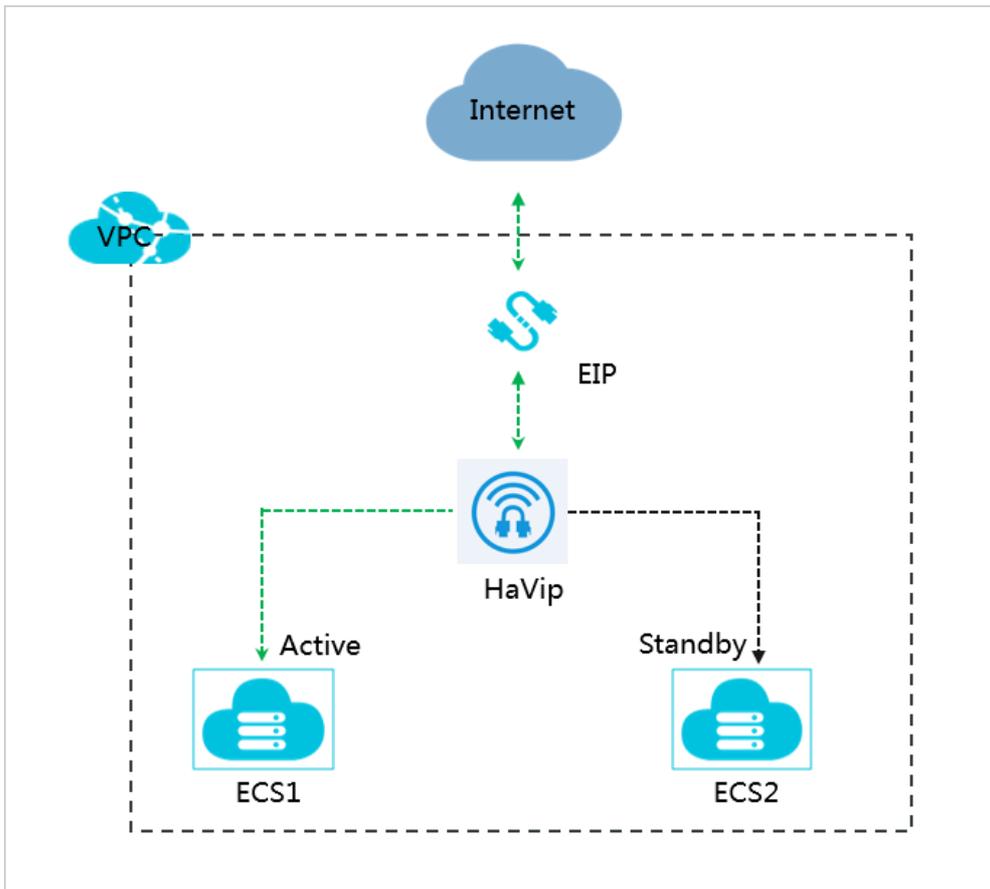
- Scenario 1: Internal-facing HA services

In the following figure, two ECS instances are assigned the same HAVIP. Keepalived is configured for the ECS instances to provide an internal-facing HA service. Other instances in the same virtual private cloud (VPC) can access this service. The HAVIP serves as the service address. If the primary ECS instance is down, the secondary ECS instance takes over. This improves the availability of your services.



- Scenario 2: Internet-facing HA services

In the following figure, two ECS instances are assigned the same HAVIP. Keepalived is configured and the HAVIP is associated with an elastic IP address (EIP) for the ECS instances to provide an Internet-facing HA service. The EIP that is associated with the HAVIP serves as the service address. If the primary ECS instance is down, the secondary ECS instance takes over. This improves the availability of your services.



Limits

Before you use HAVIPs, take note of the following limits.

| Item | Default |
|---|---------|
| Number of HAVIPs that can be created in each VPC | 5 |
| Number of HAVIPs that can be associated with each ECS instance | 5 |
| Number of HAVIPs that can be associated with each secondary ENI | 5 |
| Number of ECS instances that can be associated with each HAVIP | 2 |
| Number of secondary ENIs that can be associated with each HAVIP | 2 |
| Number of route entries that point to an HAVIP in each VPC | 5 |

| Item | Default |
|---|--|
| Whether HAVIPs support broadcast or multicast communication | Not supported <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note HAVIPs support only unicasting. To implement high availability through third-party software such as Keepalived, you must modify the configuration file to change the communication method to unicasting. </div> |

17.6.2. Create HAVIPs

High-availability virtual IP addresses (HAVIPs) are private IP addresses that can be created and released as independent resources. This topic describes how to create HAVIPs in the console.

Prerequisites

A VPC and vSwitches are created. For more information, see [Create a VPC](#) and [Work with vSwitches](#).

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **HaVip**.
3. In the top menu bar, select the region where you want to create the HAVIP.
4. On the **HaVip** page, click **Create HaVip**.
5. On the **Create a high-availability virtual IP address** page, set the following parameters and click **Submitted** to create an HAVIP.

| Parameter | Description |
|--------------------------------|--|
| Tissue | Select the organization to which the HAVIP belongs. |
| Resource Set | Select the resource set to which the HAVIP belongs. |
| Region | Select the region where you want to create the HAVIP. |
| Proprietary Network vpc | Select the VPC to which the HAVIP that you want to create belongs. |
| vswitch | Select the vSwitch to which the HAVIP that you want to create belongs. |
| Private IP Address | Specify a private IP address for the HAVIP. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note You must specify an idle private IP address that falls within the CIDR block of the vSwitch. </div> |

17.6.3. Associate HAVIPs with backend cloud resources

17.6.3.1. Associate an HAVIP with an ECS instance

This topic describes how to associate a high-availability virtual IP address (HAVIP) with an Elastic Compute Service (ECS) instance that is deployed in a virtual private clouds (VPC). After you associate an HAVIP with an ECS instance, the ECS instance can send Address Resolution Protocol (ARP) messages to advertise the HAVIP. This way, the ECS instance can use more than one private IP address. Each HAVIP can be associated with at most two ECS instances.

Prerequisites

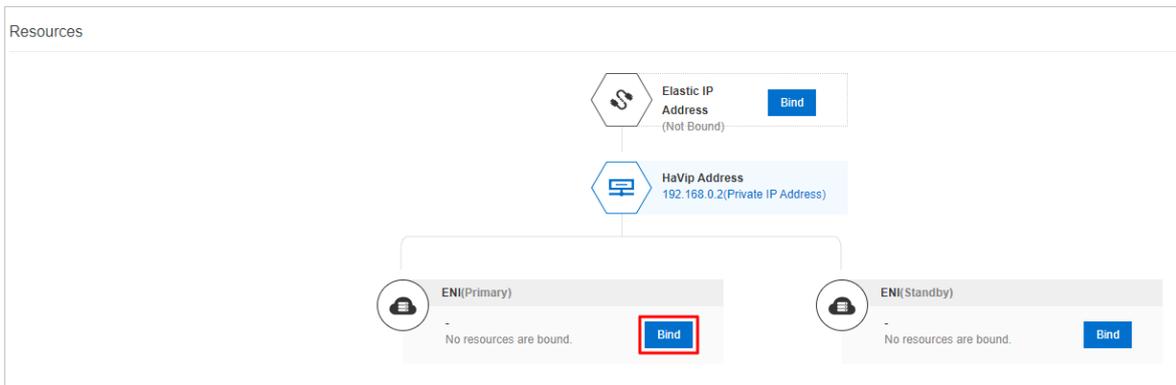
An ECS instance is created. For more information, see in the **Create an ECS instance** topic in *Quick Start of Elastic Compute Service User Guide*.

Context

You can associate an HAVIP with two ECS instances or two secondary elastic network interfaces (ENIs). However, you cannot associate an HAVIP with an ECS instance and a secondary ENI at the same time.

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **HaVip**.
3. In the top navigation bar, select the region where the HAVIP is created.
4. On the **HaVip Addresses** page, find the HAVIP that you want to manage and click **Manage** in the **Actions** column.
5. In the **Resources** section, find **ENI (Primary)** or **ENI (Standby)** and click **Bind**.



6. In the dialog box that appears, set the following parameters to associate the HAVIP with an ECS instance.

| Parameter | Description |
|----------------------|---|
| Resource Type | Select the type of resource with which you want to associate the HAVIP. Valid values: <ul style="list-style-type: none"> ◦ ECS Instances ◦ ENI In this example, ECS Instances is selected. |
| Bind Resource | Select the ECS instance with which you want to associate the HAVIP. The ECS instance that you select must meet the following requirements: <ul style="list-style-type: none"> ◦ The ECS instance is deployed in a VPC. ◦ The ECS instance and the HAVIP belong to the same vSwitch. |

7. Click **OK**.

17.6.3.2. Associate an HAVIP with an ENI

This topic describes how to associate a high-availability virtual IP address (HAVIP) with ENIs that are attached to Elastic Compute Service (ECS) instances. Then, the ECS instances can send Address Resolution Protocol (ARP) messages to advertise the HAVIP. This way, the ECS instances can use more than one private IP address.

Prerequisites

An ENI is created. For more information, see the **Create an Elastic Network Interface** topic in the **Elastic Network Interface** chapter of *Elastic Compute Service User Guide*.

Context

You can associate each HAVIP with two ECS instances or two ENIs. However, you cannot associate an HAVIP with an ECS instance and an ENI at the same time.

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **HaVip**.
3. In the top navigation bar, select the region where the HAVIP is created.
4. On the **HaVip Addresses** page, find the HAVIP that you want to associate and click its ID.
5. In the **Resources** section, find **ENI (Primary)** or **ENI (Standby)** and click **Bind**.
6. In the dialog box that appears, set the following parameters and click **OK**.

| Parameter | Description |
|----------------------|---|
| Resource Type | Select the type of resource with which you want to associate the HAVIP. Valid values: <ul style="list-style-type: none">◦ ECS Instance◦ ENI In this example, ENI is selected. |
| Bind Resource | Select the ENI with which you want to associate the HAVIP. The ENI and the HAVIP must belong to the same vSwitch. |

17.6.4. Associate HAVIPs with EIPs

This topic describes how to associate high-availability virtual IP addresses (HAVIPs) with elastic IP addresses (EIPs). After you associate an HAVIP with an EIP, the HAVIP can use the EIP to provide services over the Internet.

Prerequisites

An EIP is created. For more information, see **Create an EIP** in **Quick Start** of the *Elastic IP Address User Guide*.

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **HaVip**.
3. In the top menu bar, select the region where the HAVIP is created.
4. On the **HaVip Addresses** page, find the HAVIP that you want to manage, and choose **More > Bind EIP Address** in the **Actions** column.
5. In the dialog box that appears, select the EIP with which you want to associate the HAVIP and click **OK**.

The EIP with which you want to associate the HAVIP must meet the following requirements:

- The EIP and HAVIP are created in the same region.
- The EIP must be in the Available state.

17.6.5. Disassociate HAVIPs from backend cloud resources

17.6.5.1. Disassociate an HAVIP from an ECS instance

This topic describes how to disassociate a high-availability virtual private IP address (HAVIP) from an Elastic Compute Service (ECS) instance. After you disassociate an HAVIP from an ECS instance, the ECS instance cannot send Address Resolution Protocol (ARP) messages to advertise the HAVIP.

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **HaVip**.
3. In the top navigation bar, select the region where the HAVIP is created.
4. On the **HaVip** page, find the HAVIP that you want to manage and click **Manage** in the **Actions** column.
5. In the **Resources** section, find the ECS instance that you want to manage and click **Unbind**.
6. In the message that appears, click **OK**.

17.6.5.2. Disassociate an HAVIP from an ENI

This topic describes how to disassociate a high-availability virtual IP address (HAVIP) from an elastic network interface (ENI). After you disassociate an HAVIP from an ENI, the Elastic Compute Service (ECS) instance with which the ENI is associated cannot send Address Resolution Protocol (ARP) messages to advertise the HAVIP.

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **HaVip**.
3. In the top navigation bar, select the region where the HAVIP is created.
4. On the **HaVip** page, find the HAVIP that you want to disassociate and click **Manage** in the **Actions** column.
5. In the **Resources** section, find the ENI that you want to manage and click **Unbind**.
6. In the message that appears, click **OK**.

17.6.6. Disassociate an HAVIP from an EIP

This topic describes how to disassociate a high-availability virtual IP address (HAVIP) from an elastic IP address (EIP).

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **HaVip Addresses**.
3. Select the region to which the HAVIP belongs.
4. On the **HaVip Addresses** page, find the HAVIP that you want to manage and choose **More > Unbind with EIP** in the **Actions** column.
5. In the message that appears, click **OK**.

17.6.7. Delete an HAVIP

This topic describes how to delete a high-availability virtual IP address (HAVIP).

Prerequisites

- The HAVIP that you want to delete is not associated with an elastic IP address (EIP). If the HAVIP is associated with an EIP, disassociate the HAVIP from the EIP first. For more information, see [Disassociate HAVIPs from EIPs](#).
- The HAVIP is not associated with an Elastic Compute Service (ECS) instance. If the HAVIP is associated with an ECS instance, disassociate the HAVIP from the ECS instance. For more information, see [Disassociate an HAVIP from an ECS instance](#).
- The HAVIP is not associated with a secondary elastic network interface (ENI). If the HAVIP is associated with a secondary ENI, disassociate the HAVIP from the secondary ENI first. For more information, see [Disassociate HAVIPs from secondary ENIs](#).

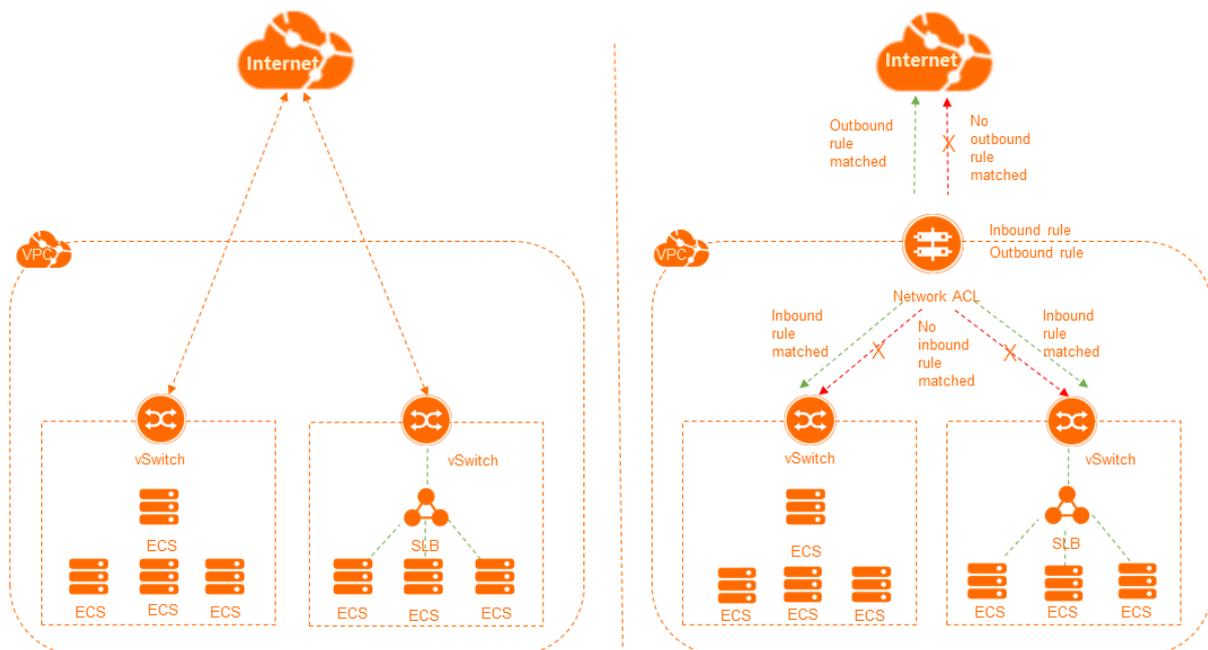
Procedure

1. Log on to the VPC console.
2. In the left-side navigation pane, click **HaVip**.
3. In the top navigation bar, select the region where the HAVIP is created.
4. On the **HaVip** page, find the HAVIP that you want to manage and choose **More > Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

17.7. Network ACLs

17.7.1. Overview

Network access control lists (ACLs) allow you to regulate access control for a virtual private cloud (VPC). You can create network ACL rules and associate a network ACL with a vSwitch. This allows you to control inbound and outbound traffic of Elastic Compute Service (ECS) instances that are associated with the vSwitch.



Features

Network ACLs have the following features:

- A network ACL is used to filter inbound and outbound network traffic of ECS instances that are associated with a vSwitch in a VPC. The network traffic forwarded to ECS instances by Server Load Balancer (SLB) instances is also filtered.
- Network ACLs are stateless. You must set both inbound and outbound rules. Otherwise, the system may fail to respond to requests.
- If you create a network ACL that does not contain a rule, all inbound and outbound access are denied.
- If a network ACL is associated with a vSwitch, the network ACL does not filter the traffic forwarded between ECS instances that are associated with the vSwitch.

Network ACL rules

You can add rules to or delete rules from a network ACL. Changes to the rules are automatically synchronized to the associated vSwitch. By default, an inbound and an outbound rule are automatically added to a newly created network ACL. These rules allow all inbound and outbound network traffic transmitted through the associated vSwitch. You can delete the default rules. The following table lists the default inbound and outbound rules.

- Default inbound rule

| Priority | Protocol | Source CIDR block | Destination port range | Action | Type |
|----------|----------|-------------------|------------------------|--------|--------|
| 1 | all | 0.0.0.0/0 | -1/-1 | Allow | Custom |

- Default outbound rule

| Priority | Protocol | Destination CIDR block | Destination port range | Action | Type |
|----------|----------|------------------------|------------------------|--------|--------|
| 1 | all | 0.0.0.0/0 | -1/-1 | Allow | Custom |

A network ACL contains the following parameters:

- Priority: A smaller value indicates a higher priority. The system attempts to match traffic requests with rules in descending order of priority that starts from the rule whose priority is 1. If a request matches a rule, the system applies the rule to the request and ignores the other rules.

For example, the following rules are added and requests destined for IP address 172.16.0.1 are sent from an ECS instance. In the following table, the requests match Rules 2 and 3. Rule 2 has a higher priority than Rule 3. Therefore, the system applies Rule 2. Based on the action of Rule 2, the requests are denied.

| Priority | Protocol | Destination CIDR block | Destination port range | Action | Type |
|----------|----------|------------------------|------------------------|--------|--------|
| 1 | all | 10.0.0.0/8 | -1/-1 | Allow | Custom |
| 2 | all | 172.16.0.0/12 | -1/-1 | Block | Custom |
| 3 | all | 172.16.0.0/12 | -1/-1 | Allow | Custom |

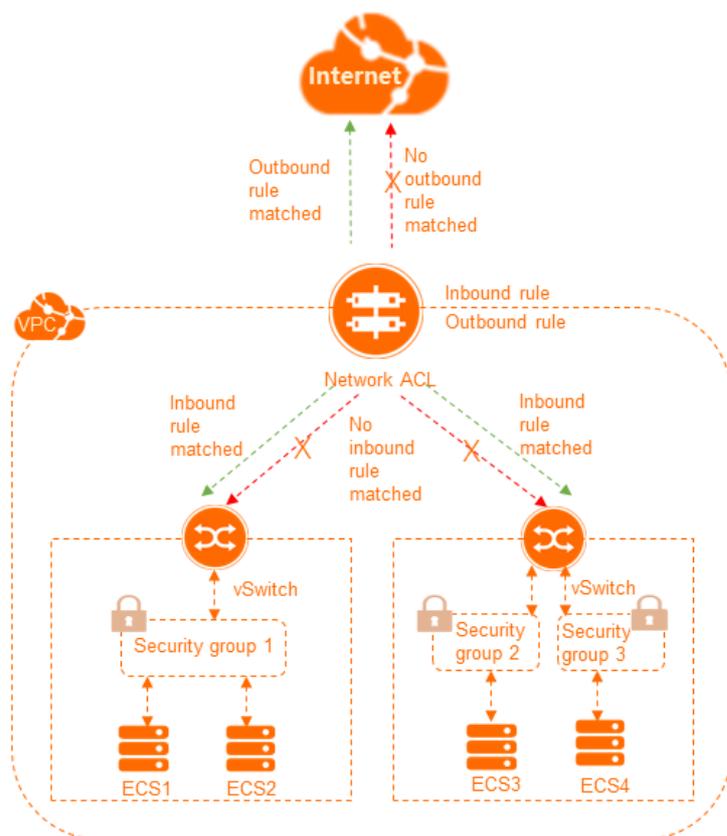
- Policy: specifies whether to allow or block specific traffic.
- Protocol: the protocol type. Available options include All, ICMP, GRE, TCP, and UDP.
- Source CIDR block: the source CIDR block from which inbound traffic is transmitted.
- Destination CIDR block: the destination CIDR block to which outbound traffic is transmitted.
- Destination port range: the range of destination ports to which the inbound rule applies.
- Destination port range: the range of destination ports to which the outbound rule applies.

Comparison between network ACLs and security groups

Network ACLs control data transmitted through associated vSwitches while security groups control data transmitted through associated ECS instances. The following table lists the differences between network ACLs and security groups.

| Network ACL | Security group |
|--|--|
| Applied to vSwitches. | Applied to instances. |
| Stateless: Returned traffic must be allowed by inbound rules. | Stateful: Returned traffic is automatically allowed and not affected by rules. |
| The system attempts to match requests with rules in descending order of priority. Not all rules are matched. | The system matches a request against all rules before a rule is applied. |
| Each vSwitch can be associated with only one network ACL. | Each ECS instance can be added to more than one security group. |

The following figure shows how network ACLs and security groups are applied to ensure network security.



Limits

Before you use network ACLs, take note of the following limits.

| Item | Default |
|---|---------|
| Number of network ACLs that can be created in each VPC | 200 |
| Number of network ACLs that can be associated with each vSwitch | 1 |

| Item | Default |
|--|---|
| Number of rules that can be added to a network ACL | <ul style="list-style-type: none"> Inbound rules: 20 Outbound rules: 20 |

Procedure

The following flowchart shows how to use a network ACL.



17.7.2. Scenarios

If you are familiar with the ports that are commonly used by ECS instances, you can specify them in access control list (ACL) rules to facilitate precise network traffic filtering. This topic describes the ports that are commonly used by ECS instances and the application scenarios of these ports.

Ports

The following table lists the ports and the services that use these ports.

| Port | Service | Description |
|------|------------|---|
| 21 | FTP | The FTP port. It is used to upload and download files. |
| 22 | SSH | The SSH port. It is used to log on to Linux instances in the command line method by using username and password pairs. |
| 23 | Telnet | The Telnet port. It is used to remotely log on to ECS instances. |
| 25 | SMTP | The SMTP port. It is used to send emails. |
| 80 | HTTP | The HTTP port. It is used to access services such as IIS, Apache, and NGINX. |
| 110 | POP3 | The POP3 port. It is used to send and receive emails. |
| 143 | IMAP | The Internet Message Access Protocol (IMAP) port. It is used to receive emails. |
| 443 | HTTPS | The HTTPS port. It is used to access services. The HTTPS protocol can implement encrypted and secure data transmission. |
| 1433 | SQL Server | The TCP port of SQL Server. It is used for SQL Server to provide external services. |
| 1434 | SQL Server | The UDP port of SQL Server. It is used to return the TCP/IP port occupied by SQL Server. |
| 1521 | Oracle | The Oracle communication port. ECS instances that run Oracle SQL must have this port open. |

| Port | Service | Description |
|------|--|--|
| 3306 | MySQL | The MySQL port. It is used for MySQL databases to provide external services. |
| 3389 | Windows Server Remote Desktop Services | The Windows Server Remote Desktop Services port. It is used to log on to a Windows instance. |
| 8080 | Proxy port | An alternative to port 80. It is commonly used for WWW proxy services. |

Custom network ACLs

Inbound rules and **Outbound rules** describe a network ACL example for VPCs that support IPv4 addresses only.

- The inbound rules in effective order 1, 2, 3, and 4 respectively allow HTTP, HTTPS, SSH, and RDP traffic to the VSwitch. Outbound response rules are those in effective order 3.
- The outbound rules in effective order 1 and 2 respectively allow HTTP and HTTPS traffic from the VSwitch. Outbound response rules are those in effective order 5.
- The inbound rule in effective order 6 denies all inbound IPv4 traffic. This rule ensures that packets that do not match any other rules are denied.
- The outbound rule in effective order 4 denies all outbound IPv4 traffic. This rule ensures that packets that do not match any other rules are denied.

 **Note** An inbound or outbound rule must correspond to an inbound or outbound rule that allows response traffic.

Inbound rules

| Effective order | Protocol | Source IP addresses | Destination port range | Action | Description |
|-----------------|----------|---------------------|------------------------|--------|--|
| 1 | TCP | 0.0.0.0/0 | 80/80 | Accept | Allows inbound HTTP traffic from any IPv4 addresses. |
| 2 | TCP | 0.0.0.0/0 | 443/443 | Accept | Allows inbound HTTPS traffic from any IPv4 addresses. |
| 3 | TCP | 0.0.0.0/0 | 22/22 | Accept | Allows inbound SSH traffic from any IPv4 addresses. |
| 4 | TCP | 0.0.0.0/0 | 3389/3389 | Accept | Allows inbound RDP traffic from any IPv4 addresses. |
| 5 | TCP | 0.0.0.0/0 | 32768/65535 | Accept | Allows inbound IPv4 traffic from the Internet. This port range is for reference only. For more information on how to select appropriate ephemeral ports, see Ephemeral ports . |
| 6 | All | 0.0.0.0/0 | -1/-1 | Drop | Denies all inbound IPv4 traffic. |

Outbound rules

| Effective order | Protocol | Destination IP addresses | Destination port range | Action | Description |
|-----------------|----------|--------------------------|------------------------|--------|---|
| 1 | TCP | 0.0.0.0/0 | 80/80 | Accept | Allows outbound IPv4 HTTP traffic from the VSwitch to the Internet. |
| 2 | TCP | 0.0.0.0/0 | 443/443 | Accept | Allows outbound IPv4 HTTPS traffic from the VSwitch to the Internet. |
| 3 | TCP | 0.0.0.0/0 | 32768/65535 | Accept | Allows outbound IPv4 traffic from the VSwitch to the Internet. This port range is for reference only. For more information on how to select appropriate ephemeral ports, see Ephemeral ports . |
| 4 | All | 0.0.0.0/0 | -1/-1 | Drop | Denies all outbound IPv4 traffic. |

Network ACLs for SLB

If the ECS instance in the VSwitch acts as the backend server of an SLB instance, you must add the following network ACL rules.

- Inbound rules

| Effective order | Protocol | Source IP addresses | Destination port range | Action | Description |
|-----------------|-----------------------|--|------------------------|--------|--|
| 1 | SLB listener protocol | Client IP addresses allowed to access the SLB instance | SLB listener port | Accept | Allows inbound traffic from specified client IP addresses. |
| 2 | Health check protocol | 100.64.0.0/10 | Health check port | Accept | Allows inbound traffic from health check IP addresses. |

- Outbound rules

| Effective order | Protocol | Destination IP addresses | Destination port range | Action | Description |
|-----------------|----------|--|------------------------|--------|---|
| 1 | All | Client IP addresses allowed to access the SLB instance | -1/-1 | Accept | Allows all outbound traffic to specified client IP addresses. |
| 2 | All | 100.64.0.0/10 | -1/-1 | Accept | Allows outbound traffic to health check IP addresses. |

Ephemeral ports

Clients use different ports to initiate requests. You can select different port ranges for network ACL rules based on the client type. The following table lists ephemeral port ranges for common clients.

| Client | Port range |
|-------------------------------|-------------|
| Linux | 32768/61000 |
| Windows Server 2003 | 1025/5000 |
| Windows Server 2008 and later | 49152/65535 |
| NAT gateway | 1024/65535 |

17.7.3. Create a network ACL

A network access control list (ACL) allows you to manage network access in a virtual private cloud (VPC). This topic describes how to create a network ACL in a VPC.

Prerequisites

A VPC is created. For more information, see [Create a VPC](#).

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **Network ACL**.
3. In the top navigation bar, select the region where you want to create the network ACL.
4. On the **Network ACL** page, click **Create Network ACL**.
5. In the **Create Network ACL** dialog box, set the following parameters, and click **OK**.

| Parameter | Description |
|---------------------|---|
| Organization | Select the organization to which the network ACL belongs. |
| Resource Set | Select the resource set to which the network ACL belongs. |
| Region | Select the region where you want to deploy the network ACL. |
| Name | Enter a name for the network ACL. The name must be 2 to 128 characters in length, and can contain digits, underscores (_), and hyphens (-). It must start with a letter. |
| Description | Enter a description for the network ACL. The description must be 2 to 256 characters in length, and can contain letters, digits, underscores (_), colons (:), and hyphens (-). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . |
| VPC | Select the VPC for which you want to create the network ACL.  Note The VPC and network ACL must be deployed in the same region. |

What's next

- [Associate a network ACL with a vSwitch](#)
- [Add an inbound rule](#)
- [Add an outbound rule](#)

17.7.4. Associate a network ACL with a vSwitch

After you create a network access control list (ACL), you can associate the network ACL with a vSwitch. This way, you can use the network ACL to manage the traffic of the Elastic Compute Service (ECS) instances in the vSwitch.

Prerequisites

Before you associate a network ACL with a vSwitch, make sure that the following requirements are met:

- A network ACL is created. For more information, see [Work with network ACLs](#).
- A vSwitch is created. The vSwitch and network ACL belong to the same virtual private cloud (VPC). For more information, see [Work with vSwitches](#).

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **Network ACL**.
3. In the top navigation bar, select the region where the network ACL is created.
4. On the **Network ACL** page, find the network ACL that you want to manage and click **Manage** in the **Actions** column.
5. On the **Resources** tab, click **Bind Resource**.
6. In the **Associate vSwitch** panel, select the vSwitch and click **OK**.

 **Note** The network ACL and vSwitch must belong to the same VPC. A vSwitch can be associated with only one network ACL.

What's next

- [Add an inbound rule](#)
- [Add an outbound rule](#)

17.7.5. Add network ACL rules

17.7.5.1. Add an inbound rule

This topic describes how to add an inbound rule to a network access control list (ACL). You can use inbound rules to control Internet or internal network traffic destined for Elastic Compute Service (ECS) instances connected to a vSwitch.

Prerequisites

A network ACL is created. For more information, see [Work with network ACLs](#).

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **Network ACL**.
3. In the top navigation bar, select the region where the network ACL is created.
4. On the **Network ACL** page, find the network ACL that you want to manage and click **Inbound Rule** in the **Actions** column.

5. On the **Inbound Rule** tab, click **Create Inbound Rule**.
6. In the **Create Inbound Rule** panel, set the following parameters and click **OK**.

| Parameter | Description |
|-------------------------------|--|
| Name | Enter a name for the inbound rule. The name must be 2 to 128 characters in length, and can contain letters, digits, underscores (_), and hyphens (-). The name must start with a letter. It cannot start with <code>http://</code> or <code>https://</code> . |
| Effective order | The order in which the inbound rule takes effect. Valid values: 1 to 20. A smaller number indicates a higher priority. |
| Action | Select an action for the inbound rule. Valid values: <ul style="list-style-type: none"> ◦ Accept: accepts network traffic that is destined for the ECS instances connected to the vSwitch. ◦ Drop: denies network traffic that is destined for the ECS instances connected to the vSwitch. |
| Protocol Type | Select a transport layer protocol. Valid values: <ul style="list-style-type: none"> ◦ ALL ◦ ICMP ◦ GRE ◦ TCP ◦ UDP |
| Source IP Addresses | The source CIDR block from which data is transmitted. Default value: 0.0.0.0/32. |
| Destination Port Range | Enter the destination port range. Valid values: 1 to 65535. Use a forward slash (/) to separate the highest and lowest values in a port range, for example, 1/200 or 80/80. -1/-1 indicates that all ports are available. You cannot set the value to -1/-1. |

17.7.5.2. Add an outbound rule

This topic describes how to add an outbound rule to a network access control list (ACL). You can use outbound rules to control how Elastic Compute Service (ECS) instances connected to a vSwitch access the Internet or other internal networks.

Prerequisites

A network ACL is created. For more information, see [Work with network ACLs](#).

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **Network ACL**.
3. In the top navigation bar, select the region where the network ACL is created.
4. On the **Network ACL** page, find the network ACL and click **Outbound Rule** in the **Actions** column.

5. On the **Outbound Rule** tab, click **Create Outbound Rule**.
6. In the **Create Outbound Rule** panel, set the following parameters and click **OK**.

| Parameter | Description |
|---------------------------------|---|
| Name | Enter a name for the outbound rule. The name must be 2 to 128 characters in length, and can contain letters, digits, underscores (_), and hyphens (-). The name must start with a letter. It cannot start with <code>http://</code> or <code>https://</code> . |
| Effective order | The order in which the outbound rule takes effect. Valid values: 1 to 20. A smaller number indicates a higher priority. |
| Policy | Select an action for the outbound rule. Valid values: <ul style="list-style-type: none"> ◦ Accept: allows ECS instances connected to the vSwitch to access the Internet or other internal networks. ◦ Drop: forbids ECS instances connected to the vSwitch to access the Internet or other internal networks. |
| Protocol Type | Select a transport layer protocol. Valid values: <ul style="list-style-type: none"> ◦ ALL ◦ ICMP ◦ GRE ◦ TCP ◦ UDP |
| Destination IP Addresses | The destination CIDR block to which data is transmitted. Default value: 0.0.0.0/32. |
| Destination Port Range | Enter the destination port range. Valid values: 1 to 65535. Use a forward slash (/) to separate the highest and lowest values in a port range, for example, 1/200 or 80/80. -1/-1 indicates that all ports are available. You cannot set the value to -1/-1. |

17.7.5.3. Change the priority of a network ACL rule

Rules added to network access control lists (ACLs) take effect in descending order of priority. A smaller value indicates a higher priority. You can change the priority of a network ACL rule to meet your business requirements.

Change the priority of an inbound rule

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **Network ACL**.
3. In the top navigation bar, select the region where the network ACL is created.
4. On the **Network ACL** page, find the network ACL that you want to manage and click **Manage** in the **Actions** column.
5. Click the **Inbound Rule** tab and click **Sort**.
6. In the **Sort** panel, change the priority of the rule by dragging and dropping the rule. Then, click **OK**.

 **Note** Rules are listed in descending order of priority.

Change the priority of an outbound rule

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **Network ACL**.
3. In the top navigation bar, select the region where the network ACL is created.
4. On the **Network ACL** page, find the network ACL and click **Manage** in the **Actions** column.
5. Click the **Outbound Rule** tab and click **Sort**.
6. In the **Sort** panel, change the priority of the rule by dragging and dropping the rule. Then, click **OK**.

 **Note** Rules are listed in descending order of priority.

17.7.6. Disassociate a network ACL from a vSwitch

This topic describes how to disassociate a network access control list (ACL) from a vSwitch. After you disassociate a network ACL from a vSwitch, the network ACL no longer controls the traffic of Elastic Compute Service (ECS) instances that belong to the vSwitch.

Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **Network ACL**.
3. In the top navigation bar, select the region where the network ACL is created.
4. On the **Network ACL** page, find the network ACL and click **Manage** in the **Actions** column.
5. On the **Resources** tab, find the vSwitch and click **Unbind** in the **Actions** column.
6. In the **Unbind Network ACL** message, click **OK**.

17.7.7. Delete a network ACL

This topic describes how to delete a network access control list (ACL).

Prerequisites

Make sure that the network ACL is not associated with a vSwitch. If the network ACL is associated with a vSwitch, disassociate the network ACL from the vSwitch first. For more information, see [Disassociate a VSwitch from a network ACL](#).

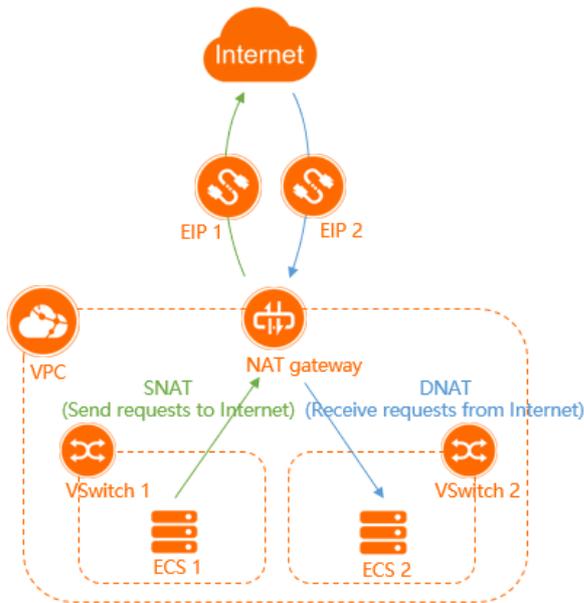
Procedure

1. [Log on to the VPC console](#).
2. In the left-side navigation pane, click **Network ACL**.
3. In the top navigation bar, select the region where the network ACL is created.
4. On the **Network ACL** page, find the network ACL that you want to delete and click **Delete** in the **Actions** column.
5. In the **Delete Network ACL** message, click **OK**.

18.NAT Gateway

18.1. What is NAT Gateway?

NAT gateways are enterprise-class gateways that provide the Source Network Address Translation (SNAT) and Destination Network Address Translation (DNAT) features. Each NAT gateway provides a throughput capacity of up to 100 Gbit/s. NAT gateways also support cross-zone disaster recovery.



Features

You must associate public IP addresses with NAT gateways so that the NAT gateways can function as expected. After you create a NAT gateway, you can associate elastic IP addresses (EIPs) with the NAT gateway.

NAT gateways provide the SNAT and DNAT features.

| Feature | Description |
|---------|---|
| SNAT | SNAT allows ECS instances that are deployed in a VPC to access the Internet when no public IP address is associated with these ECS instances. |
| DNAT | DNAT maps the EIPs of a NAT gateway to ECS instances. This way, the ECS instances can receive requests from the Internet. |

18.2. Log on to the NAT Gateway console

This topic describes how to log on to the Apsara Uni-manager Management Console to manage your NAT gateways. The Google Chrome browser is used as an example.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

Note When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- o Uppercase or lowercase letters
- o Digits
- o Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top menu bar, choose **Products > Networking > Virtual Private Cloud**.
5. In the left-side navigation pane, choose **Internet Access > NAT Gateway**.

18.3. Quick Start

18.3.1. Overview

This topic describes how to configure Source Network Address Translation (SNAT) and Destination Network Address Translation (DNAT). You can configure SNAT and DNAT to enable ECS instances in a Virtual Private Cloud (VPC) network to communicate with the Internet through a NAT gateway.

Prerequisites

Before you start, make sure that the following conditions are met:

- A VPC network is created. For more information, see *VPC User Guide Create a VPC network* in the **Quick Start** chapter in the VPC User Guide.
- An ECS instance is created in the VPC network. For more information, see *ECS User Guide Create an instance* in the **Quick Start** chapter in the ECS User Guide.
- An elastic IP address (EIP) is created. For more information, see *EIP User Guide Create an EIP* in the **Quick Start** chapter in the EIP User Guide.

Procedure

In this topic, an ECS instance that is not associated with any public IP addresses in a VPC network is used as an example. The following flowchart shows how to associate an EIP with a NAT gateway:



1 Create a NAT Gateway **2** Associate an EIP **3** Create a DNAT entry **4** Create an SNAT entry

- Region
- VPC
- Specification
- Select an EIP
- Public IP address
- Private IP address
- Port settings
- VSwitch granularity
- ECS granularity

1. Create a NAT gateway

A NAT gateway is an enterprise-class gateway that provides NAT proxy services. Before you configure SNAT and DNAT entries, you must create a NAT gateway.

For more information, see [Create a NAT gateway](#).

2. Associate an EIP to a NAT gateway

A NAT gateway functions as expected only after it is associated with a public IP address. After you create a NAT gateway, you can associate it with an EIP.

For more information, see [Associate an EIP with a NAT Gateway](#).

3. Create a DNAT entry

This topic describes how to create a Destination Network Address Translation (DNAT) entry. DNAT maps public IP addresses to Elastic Compute Service (ECS) instances in a Virtual Private Cloud (VPC) network. This way, the ECS instances can receive requests sent over the Internet. DNAT supports port mapping and IP mapping.

For more information, see [Create a DNAT entry](#).

4. Create an SNAT entry

This topic describes how to create a Source Network Address Translation (SNAT) entry. SNAT allows Elastic Compute Service (ECS) instances in a Virtual Private Cloud (VPC) network to access the Internet without using public IP addresses.

For more information, see [Create a SNAT entry](#).

18.3.2. Create a NAT gateway

A NAT gateway is an enterprise-class gateway that provides NAT proxy services. Before you configure SNAT and DNAT entries, you must create a NAT gateway.

Prerequisites

A VPC network is created. For more information, see [VPC User Guide](#) **Create a VPC network** in the **Quick Start** chapter in the VPC User Guide.

Procedure

1. [Log on to the NAT Gateway console](#).
2. On the **NAT Gateway** page, click **Create NAT Gateway**.
3. Set the parameters for the NAT gateway based on the following information, and then click **Submit**.

| Parameter | Description |
|---------------------|---|
| Organization | The organization to which the NAT gateway belongs. |
| Resource set | The resource set to which the NAT gateway belongs. |
| Region | The region where the NAT gateway is deployed. |
| VPC | <p>The VPC network to which the NAT gateway belongs.</p> <p>If you cannot find the target VPC network in the list, perform the following operations:</p> <ul style="list-style-type: none"> ◦ Check whether the VPC network is already associated with a NAT gateway. Each VPC network can be associated with only one NAT gateway. ◦ Check whether the VPC network has a custom route entry with the destination CIDR block set to 0.0.0.0/0. If such a custom route entry exists, delete it. ◦ Check whether the RAM user is authorized to access the VPC network. If the RAM user is not authorized, contact your Alibaba Cloud account owner to grant permissions. |

| Parameter | Description |
|---------------|---|
| Specification | <p>Select the size of the NAT gateway. Valid values:</p> <ul style="list-style-type: none"> ◦ Small: supports up to 10,000 SNAT connections. ◦ Medium: supports up to 50,000 SNAT connections. ◦ Large: supports up to 200,000 SNAT connections. ◦ Super Large: supports up to 1,000,000 SNAT connections. <div style="background-color: #e0f2f7; padding: 5px; border: 1px solid #ccc;"> <p>Note The size of a NAT gateway determines the maximum number of SNAT connections, but it does not affect the maximum number of DNAT connections.</p> </div> |
| Parameter | <p>Enter a name for the NAT gateway.</p> <p>The name must be 2 to 128 characters in length and can contain letters, digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (.). It must start with a letter or Chinese character but cannot start with <code>http://</code> or <code>https://</code>.</p> |

18.3.3. Associate an EIP with a NAT gateway

A NAT gateway works as expected only after you associate an elastic IP address (EIP) with it. This topic describes how to associate an EIP with a NAT gateway.

Procedure

1. [Log on to the NAT Gateway console.](#)
2. In the top navigation bar, select the region where the NAT gateway is deployed.
3. On the **NAT Gateway** page, find the NAT gateway with which you want to associate an EIP, and choose **⋮ >** **Bind Elastic IP Address** in the **Actions** column.
4. In the **Associate EIP** dialog box, set the following parameters, and click **OK**.

| Parameter | Description |
|-----------------|--|
| Usable EIP List | Select the EIP that is used to communicate with the Internet. |
| vSwitch | <p>Select the vSwitch to which you want to add SNAT entries.</p> <p>After you select a vSwitch, the system automatically adds SNAT entries to the vSwitch. Then, cloud services in the vSwitch can access the Internet. You can skip this step and manually add SNAT entries after you associate an EIP with the NAT gateway. For more information, see Create a SNAT entry.</p> |

18.3.4. Create a DNAT entry

This topic describes how to create a DNAT entry. DNAT can map public IP addresses of NAT gateways to Elastic Compute Service (ECS) instances. This way, ECS instances can provide services over the Internet. DNAT supports port mapping and IP mapping.

Procedure

1. [Log on to the NAT Gateway console.](#)
2. In the top navigation bar, select the region where the NAT gateway is deployed.

3. On the **NAT Gateway** page, find the NAT gateway that you want to manage, and click **Configure DNAT** in the **Actions** column.
4. On the **DNAT Table** page, click **Create DNAT Entry**.
5. In the **Create DNAT Entry** panel, set the following parameters and click **OK**.

| Parameter | Description |
|---------------------------|---|
| Public IP Address | <p>Select an available public IP address.</p> <p> Note If a public IP address is already used to create an SNAT entry, the public IP address cannot be used to create a DNAT entry.</p> |
| Private IP Address | <p>Specify the ECS instance that uses the DNAT entry to communicate with the Internet. You can specify the private IP address of the ECS instance in the following ways:</p> <ul style="list-style-type: none"> ◦ Auto Fill: Select the ECS instance from the drop-down list. ◦ Manually Input: Enter the private IP address of the ECS instance. <p> Note This private IP address must fall within the CIDR block of the virtual private cloud (VPC). You can also enter the private IP address of an existing ECS instance.</p> |
| Port Settings | <p>Select a DNAT mapping method:</p> <ul style="list-style-type: none"> ◦ All: This method uses IP mapping. All requests destined for the elastic IP address (EIP) are forwarded to the ECS instance. ◦ Specific Port: This method uses port mapping. The NAT gateway forwards requests that use the specified protocol and port to the specified port of the ECS instance. <p>After you select Specific Port, specify the Public Port (the external port), Private Port (the internal port), and IP Protocol (the protocol over which data is transferred).</p> |
| Entry Name | <p>Enter a name for the DNAT entry.</p> <p>The name must be 2 to 128 characters in length, and can contain digits, underscores (_), and hyphens (-). It must start with a letter.</p> |

18.3.5. Create an SNAT entry

This topic describes how to create an SNAT entry. SNAT can provide the proxy service for Elastic Compute Service (ECS) instances. ECS instances that have no IP address assigned in virtual private clouds (VPCs) can access the Internet by using SNAT.

Procedure

1. [Log on to the NAT Gateway console](#).
2. In the top navigation bar, select the region where the NAT gateway is deployed.
3. On the **NAT Gateway** page, find the NAT gateway and click **Configure SNAT** in the **Actions** column.
4. On the **SNAT Table** page, click **Create SNAT Entry**.
5. In the **Create SNAT Entry** panel, set the following parameters and click **OK**.

| Parameter | Description |
|--------------------------------|--|
| VSwitch Granularity | |
| VSwitch | <p>Select a vSwitch in the VPC. All ECS instances in the vSwitch can access the Internet by using SNAT.</p> <p>Note SNAT entries do not take effect on ECS instances that are assigned public IP addresses. For example, an ECS instance may be assigned a static public IP address, associated with an elastic IP address (EIP), or configured with DNAT IP mapping. Such an ECS instance uses the public IP address instead of the SNAT entry to access the Internet.</p> |
| VSwitch CIDR Block | The CIDR block of the selected vSwitch. |
| Public IP Address | <p>Select the EIP that is used to access the Internet.</p> <p>You can select one or more EIPs. You can use multiple EIPs to create an SNAT IP address pool.</p> <p>Note An EIP that is already used in a DNAT entry cannot be used in an SNAT entry.</p> |
| ECS Granularity | |
| Available ECS Instances | <p>Select an ECS instance in the VPC.</p> <p>The ECS instance can access the Internet by using the specified EIP. Make sure that the following requirements are met:</p> <ul style="list-style-type: none"> The ECS instance is in the Running state. The ECS instance is not assigned an EIP or a static public IP address. |
| ECS CIDR Block | The CIDR block of the ECS instance. |
| Public IP Address | <p>Select the EIP that is used to access the Internet.</p> <p>You can select one or more EIPs. You can use multiple EIPs to create an SNAT IP address pool.</p> <p>The maximum bandwidth supported by each EIP in an SNAT IP address pool is 200 Mbit/s. To make full use of your EIP bandwidth plan and avoid port conflicts caused by insufficient EIPs, we recommend that you add EIPs to the SNAT address pool based on the following rules:</p> <ul style="list-style-type: none"> If the maximum bandwidth value of the EIP bandwidth plan is 1,024 Mbit/s, specify at least five EIPs in each SNAT entry. If the maximum bandwidth value of the EIP bandwidth plan exceeds 1,024 Mbit/s, specify at least one additional EIP in each SNAT entry for every incremental 200 Mbit/s. <p>Note An EIP that is already used in a DNAT entry cannot be used in an SNAT entry.</p> |

18.4. Manage a NAT gateway

18.4.1. Overview

NAT gateways are enterprise-class gateways that support SNAT and DNAT features. Each NAT gateway provides a forwarding capacity of 10 Gbit/s. NAT gateways support cross-zone disaster recovery.

Sizes of NAT gateways

NAT gateways are available in multiple sizes, including small, middle, large, and super large-1. The size of a NAT gateway determines the SNAT performance, which includes the maximum number of connections and the number of new connections per second. However, the size of a NAT gateway does not affect the DNAT performance. The following table describes available sizes of NAT gateways.

| Size | Maximum number of SNAT connections | Number of new SNAT connections per second |
|---------------|------------------------------------|---|
| Small | 10,000 | 1,000 |
| Medium | 50,000 | 5,000 |
| Large | 200,000 | 10,000 |
| Super Large-1 | 1,000,000 | 50,000 |

When you specify the size of a NAT gateway, take note of the following limits:

- CloudMonitor monitors only the maximum number of SNAT connections for NAT gateways. CloudMonitor does not monitor the number of new SNAT connections per second.
- The timeout period of SNAT connections on NAT gateways is 900 seconds.
- To avoid timeouts of SNAT connections caused by network congestion or Internet jitter, make sure that your applications support automatic reconnection.
- NAT gateways do not support packet fragmentation.
- If you use the same destination public IP address and port, the maximum number of concurrent connections is based on the number of elastic IP addresses (EIPs) that are associated with the NAT gateway. Each EIP that is associated with the NAT gateway supports up to 55,000 concurrent connections. If N EIPs are associated with the NAT gateway, the maximum number of concurrent connections that the NAT gateway supports is $N \times 55,000$.

18.4.2. Create a NAT gateway

A NAT gateway is an enterprise-class gateway that provides NAT proxy services. Before you configure SNAT and DNAT entries, you must create a NAT gateway.

Prerequisites

A VPC network is created. For more information, see *VPC User Guide* **Create a VPC network** in the **Quick Start** chapter in the VPC User Guide.

Procedure

1. [Log on to the NAT Gateway console](#).
2. On the **NAT Gateway** page, click **Create NAT Gateway**.
3. Set the parameters for the NAT gateway based on the following information, and then click **Submit**.

| Parameter | Description |
|---------------------|--|
| Organization | The organization to which the NAT gateway belongs. |

| Parameter | Description |
|---------------|--|
| Resource set | The resource set to which the NAT gateway belongs. |
| Region | The region where the NAT gateway is deployed. |
| VPC | <p>The VPC network to which the NAT gateway belongs.</p> <p>If you cannot find the target VPC network in the list, perform the following operations:</p> <ul style="list-style-type: none"> Check whether the VPC network is already associated with a NAT gateway. Each VPC network can be associated with only one NAT gateway. Check whether the VPC network has a custom route entry with the destination CIDR block set to 0.0.0.0/0. If such a custom route entry exists, delete it. Check whether the RAM user is authorized to access the VPC network. If the RAM user is not authorized, contact your Alibaba Cloud account owner to grant permissions. |
| Specification | <p>Select the size of the NAT gateway. Valid values:</p> <ul style="list-style-type: none"> Small: supports up to 10,000 SNAT connections. Medium: supports up to 50,000 SNAT connections. Large: supports up to 200,000 SNAT connections. Super Large: supports up to 1,000,000 SNAT connections. <div style="background-color: #e6f2ff; padding: 5px; border: 1px solid #d9e1f2;"> <p> Note The size of a NAT gateway determines the maximum number of SNAT connections, but it does not affect the maximum number of DNAT connections.</p> </div> |
| Parameter | <p>Enter a name for the NAT gateway.</p> <p>The name must be 2 to 128 characters in length and can contain letters, digits, underscores (_), hyphens (-), periods (.), colons (:), and commas (,). It must start with a letter or Chinese character but cannot start with <code>http://</code> or <code>https://</code>.</p> |

18.4.3. Modify a NAT gateway

This topic describes how to modify the name and description of a NAT gateway.

Procedure

- Log on to the NAT Gateway console.
- In the top navigation bar, select the region where the NAT gateway is deployed.
- On the NAT Gateway page, find the NAT gateway that you want to manage, and then click **Manage** in the **Actions** column.
- On the NAT Gateway Details tab, click **Edit** next to the name. In the dialog box that appears, enter a new name for the NAT gateway, and then click **OK**.
The name must be 2 to 128 characters in length, and can contain digits, underscores (_), and hyphens (-). It must start with a letter or Chinese character.
- Click **Edit** next to the description. In the dialog box that appears, enter a new description, and then click **OK**.
The description must be 2 to 256 characters in length. It cannot start with `http://` or `https://`.

18.4.4. Delete a NAT gateway

You can delete NAT gateways that are billed on a pay-as-you-go basis. You cannot delete subscription NAT gateways.

Prerequisites

Before you delete a NAT gateway, make sure that the following conditions are met:

- The NAT gateway is not associated with an EIP. If the NAT gateway is associated with an EIP, disassociate the EIP first. For more information, see [Disassociate EIPs from a NAT gateway](#).
- The DNAT table is empty. If the DNAT table contains DNAT entries, delete these entries first. For more information, see [Delete a DNAT entry](#).
- The SNAT table is empty. If the DNAT table contains SNAT entries, delete these entries first. For more information, see [Delete a SNAT entry](#).

Procedure

1. [Log on to the NAT Gateway console](#).
2. In the top navigation bar, select the region where the NAT gateway is deployed.
3. On the **NAT Gateways** page, find the target NAT gateway, and click  > **Delete** in the **Actions** column.
4. In the dialog box that appears, click **OK**.

 **Note** If you select **Delete (Delete NAT gateway and resources)**, the DNAT and SNAT entries of the NAT gateway are deleted automatically. The EIP associated with the NAT gateway is also disassociated.

18.5. Manage EIPs

18.5.1. Associate an EIP with a NAT gateway

A NAT gateway works as expected only after it is associated with an elastic IP address (EIP). This topic describes how to associate an EIP with a NAT gateway.

Prerequisites

Before you associate an EIP with a NAT gateway, make sure that the following requirements are met:

- A NAT gateway is created. For more information, see [Create a NAT gateway](#).
- An EIP is purchased. For more information, see the *Create an EIP* topic in **Quick Start** of **Elastic IP Address User Guide**.

Procedure

1. [Log on to the NAT Gateway console](#).
2. In the top navigation bar, select the region where the NAT gateway is deployed.
3. On the **NAT Gateway** page, find the NAT gateway with which you want to associate an EIP, and choose  > **Bind Elastic IP Address** in the **Actions** column.
4. In the **Associate EIP** dialog box, set the following parameters, and click **OK**.

| Parameter | Description |
|-----------------|---|
| Usable EIP list | Select the EIP that is used to communicate with the Internet. |

| Parameter | Description |
|-----------|--|
| vSwitch | Select the vSwitch to which you want to add SNAT entries. After you select a vSwitch, the system automatically adds SNAT entries to the vSwitch. Then, cloud services in the vSwitch can access the Internet. You can skip this step and manually add SNAT entries after you associate an EIP with the NAT gateway. |

18.5.2. Disassociate an EIP from a NAT gateway

This topic describes how to disassociate an elastic IP address (EIP) from a NAT gateway. After an EIP is disassociated from a NAT gateway, the NAT gateway can no longer communicate with the Internet by using the EIP.

Prerequisites

Make sure that the EIP to be disassociated is not used in an SNAT entry or a DNAT entry. If the EIP is used in an SNAT or a DNAT entry, delete the SNAT or DNAT entry first. For more information, see [Delete a SNAT entry](#) and [Delete a DNAT entry](#).

Procedure

1. [Log on to the NAT Gateway console](#).
2. In the top navigation bar, select the region where the NAT gateway is deployed.
3. On the **NAT Gateway** page, find the NAT gateway, and choose  > **Unbind Elastic IP Address** in the **Actions** column.
4. In the **Unbind Elastic IP Address** panel, select the EIP that you want to disassociate, and click **OK**.

18.6. Manage a DNAT table

18.6.1. DNAT overview

NAT Gateway supports the Destination Network Address Translation (DNAT) feature. You can create DNAT entries to map public IP addresses to ECS instances in a Virtual Private Cloud (VPC) network. This way, the ECS instances can receive requests from the Internet.

DNAT entries

You can configure port mapping when you create a DNAT entry. After the DNAT entry is created, requests destined for the specified public IP address are forwarded to the ECS instances within a VPC network based on the port mapping rule.

Each DNAT entry consists of the following elements:

- **Public IP** : the Elastic IP Address (EIP) bound to the NAT gateway.
- **Private IP** : the private IP address of the ECS instance in the Virtual Private Cloud.
- **Public port** : the external port for port forwarding.
- **Private Network Port** : the internal port for port forwarding.
- **Protocol type** : the protocol type of the forwarding port.

Port mapping and IP mapping

The DNAT feature supports port mapping and IP mapping:

- Port Mapping

After port mapping is configured, a NAT gateway forwards requests destined for a public IP address to the specified ECS instance based on the specified protocol and ports. The following DNAT entries are used as examples:

- Entry 1: The NAT gateway forwards requests destined for TCP port 80 of ECS instance 1.1.XX.XX to TCP port 80 of ECS instance 192.168.1.1.
- Entry 2: The NAT gateway forwards requests destined for UDP port 8080 of ECS instance 2.2.XX.XX to UDP port 8000 of ECS instance 192.168.1.2.

| DNAT entry | Public IP address | External port | Private IP address | Internal port | TLS version |
|------------|-------------------|---------------|--------------------|---------------|-------------|
| Entry 1 | 1.1.XX.XX | 80 | 192.168.1.1 | 80 | TCP |
| Entry 2 | 2.2.XX.XX | 8080 | 192.168.1.2 | 8000 | UDP |

- IP mapping

After IP mapping is configured, a NAT gateway forwards all requests destined for a public IP address to the specified ECS instance. The following entry is used as an example:

Entry 3: The NAT gateway forwards requests destined for ECS instance 3.3.XX.XX to ECS instance 192.168.1.3.

| DNAT entry | Public IP address | External port | Private IP address | Internal port | TLS version |
|------------|-------------------|---------------|--------------------|---------------|-------------|
| Entry 3 | 3.3.XX.XX | Any | 192.168.1.3 | Any | Any |

18.6.2. Create a DNAT entry

This topic describes how to create a Destination Network Address Translation (DNAT) entry. DNAT maps a public IP address to an Elastic Compute Service (ECS) instance in a Virtual Private Cloud (VPC) network. This allows the ECS instance to receive requests sent over the Internet. DNAT supports port mapping and IP mapping.

Prerequisites

A NAT gateway is created and associated with an elastic IP address (EIP). For more information, see [Create a NAT gateway](#) and [Associate an EIP with a NAT Gateway](#).

Procedure

1. [Log on to the NAT Gateway console](#).
2. In the top navigation bar, select the region where the NAT gateway is deployed.
3. On the **NAT Gateway** page, find the NAT gateway that you want to manage, and click **Configure DNAT** in the **Actions** column.
4. On the **DNAT Table** page, click **Create DNAT Entry**.
5. In the **Create DNAT Entry** panel, set the following parameters and click **OK**.

| Parameter | Description |
|-------------------|---|
| Public IP Address | Select an available public IP address. <div style="background-color: #e6f2ff; padding: 5px;"> ? Note If a public IP address is already used to create an SNAT entry, the public IP address cannot be used to create a DNAT entry. </div> |

| Parameter | Description |
|--------------------|---|
| Private IP Address | <p>Specify the ECS instance that uses the DNAT entry to communicate with the Internet. You can specify the private IP address of the ECS instance in the following ways:</p> <ul style="list-style-type: none"> ◦ Auto Fill: Select the ECS instance from the drop-down list. ◦ Manually Input: Enter the private IP address of the ECS instance. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note This private IP address must fall within the CIDR block of the virtual private cloud (VPC). You can also enter the private IP address of an existing ECS instance.</p> </div> |
| Port Settings | <p>Select a DNAT mapping method:</p> <ul style="list-style-type: none"> ◦ All: This method uses IP mapping. All requests destined for the elastic IP address (EIP) are forwarded to the ECS instance. ◦ Specific Port: This method uses port mapping. The NAT gateway forwards requests that use the specified protocol and port to the specified port of the ECS instance. <p>After you select Specific Port, specify the Public Port (the external port), Private Port (the internal port), and IP Protocol (the protocol over which data is transferred).</p> |
| Entry Name | <p>Enter a name for the DNAT entry.</p> <p>The name must be 2 to 128 characters in length, and can contain digits, underscores (_), and hyphens (-). It must start with a letter.</p> |

18.6.3. Modify a DNAT entry

This topic describes how to modify the public IP address, private IP address, or port of a DNAT entry.

Procedure

1. Log on to the NAT Gateway console.
2. In the top navigation bar, select the region where the NAT gateway is deployed.
3. On the NAT Gateway page, find the NAT gateway that you want to manage, and click **Configure DNAT** in the **Actions** column.
4. On the DNAT Table page, find the DNAT entry, and click **Edit** in the **Actions** column.
5. In the Edit DNAT Entry panel, modify the public IP address, private IP address, or port, and then click **OK**.

18.6.4. Delete a DNAT entry

This topic describes how to delete a Destination Network Address Translation (DNAT) entry. If you no longer need an Elastic Compute Service (ECS) instance to receive requests sent over the Internet, you can delete the DNAT entry of the ECS instance.

Procedure

1. Log on to the NAT Gateway console.
2. In the top navigation bar, select the region where the NAT gateway is deployed.
3. On the NAT Gateway page, find the NAT gateway that you want to manage, and click **Configure DNAT** in the **Actions** column.

4. On the **DNAT table** page, find the target DNAT entry, and click **Remove** in the **Actions** column.
5. In the dialog box that appears, click **OK**.

18.7. Manage an SNAT table

18.7.1. SNAT table overview

NAT Gateway supports Source Network Address Translation (SNAT). SNAT allows Elastic Compute Service (ECS) instances in a Virtual Private Cloud (VPC) network to access the Internet without using public IP addresses.

SNAT entries

You can create SNAT entries in an SNAT table to allow ECS instances to access the Internet.

An SNAT entry consists of the following elements:

- **VSwitch or ECS instance:** the VSwitch or ECS instance that requires the SNAT proxy service.
- **Public IP address:** the public IP address used to access the Internet.

VSwitch granularity and ECS granularity

SNAT entries can be created based on the following granularity to enable ECS instances in a VPC network to access the Internet.

- VSwitch granularity

You can select the VSwitch granularity to create an SNAT entry. The NAT gateway provides proxy service for an ECS instance attached to the specified VSwitch by using a specified public IP address when the instance sends requests to the Internet. By default, all ECS instances attached to the VSwitch can use the specified public IP address to access the Internet.

 **Note** SNAT entries do not take effect on ECS instances that are assigned public IP addresses. For example, an ECS instance may be assigned static public IP address, associated with an elastic IP address (EIP) or has a Destination Network Address Translation (DNAT) IP mapping configured. These ECS instances use the public IP addresses instead of the SNAT entries to access the Internet.

- ECS granularity

If you select the ECS granularity to create an SNAT entry, the specified ECS instance uses the specified public IP address to access the Internet. The NAT gateway provides proxy service (SNAT) for a specified ECS instance by using a specified public IP address when the instance sends requests to the Internet.

18.7.2. Create an SNAT entry

This topic describes how to create a Source Network Address Translation (SNAT) entry. SNAT allows Elastic Compute Service (ECS) instances in a Virtual Private Cloud (VPC) network to access the Internet without using public IP addresses.

Prerequisites

Before you create an SNAT entry, make sure that the following requirements are met:

- A NAT gateway is created and associated with an elastic IP address (EIP). For more information, see [Create a NAT gateway](#) and [Associate an EIP with a NAT Gateway](#).
- To create an SNAT entry with VSwitch granularity, make sure that the VSwitch is created and associated with the NAT gateway in a VPC network.
- To create an SNAT entry with ECS granularity, make sure that the ECS instance is created and associated with the NAT gateway in a VPC network.

Procedure

1. Log on to the NAT Gateway console.
2. In the top navigation bar, select the region where the NAT gateway is deployed.
3. On the NAT Gateway page, find the NAT gateway and click **Configure SNAT** in the **Actions** column.
4. On the **SNAT Table** page, click **Create SNAT Entry**.
5. In the **Create SNAT Entry** panel, set the following parameters and click **OK**.

| Parameter | Description |
|--------------------------------|--|
| VSwitch Granularity | |
| VSwitch | <p>Select a vSwitch in the VPC. All ECS instances in the vSwitch can access the Internet by using SNAT.</p> <p>Note SNAT entries do not take effect on ECS instances that are assigned public IP addresses. For example, an ECS instance may be assigned a static public IP address, associated with an elastic IP address (EIP), or configured with DNAT IP mapping. Such an ECS instance uses the public IP address instead of the SNAT entry to access the Internet.</p> |
| VSwitch CIDR Block | The CIDR block of the selected vSwitch. |
| Public IP Address | <p>Select the EIP that is used to access the Internet.</p> <p>You can select one or more EIPs. You can use multiple EIPs to create an SNAT IP address pool.</p> <p>Note An EIP that is already used in a DNAT entry cannot be used in an SNAT entry.</p> |
| ECS Granularity | |
| Available ECS Instances | <p>Select an ECS instance in the VPC.</p> <p>The ECS instance can access the Internet by using the specified EIP. Make sure that the following requirements are met:</p> <ul style="list-style-type: none"> ◦ The ECS instance is in the Running state. ◦ The ECS instance is not assigned an EIP or a static public IP address. |
| ECS CIDR Block | The CIDR block of the ECS instance. |

| Parameter | Description |
|-------------------|--|
| Public IP Address | <p>Select the EIP that is used to access the Internet.</p> <p>You can select one or more EIPs. You can use multiple EIPs to create an SNAT IP address pool.</p> <p>The maximum bandwidth supported by each EIP in an SNAT IP address pool is 200 Mbit/s. To make full use of your EIP bandwidth plan and avoid port conflicts caused by insufficient EIPs, we recommend that you add EIPs to the SNAT address pool based on the following rules:</p> <ul style="list-style-type: none"> ◦ If the maximum bandwidth value of the EIP bandwidth plan is 1,024 Mbit/s, specify at least five EIPs in each SNAT entry. ◦ If the maximum bandwidth value of the EIP bandwidth plan exceeds 1,024 Mbit/s, specify at least one additional EIP in each SNAT entry for every incremental 200 Mbit/s. <div style="background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note An EIP that is already used in a DNAT entry cannot be used in an SNAT entry.</p> </div> |

18.7.3. Modify an SNAT entry

This topic describes how to modify the public IP address in an SNAT entry.

Procedure

1. [Log on to the NAT Gateway console](#).
2. In the top navigation bar, select the region where the NAT gateway is deployed.
3. On the **NAT Gateway** page, find the NAT gateway and click **Configure SNAT** in the **Actions** column.
4. On the **SNAT Table** page, find the SNAT entry, and click **Edit** in the **Actions** column.
5. In the **Edit SNAT Entry** panel, change the public IP address and click **OK**.

18.7.4. Delete a SNAT entry

This topic describes how to delete a Source Network Address Translation (SNAT) entry. You can delete the SNAT entry if the Elastic Compute Service (ECS) instances without public IP addresses in a Virtual Private Cloud (VPC) network no longer need the SNAT service to access the Internet.

Procedure

1. [Log on to the NAT Gateway console](#).
2. In the top navigation bar, select the region where the NAT gateway is deployed.
3. On the **NAT Gateway** page, find the NAT gateway and click **Configure SNAT** in the **Actions** column.
4. On the **SNAT table** page, find the target SNAT entry, and click **Remove** in the **Actions** column.
5. In the dialog box that appears, click **OK**.

18.8. NAT service plan

18.8.1. Create a NAT service plan

You can associate an elastic IP address (EIP) or a NAT service plan to a NAT gateway. However, you can choose only one of them for the NAT gateway. If you want to associate a NAT service plan with the NAT gateway, you must create a NAT service plan first. Then, you can configure SNAT or DNAT for the NAT gateway. A NAT service plan consists of public IP addresses and Internet bandwidth.

Procedure

1. [Log on to the NAT Gateway console.](#)
2. On the **NAT Gateways** page, find the target NAT gateway and choose **Purchase NAT Bandwidth Package** in the **Internet Shared Bandwidth** column.
3. On the **Bandwidth Package Details** page, click **Purchase**.
4. On the **NAT Bandwidth Package** page, set the following parameters, and click **Submit**.

| Parameter | Description |
|---------------------------|---|
| Region | Indicates the region for which the NAT service plan is purchased. |
| Billing methods | Select the billing method of the NAT service plan. Only By Bandwidth is supported. |
| Bandwidth (Mbit/s) | Enter a bandwidth value for the NAT service plan that you want to purchase. The maximum value is 5000 Mbit/s. |
| Name | Enter a name for the NAT service plan. The name must be 2 to 128 characters in length and can contain letters, digits, underscores (_), and hyphens (-). It must start with a letter or Chinese character. |
| Description | Enter a description for the NAT service plan. The description must be 2 to 256 characters in length and cannot start with <code>http://</code> or <code>https://</code> . |
| Quantity | Enter the number of NAT bandwidth plans that you want to purchase. |

18.8.2. Modify the bandwidth of a NAT service plan

This topic describes how to modify the bandwidth of a NAT bandwidth plan. The modification takes effect immediately.

Procedure

1. [Log on to the NAT Gateway console.](#)
2. On the **NAT Gateways** page, find the target NAT gateway and click the ID of the NAT service plan in the **Internet Shared Bandwidth** column.
3. On the **Bandwidth Package Details** page, click the target NAT service plan, and then choose **Modify Bandwidth**.
4. On the **Modify Bandwidth** page, modify the bandwidth, and then click **Submit**.

Each NAT bandwidth plan supports a maximum of 5,000 Mbit/s in bandwidth.

18.8.3. Add an IP address

This topic describes how to add IP addresses to a NAT service plan. The added IP addresses can be used to create Source Network Address Translation (SNAT) and Destination Network Address Translation (DNAT) rules.

Procedure

1. [Log on to the NAT Gateway console.](#)
2. On the **NAT Gateways** page, find the target NAT gateway and click the ID of the NAT service plan in the **Internet Shared Bandwidth** column.
3. On the **Bandwidth Package Details** page, click the target NAT service plan, and then choose **Add IP Address**.
4. On the **Modify IP Addresses** page, enter the number of IP addresses to be added, and then click **Submit**.

18.8.4. Release an IP address

This topic describes how to release IP addresses in a NAT service plan. The NAT service plan must contain at least one IP address.

Prerequisites

Before you release an IP address in the NAT service plan, make sure that the IP address is not used in Source Network Address Translation (SNAT) and Destination Network Address Translation (DNAT) entries. If the IP address is used in an SNAT or DNAT entry, delete the SNAT or DNAT entry first. For more information, see [Delete a DNAT entry](#) and [Delete a SNAT entry](#).

Procedure

1. [Log on to the NAT Gateway console.](#)
2. On the **NAT Gateways** page, find the target NAT gateway and click the ID of the NAT service plan in the **Internet Shared Bandwidth** column.
3. On the **Bandwidth Package Details** page, click the target NAT service plan.
4. In the **Public IP List** section, find the target IP address, and click **Release** in the **Actions** column.
5. In the **Release IP** dialog box, click **OK**.

18.8.5. Delete a NAT service plan

This topic describes how to delete a service plan.

Prerequisites

Before you start, make sure that the following requirements are met:

- Delete the IP addresses that are used in Destination Network Address Translation (DNAT) entries. For more information, see [Delete a DNAT entry](#).
- Delete the IP addresses that are used for Source Network Address Translation (SNAT) entries. For more information, see [Delete a SNAT entry](#).

Procedure

1. [Log on to the NAT Gateway console.](#)
2. On the **NAT Gateways** page, find the target NAT gateway and click the ID of the NAT service plan in the **Internet Shared Bandwidth** column.
3. On the **Bandwidth Package Details** page, find the target NAT service plan and click **Delete**.
4. In the **Delete Shared Internet Shared Bandwidth** dialog box, click **OK**.

18.9. Anti-DDoS Origin Basic

A distributed denial-of-service (DDoS) attack is a malicious network attack against one or more systems, which can crash the targeted network. Alibaba Cloud provides up to 5 Gbit/s of basic anti-DDoS protection for a NAT gateway free of charge. Anti-DDoS Origin Basic can effectively prevent DDoS attacks.

How Anti-DDoS Origin Basic works

After you enable Anti-DDoS Origin Basic, traffic from the Internet must pass through Alibaba Cloud Security before the traffic arrives at the NAT gateway. Anti-DDoS Origin Basic scrubs and filters common DDoS attacks at Alibaba Cloud Security. Anti-DDoS Origin Basic protects your services against attacks such as SYN floods, UDP floods, ACK floods, ICMP floods, and DNS Query floods.

Anti-DDoS Origin Basic specifies the traffic scrubbing and blackhole triggering thresholds based on the bandwidth limit of the elastic IP address (EIP) that is associated with the NAT gateway. When the inbound traffic reaches the threshold, traffic scrubbing or blackhole is triggered:

- **Traffic scrubbing:** When the attack traffic from the Internet exceeds the scrubbing threshold or matches the attack traffic pattern, Alibaba Cloud Security starts to scrub the attack traffic. Traffic scrubbing includes packet filtering, bandwidth capping, and traffic throttling.
- **Blackhole:** When the attack traffic from the Internet exceeds the blackhole triggering threshold, blackhole is triggered and all inbound traffic is dropped.

Traffic scrubbing and blackhole triggering thresholds

The following table describes the methods that are used to calculate the traffic scrubbing and blackhole triggering thresholds on NAT gateways.

| Bandwidth limit of the EIP | Traffic scrubbing threshold (bit/s) | Traffic scrubbing threshold (pps) | Default blackhole triggering threshold |
|-----------------------------------|-------------------------------------|-----------------------------------|--|
| Lower than or equal to 800 Mbit/s | 800 Mbit/s | 120,000 | 1.5 Gbit/s |
| Higher than 800 Mbit/s | Predefined bandwidth | Predefined bandwidth × 150 | Predefined bandwidth × 2 |

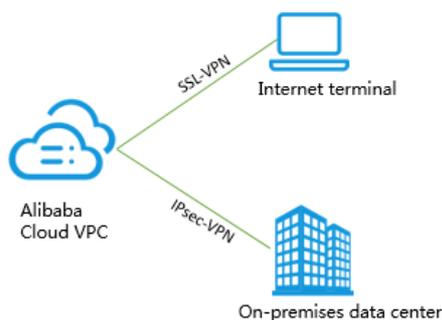
If the bandwidth limit of the EIP is 1,000 Mbit/s, the traffic scrubbing threshold (bit/s) is 1,000 Mbit/s, the traffic scrubbing threshold (pps) is 150,000, and the default blackhole triggering threshold is 2 Gbit/s.

19.VPN Gateway

19.1. What is VPN Gateway?

VPN Gateway is an Internet-based service that allows you to connect enterprise data centers, office networks, or Internet-facing terminals to Alibaba Cloud Virtual Private Cloud (VPC) networks through secure and reliable connections. VPN Gateway supports both IPsec-VPN connections and SSL-VPN connections.

Note The Alibaba Cloud VPN Gateway service complies with the local regulations and policies. VPN Gateway does not provide Internet access services.



Features

VPN Gateway supports the following features:

- IPsec-VPN

Route-based IPsec-VPN allows you to route network traffic in multiple ways, and also facilitates the configuration and maintenance of VPN policies.

You can use IPsec-VPN to connect an on-premises data center to a VPC network or connect two VPC networks. IPsec-VPN supports the IKEv1 and IKEv2 protocols. Any devices that support these two protocols can connect to Alibaba Cloud VPN Gateway, such as devices manufactured by Huawei, H3C, Hillstone, Sangfor, Cisco ASA, Juniper, SonicWall, Nokia, IBM, and Ixia.

- SSL-VPN

SSL-VPN is implemented based on the OpenVPN framework. You can create an SSL-VPN connection to connect a remote client to applications and services deployed in a VPC network. After you deploy your applications or services, you only need to import the certificate to the client to initiate a connection.

Benefits

VPN Gateway offers the following benefits:

- High security: You can use the IKE and IPsec protocols to encrypt data for secure and reliable data transmission.
- High availability: VPN Gateway adopts the hot-standby architecture to achieve failover within a few seconds, session persistence, and zero service downtime.
- Cost-effectiveness: The encrypted Internet connections provided by VPN Gateway are more cost-effective than leased lines.
- Ease of use: VPN Gateway is a ready-to-use service. VPN gateways start to work immediately after they are deployed.

19.2. Log on to the VPN Gateway console

This topic describes how to log on to the Apsara Uni-manager Management Console. You can manage your VPN gateways in the console. The Google Chrome browser is used as an example.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

Note When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top menu bar, choose **Products > Networking > Virtual Private Cloud**.
5. In the left-side navigation pane, choose **Cross-Network Connections > VPN**.

19.3. Get started with IPsec-VPN

19.3.1. Connect a data center to a VPC

This topic describes how to connect a data center to a virtual private cloud (VPC) by using an IPsec-VPN connection.

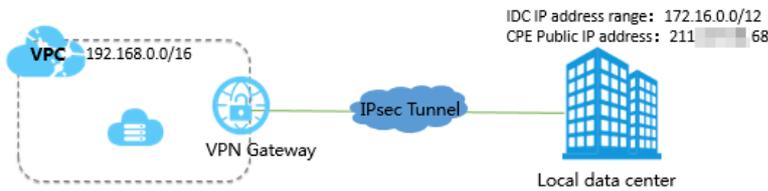
Prerequisites

Before you start, make sure that the following requirements are met:

- The gateway device in the data center is checked. VPN Gateway supports the standard IKEv1 and IKEv2 protocols. Gateway devices that support these two protocols can connect to VPN gateways on Alibaba Cloud, such as gateway devices that are manufactured by H3C, 华为、Hillstone, Sangfor, Cisco ASA, Juniper, SonicWall, Nokia, IBM, or Ixia.
- A static public IP address is assigned to the gateway device in the data center.
- The CIDR block of the data center does not overlap with that of the VPC.

Context

An enterprise has created a VPC on Alibaba Cloud. The CIDR block of the VPC is 192.168.0.0/16. The CIDR block of the data center is 172.16.0.0/12. The static public IP address for the gateway device in the data center is 211.XX.XX.68. To meet business requirements, the enterprise needs to connect the data center to the VPC.



As shown in the preceding figure, you can establish an IPsec-VPN connection between the data center and the VPC.

Step 1: Create a VPN gateway

Take the following steps to create a VPN gateway:

1. [Log on to the VPN Gateway console.](#)
2. In the left-side navigation pane, choose **VPN > VPN Gateways**.
3. On the **VPN Gateways** page, click **Create VPN Gateway**.
4. On the **Create VPN Gateway** page, set the following parameters for the VPN gateway, and then click **Submit**.
 - **Organization:** Select the organization to which the VPN gateway belongs.
 - **Resource Set:** Select the resource set to which the VPN gateway belongs.
 - **Region:** Select the region where you want to deploy the VPN gateway.

Note Make sure that the VPN gateway and the VPC are deployed in the same region.

- **Name:** Enter a name for the VPN gateway.
The name must be 2 to 128 characters in length, and can contain letters, digits, periods (.), underscores (_), and hyphens (-). It must start with a letter, and cannot start with `http://` or `https://`.
- **VPC:** Select the VPC to be associated with the VPN gateway.
- **vSwitch:** Select the vSwitch to which you want to attach the VPN gateway.
- **Bandwidth:** Select a maximum bandwidth value for the VPN gateway. The bandwidth is used for data transfer over the Internet.
- **IPsec-VPN:** Specify whether to enable IPsec-VPN for the VPN gateway. In this example, **Enable** is selected.
After IPsec-VPN is enabled, you can create a secure IPsec tunnel to connect a data center to a VPC, or connect two VPCs.
- **SSL-VPN:** Specify whether to enable SSL-VPN. In this example, **Disable** is selected.
SSL-VPN connections are point-to-site connections. SSL-VPN allows you to connect a client to Alibaba Cloud without the need to configure a customer gateway.
- **SSL Connections:** Select the maximum number of concurrent SSL connections that the VPN gateway supports.

Note This parameter is available only after SSL-VPN is enabled.

5. Go to the **VPN Gateways** page to view the newly created VPN gateway.

The newly created VPN gateway is in the **Preparing** state. The VPN gateway changes to the **Normal** state after about 2 minutes. The **Normal** state indicates that the VPN gateway is initialized and ready for use.

 **Note** It takes about 1 to 5 minutes to create a VPN gateway.

Step 2: Create a customer gateway

Perform the following operations to create a customer gateway.

1. In the left-side navigation pane, choose **VPN > Customer Gateways**.
2. Select the region where you want to create the customer gateway.
3. On the **Customer Gateways** page, click **Create Customer Gateway**.
4. On the **Create Customer Gateway** page, set the following parameters, and click **Submit**.
 - **Organization**: Select the organization to which the customer gateway belongs.
 - **Resource Set**: Select the resource set to which the customer gateway belongs.
 - **Region**: Select the region where you want to deploy the customer gateway.

 **Note** Make sure that the customer gateway and the VPN gateway to be connected are deployed in the same region.

- **Zone**: Select the zone where you want to deploy the customer gateway.
- **Name**: Enter a name for the customer gateway.

The name must be 2 to 128 characters in length, and can contain digits, hyphens (-), and underscores (_). It must start with a letter and cannot start with `http://` or `https://`.
- **IP Address**: Enter the public IP address of the gateway device in the data center that you want to connect to the VPC. In this example, `211.XX.XX.68` is entered.
- **Description**: Enter a description for the customer gateway.

The description must be 2 to 256 characters in length, and can contain digits, hyphens (-), underscores (_), full width periods (。), full width commas (。), and full width colons (:). The description must start with a letter and cannot start with `http://` or `https://`.

Step 3: Create an IPsec-VPN connection

Perform the following operations to create an IPsec-VPN connection:

1. In the left-side navigation pane, choose **VPN > IPsec Connections**.
2. Select the region where you want to create an IPsec-VPN connection.
3. On the **IPsec Connections** page, click **Create IPsec Connection**.
4. On the **Create IPsec Connection** page, set the following parameters for the IPsec-VPN connection, and click **Submit**.
 - **Organization**: Select the organization to which the IPsec-VPN connection belongs.
 - **Resource Set**: Select the resource set to which the IPsec-VPN connection belongs.
 - **Region**: Select the region to which the IPsec-VPN connection belongs.
 - **Zone**: Select the zone to which the IPsec-VPN connection belongs.
 - **Name**: Enter a name for the IPsec-VPN connection.
 - **VPN Gateway**: Select the VPN gateway you created.
 - **Customer Gateway**: Select the customer gateway to be connected through the IPsec-VPN connection.
 - **Source CIDR Block**: Enter the CIDR block of the VPC with which the selected VPN gateway is associated. In this example, `192.168.0.0/16` is entered.
 - **Destination CIDR Block**: Enter the CIDR block of the data center. In this example, `172.16.0.0/12` is entered.

- **Immediate Effect**: Specify whether to start connection negotiations immediately.
 - **Yes**: immediately starts negotiations after you complete the configuration.
 - **No**: starts negotiations when traffic is detected.
- **Advanced Configuration**: Use **default** advanced configurations.
A pre-shared key is automatically generated by default.

Step 4: Load the configurations of the IPsec-VPN connection to the on-premises gateway device

Perform the following operations to load the configurations of the IPsec-VPN connection to the on-premises gateway device:

1. In the left-side navigation pane, choose **VPN > IPsec Connections**.
2. Select the region where the IPsec-VPN connection is created.
3. On the **IPsec Connections** page, find the IPsec-VPN connection you want to manage, and then choose **More > Download Configuration** in the **Actions** column.
4. Load the configurations of the IPsec-VPN connection to the on-premises gateway device based on the configuration requirements of the gateway device. For more information, consult the vendor of the gateway device.

RemoteSubnet and LocalSubnet in the downloaded configurations are opposite to RemoteSubnet and LocalSubnet that you specify when you create an IPsec-VPN connection. For a VPN gateway, RemoteSubnet refers to the CIDR block of the data center and LocalSubnet refers to the CIDR block of the VPC. For a customer gateway, LocalSubnet refers to the CIDR block of the data center and RemoteSubnet refers to the CIDR block of the VPC.

Step 5: Configure routes for the VPN gateway

Take the following steps to configure routes for the VPN gateway:

1. In the left-side navigation pane, choose **VPN > VPN Gateways**.
2. Select the region where the VPN gateway is deployed.
3. On the **VPN Gateways** page, find the VPN gateway for which you want to configure routes, and click its ID in the **Instance ID/Name** column.
4. In the **Destination-based routing** tab, click **Add Route Entry**.
5. In the **Add Route Entry** dialog box, set the following parameters, and click **OK**.
 - **Destination CIDR Block**: Enter the CIDR block of the data center. In this example, **172.16.0.0/12** is entered.
 - **Next Hop Type**: Select **IPsec Connection**.
 - **Next Hop**: Select an IPsec-VPN connection.
 - **Publish to VPC**: Specify whether to automatically advertise new route entries to the VPC route table. In this example, **Yes** is selected.
 - **Weight**: Select a weight. In this example, **100** is selected.

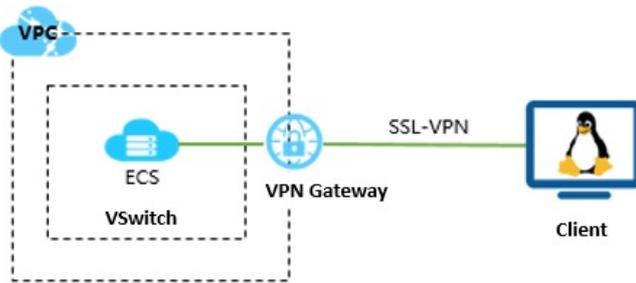
Step 6: Test the connectivity

Log on to an Elastic Compute Service (ECS) instance that is not assigned a public IP address in the VPC. Run the **ping** command to ping the private IP address of a server in the data center to test the connectivity.

19.4. Get started with SSL-VPN

19.4.1. Connect a Linux client to a VPC

This topic describes how to connect a Linux client to a virtual private cloud (VPC) through SSL-VPN connections.



Prerequisites

Before you start, make sure that the following requirements are met:

- A VPC is created.
- The CIDR block of the VPC must be different from that of the on-premises device.
- Your client can access the Internet.

Step 1: Create a VPN gateway

Perform the following operations to create a VPN gateway:

1. [Log on to the VPN Gateway console.](#)
2. In the left-side navigation pane, choose **VPN > VPN Gateways**.
3. On the **VPN Gateways** page, click **Create VPN Gateway**.
4. On the **Create VPN Gateway** page, set the following parameters for the VPN gateway, and then click **Submit**.
 - **Organization**: Select the organization to which the VPN gateway belongs.
 - **Resource Set**: Select the resource set to which the VPN gateway belongs.
 - **Region**: Select the region where you want to deploy the VPN gateway.

Note Make sure that the VPC and the VPN gateway for the VPC are deployed in the same region.

- **Instance Name**: Enter a name for the VPN gateway.
The name must be 2 to 128 characters in length and can contain letters, digits, periods (.), colons (:), underscores (_), and hyphens (-). It must start with a letter and cannot start with `http://` or `https://`.
 - **VPC**: Select the VPC network to be associated with the VPN gateway.
 - **vSwitch**: Select the vSwitch to which you want to attach the VPN gateway.
 - **Bandwidth**: Specify the maximum bandwidth of the VPN gateway. The bandwidth is used for data transfer over the Internet.
 - **IPsec-VPN**: Specify whether to enable IPsec-VPN for the VPN gateway. In this example, **Disable** is selected.
 - **SSL-VPN**: Specify whether to enable SSL-VPN for the VPN gateway. In this example, **Enable** is selected.
SSL-VPN connections are point-to-site connections. SSL-VPN allows you to connect a client to Alibaba Cloud without configuring a gateway for the client.
 - **SSL Connections**: Specify the maximum number of concurrent SSL connections that the VPN gateway supports.
5. Return to the **VPN Gateways** page to view the VPN gateway.

The newly created VPN gateway is in the **Preparing** state. The VPN gateway changes to the **Normal** state after about 2 minutes. The **Normal** state indicates that the VPN gateway is initialized and ready for use.

 **Note** It takes about 1 to 5 minutes to create a VPN gateway.

Step 2: Create an SSL server

Perform the following steps to create an SSL server:

1. In the left-side navigation pane, choose **VPN > SSL Servers**.
2. In the top navigation bar, select the region where you want to create the SSL server.
3. On the **SSL Servers** page, click **Create SSL Server**.
4. On the **Create SSL Server** page, set the following parameters for the SSL server, and then click **Submit**.
 - **Organization**: Select the organization to which the SSL server belongs.
 - **Resource Set**: Select the resource set to which the SSL server belongs.
 - **Region**: Select the region where you want to deploy the SSL server.
 - **Zone**: Select the zone where you want to deploy the SSL server.
 - **Name**: Enter a name for the SSL server.
 - **VPN Gateway**: Select a VPN gateway from the drop-down list.
 - **Local Network**: Enter the CIDR block of the network to which you want to connect. Click  to add more CIDR blocks. You can add the CIDR block of a VPC, a vSwitch, and an on-premises network.
 - **Client Subnet**: Enter the client CIDR block that the client uses to connect to the SSL server.
 - **Advanced Configuration**: Use **default** advanced configurations.

Step 3: Create and download an SSL client certificate

1. In the left-side navigation pane, choose **VPN > SSL Clients**.
2. In the top navigation bar, select the region where the SSL client is deployed.
3. On the **SSL Clients** page, click **Create Client Certificate**.
4. On the **Create SSL Client Certificate** page, set the following parameters for the SSL client, and then click **Submit**.
 - **Organization**: Select the organization to which the SSL client certificate belongs.
 - **Resource Set**: Select the resource set to which the SSL client certificate belongs.
 - **Region**: Select the region where you want to create the SSL client certificate.
 - **Zone**: Select the zone where you want to create the SSL client certificate.
 - **Name**: Enter a name for the SSL client certificate.
 - **VPN Gateway**: Select the VPN gateway to be associated with the SSL client certificate.
 - **SSL Server**: Select the SSL server to be associated with the SSL client certificate.
5. On the **SSL Clients** page, find the SSL client certificate and click **Download** in the **Actions** column.

The SSL client certificate is downloaded to your on-premises device.

Step 4: Configure the client

Perform the following steps to configure the Linux client:

1. Run the following command to install the OpenVPN client:

```
yum install -y openvpn
```

2. Extract the downloaded SSL client certificate and copy it to the `/etc/openvpn/conf/` directory.

3. Run the following command to launch OpenVPN:

```
openvpn --config /etc/openvpn/conf/config.ovpn --daemon
```

Step 5: Test the connectivity

Run the `ping` command to test the connectivity between the client and an Elastic Compute Service (ECS) instance in the VPC.

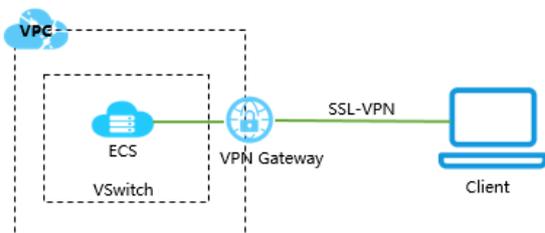
Note Make sure that the security group rules of the ECS instance allow remote access from Linux clients.

19.4.2. Connect a Windows client to a VPC

This topic describes how to connect a Windows client to a virtual private cloud (VPC) through SSL-VPN connections.

Context

The following scenario is used as an example.



Prerequisites

Before you start, make sure that the following requirements are met:

- A VPC is created.
- The CIDR block of the VPC must be different from that of the on-premises device.
- Your client can access the Internet.

Step 1: Create a VPN gateway

Perform the following operations to create a VPN gateway:

1. [Log on to the VPN Gateway console.](#)
2. In the left-side navigation pane, choose **VPN > VPN Gateways**.
3. On the **VPN Gateways** page, click **Create VPN Gateway**.
4. On the **Create VPN Gateway** page, set the following parameters for the VPN gateway, and then click **Submit**.
 - **Organization:** Select the organization to which the VPN gateway belongs.
 - **Resource Set:** Select the resource set to which the VPN gateway belongs.
 - **Region:** Select the region where you want to deploy the VPN gateway.

Note Make sure that the VPC and the VPN gateway for the VPC are deployed in the same region.

- **Instance Name:** Enter a name for the VPN gateway.

The name must be 2 to 128 characters in length and can contain letters, digits, periods (.), colons (:), underscores (_), and hyphens (-). It must start with a letter and cannot start with `http://` or `https://`.

- **VPC:** Select the VPC network to be associated with the VPN gateway.

- **vSwitch**: Select the vSwitch to which you want to attach the VPN gateway.
 - **Bandwidth**: Specify the maximum bandwidth of the VPN gateway. The bandwidth is used for data transfer over the Internet.
 - **IPsec-VPN**: Specify whether to enable IPsec-VPN for the VPN gateway. In this example, **Disable** is selected.
 - **SSL-VPN**: Specify whether to enable SSL-VPN for the VPN gateway. In this example, **Enable** is selected.
SSL-VPN connections are point-to-site connections. SSL-VPN allows you to connect a client to Alibaba Cloud without configuring a gateway for the client.
 - **SSL Connections**: Specify the maximum number of concurrent SSL connections that the VPN gateway supports.
5. Return to the **VPN Gateways** page to view the VPN gateway.

The newly created VPN gateway is in the **Preparing** state. The VPN gateway changes to the **Normal** state after about 2 minutes. The **Normal** state indicates that the VPN gateway is initialized and ready for use.

 **Note** It takes about 1 to 5 minutes to create a VPN gateway.

Step 2: Create an SSL server

Perform the following steps to create an SSL server:

1. In the left-side navigation pane, choose **VPN > SSL Servers**.
2. In the top navigation bar, select the region where you want to create the SSL server.
3. On the **SSL Servers** page, click **Create SSL Server**.
4. On the **Create SSL Server** page, set the following parameters for the SSL server, and then click **Submit**.
 - **Organization**: Select the organization to which the SSL server belongs.
 - **Resource Set**: Select the resource set to which the SSL server belongs.
 - **Region**: Select the region where you want to deploy the SSL server.
 - **Zone**: Select the zone where you want to deploy the SSL server.
 - **Name**: Enter a name for the SSL server.
 - **VPN Gateway**: Select a VPN gateway from the drop-down list.
 - **Local Network**: Enter the CIDR block of the network to which you want to connect. Click  to add more CIDR blocks. You can add the CIDR block of a VPC, a vSwitch, and an on-premises network.
 - **Client Subnet**: Enter the client CIDR block that the client uses to connect to the SSL server.
 - **Advanced Configuration**: Use **default** advanced configurations.

Step 3: Create and download an SSL client certificate

1. In the left-side navigation pane, choose **VPN > SSL Clients**.
2. In the top navigation bar, select the region where the SSL client is deployed.
3. On the **SSL Clients** page, click **Create Client Certificate**.
4. On the **Create SSL Client Certificate** page, set the following parameters for the SSL client, and then click **Submit**.
 - **Organization**: Select the organization to which the SSL client certificate belongs.
 - **Resource Set**: Select the resource set to which the SSL client certificate belongs.
 - **Region**: Select the region where you want to create the SSL client certificate.
 - **Zone**: Select the zone where you want to create the SSL client certificate.
 - **Name**: Enter a name for the SSL client certificate.

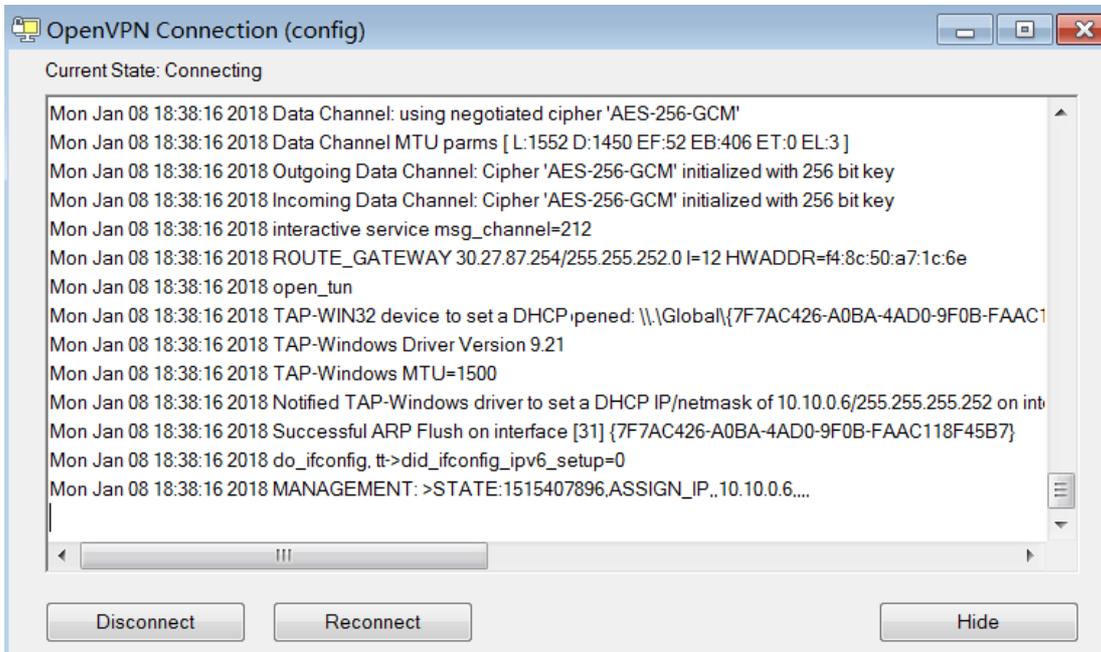
- **VPN Gateway:** Select the VPN gateway to be associated with the SSL client certificate.
 - **SSL Server:** Select the SSL server to be associated with the SSL client certificate.
5. On the **SSL Clients** page, find the SSL client certificate and click **Download** in the **Actions** column.
- The SSL client certificate is downloaded to your on-premises device.

Step 4: Configure the client

Perform the following steps to configure the Windows client:

 **Notice** You must run the client as an administrator.

1. Download and install all the OpenVPN client.
Download [OpenVPN](#).
2. Extract the downloaded SSL client certificate and copy it to the *OpenVPN\config* directory.
In this example, the certificate is copied to the *C:\Program Files\OpenVPN\config* directory. You must copy the certificate to the directory where the OpenVPN client is installed.
3. Start the OpenVPN client and click **Connect** to initiate a connection.



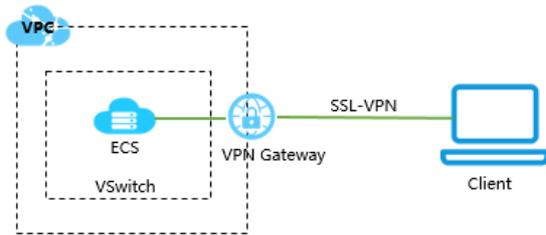
Step 5: Test the connectivity

Run the **ping** command to test the connectivity between the client and an Elastic Compute Service (ECS) instance in the VPC.

 **Note** Make sure that the security group rules of the ECS instance allow remote access from Windows clients.

19.4.3. Connect a macOS client to a VPC

This topic describes how to connect a macOS client to a virtual private cloud (VPC) through SSL-VPN connections.



Prerequisites

Before you start, make sure that the following requirements are met:

- A VPC is created.
- The CIDR block of the VPC must be different from that of the on-premises device.
- Your client can access the Internet.

Step 1: Create a VPN gateway

Perform the following operations to create a VPN gateway:

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > VPN Gateways**.
3. On the **VPN Gateways** page, click **Create VPN Gateway**.
4. On the **Create VPN Gateway** page, set the following parameters for the VPN gateway, and then click **Submit**.
 - **Organization**: Select the organization to which the VPN gateway belongs.
 - **Resource Set**: Select the resource set to which the VPN gateway belongs.
 - **Region**: Select the region where you want to deploy the VPN gateway.

Note Make sure that the VPC and the VPN gateway for the VPC are deployed in the same region.

- **Instance Name**: Enter a name for the VPN gateway.
The name must be 2 to 128 characters in length and can contain letters, digits, periods (.), colons (:), underscores (_), and hyphens (-). It must start with a letter and cannot start with `http://` or `https://`.
 - **VPC**: Select the VPC network to be associated with the VPN gateway.
 - **vSwitch**: Select the vSwitch to which you want to attach the VPN gateway.
 - **Bandwidth**: Specify the maximum bandwidth of the VPN gateway. The bandwidth is used for data transfer over the Internet.
 - **IPsec-VPN**: Specify whether to enable IPsec-VPN for the VPN gateway. In this example, **Disable** is selected.
 - **SSL-VPN**: Specify whether to enable SSL-VPN for the VPN gateway. In this example, **Enable** is selected.
SSL-VPN connections are point-to-site connections. SSL-VPN allows you to connect a client to Alibaba Cloud without configuring a gateway for the client.
 - **SSL Connections**: Specify the maximum number of concurrent SSL connections that the VPN gateway supports.
5. Return to the **VPN Gateways** page to view the VPN gateway.

The newly created VPN gateway is in the **Preparing** state. The VPN gateway changes to the **Normal** state after about 2 minutes. The **Normal** state indicates that the VPN gateway is initialized and ready for use.

Note It takes about 1 to 5 minutes to create a VPN gateway.

Step 2: Create an SSL server

Perform the following steps to create an SSL server:

1. In the left-side navigation pane, choose **VPN > SSL Servers**.
2. In the top navigation bar, select the region where you want to create the SSL server.
3. On the **SSL Servers** page, click **Create SSL Server**.
4. On the **Create SSL Server** page, set the following parameters for the SSL server, and then click **Submit**.
 - **Organization**: Select the organization to which the SSL server belongs.
 - **Resource Set**: Select the resource set to which the SSL server belongs.
 - **Region**: Select the region where you want to deploy the SSL server.
 - **Zone**: Select the zone where you want to deploy the SSL server.
 - **Name**: Enter a name for the SSL server.
 - **VPN Gateway**: Select a VPN gateway from the drop-down list.
 - **Local Network**: Enter the CIDR block of the network to which you want to connect. Click  to add more CIDR blocks. You can add the CIDR block of a VPC, a vSwitch, and an on-premises network.
 - **Client Subnet**: Enter the client CIDR block that the client uses to connect to the SSL server.
 - **Advanced Configuration**: Use **default** advanced configurations.

Step 3: Create and download an SSL client certificate

1. In the left-side navigation pane, choose **VPN > SSL Clients**.
2. In the top navigation bar, select the region where the SSL client is deployed.
3. On the **SSL Clients** page, click **Create Client Certificate**.
4. On the **Create SSL Client Certificate** page, set the following parameters for the SSL client, and then click **Submit**.
 - **Organization**: Select the organization to which the SSL client certificate belongs.
 - **Resource Set**: Select the resource set to which the SSL client certificate belongs.
 - **Region**: Select the region where you want to create the SSL client certificate.
 - **Zone**: Select the zone where you want to create the SSL client certificate.
 - **Name**: Enter a name for the SSL client certificate.
 - **VPN Gateway**: Select the VPN gateway to be associated with the SSL client certificate.
 - **SSL Server**: Select the SSL server to be associated with the SSL client certificate.
5. On the **SSL Clients** page, find the SSL client certificate and click **Download** in the **Actions** column.

The SSL client certificate is downloaded to your on-premises device.

Step 4: Configure the client

Perform the following steps to configure the macOS client:

1. Run the following command to install the OpenVPN client:

```
brew install openvpn
```

 **Note** Make sure that homebrew is installed before you install OpenVPN.

2. Extract the certificate that you downloaded in Step 3 and copy it to the directory where the OpenVPN client is installed. Then, initiate an SSL-VPN connection.
 - i. Back up the default configuration file.

- ii. Run the following command to delete the default configuration file:

```
rm /usr/local/etc/openssl/*
```

- iii. Run the following command to copy the file to the configuration directory:

```
cp cert_location /usr/local/etc/openssl/
```

In the preceding command, replace `cert_location` with the directory where the certificate is downloaded in Step 3. For example: `/Users/example/Downloads/certs6.zip`.

- iv. Run the following command to decompress the SSL client certificate package:

```
cd /usr/local/etc/openssl/
unzip /usr/local/etc/openssl/certs6.zip
```

- v. Run the following command to initiate a connection:

```
sudo /usr/local/opt/openssl/sbin/openssl --config /usr/local/etc/openssl/config.ovpn
```

Step 5: Test the connectivity

Run the `ping` command to test the connectivity between the client and an Elastic Compute Service (ECS) instance in the VPC.

 **Note** Make sure that the security group rules of the ECS instance allow remote access from macOS clients.

19.5. Manage a VPN Gateway

19.5.1. Create a VPN gateway

This topic describes how to create a VPN gateway. You must create a VPN gateway before you can use the IPsec-VPN and SSL-VPN features. After you create a VPN gateway, a public IP address is assigned to the VPN gateway.

Procedure

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > VPN Gateways**.
3. On the **VPN Gateways** page, click **Create VPN Gateway**.
4. On the **Create VPN Gateway** page, set the following parameters for the VPN gateway, and then click **Submit**.

| Parameter | Description |
|---------------------|--|
| Organization | Select the organization to which the VPN gateway belongs. |
| Resource Set | Select the resource set to which the VPN gateway belongs. |
| Region | Select the region where you want to create the VPN gateway. You can create IPsec-VPN connections on VPN gateways to connect a data center to a VPC or connect two VPCs. Make sure that the VPC and the VPN gateway associated with the VPC are deployed in the same region. |

| Parameter | Description |
|-----------------|--|
| Instance Name | <p>Enter a name for the VPN gateway.</p> <p>The name must be 2 to 128 characters in length, and can contain letters, digits, periods (.), underscores (_), and hyphens (-). It must start with a letter, and cannot start with <code>http://</code> or <code>https://</code>.</p> |
| VPC | Select the VPC network to be associated with the VPN gateway. |
| vSwitch | Select the vSwitch to which you want to attach the VPN gateway. |
| Bandwidth | Select a maximum bandwidth value for the VPN gateway. The bandwidth is used for data transfer over the Internet. |
| IPsec-VPN | <p>Specify whether to enable IPsec-VPN for the VPN gateway. The default value is Enable IPsec.</p> <p>After IPsec-VPN is enabled, you can create a secure IPsec tunnel to connect a data center to a VPC, or connect two VPCs.</p> |
| SSL-VPN | <p>Specify whether to enable SSL-VPN for the VPN gateway. The default value is Disable.</p> <p>SSL-VPN connections are point-to-site connections. SSL-VPN allows you to connect a client to Alibaba Cloud without the need to configure a customer gateway.</p> |
| SSL Connections | <p>Select the maximum number of concurrent SSL connections that the VPN gateway supports.</p> <div style="background-color: #e0f2f1; padding: 5px; border: 1px solid #ccc;"> <p> Note You can set this parameter only after SSL-VPN is enabled.</p> </div> |

19.5.2. Modify a VPN gateway

This topic describes how to modify the name and description of a VPN gateway.

Prerequisites

A VPN gateway is created. For more information, see [Create a VPN gateway](#).

Procedure

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > VPN Gateways**.
3. Select the region where the VPN gateway is deployed.
4. On the **VPN Gateways** page, find the target VPN gateway, and click the  icon in the **Instance ID/Name** column. In the dialog box that appears, enter a new name and click **OK**.
The name must be 2 to 100 characters in length, and can contain letters, digits, underscores (_), and hyphens (-). It must start with a letter or Chinese character.
5. Click the  icon in the **Description** column. In the dialog box that appears, enter a new description and click **OK**.
The description must be 2 to 256 characters in length, and cannot start with `http://` or `https://`.

19.5.3. Configure routes of a VPN Gateway

19.5.3.1. Route overview

After you create an IPsec-VPN connection by using a VPN gateway, you must add a route to the VPN gateway.

Route-based IPsec-VPN allows you to route network traffic in multiple ways, and facilitates the configuration and maintenance of VPN policies.

You can add the following two types of route to a VPN gateway:

- Policy-based routes.
- Destination-based routes.

Policy-based routes

Policy-based routes forward traffic based on source and destination IP addresses.

For more information, see [Add a policy-based route entry](#).

 **Note** Policy-based routes take precedence over destination-based routes.

Destination-based routes

Destination-based routes forward traffic to specified destination IP addresses.

For more information, see [Add a destination-based route entry](#).

19.5.3.2. Work with a policy-based route

A policy-based route forwards traffic based on source and destination IP addresses. This topic describes how to create, advertise, modify, and delete a policy-based route.

Prerequisites

An IPsec-VPN connection is created. For more information, see [Create an IPsec-VPN connection](#).

Add a policy-based route

After you create an IPsec-VPN connection, you can create a policy-based route for the IPsec-VPN connection.

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > VPN Gateways**.
3. Select the region where the VPN gateway is deployed.
4. On the **VPN Gateways** page, find the VPN gateway you want to manage and click the gateway ID in the **Instance ID/Name** column.
5. Click the **Policy-based Routing** tab, and then click **Add Route Entry**.
6. In the **Add Route Entry** panel, set the following parameters and click **OK**.

| Parameter | Description |
|-------------------------------|--|
| Destination CIDR block | The private CIDR block that you want to access. |
| Source CIDR Block | The private CIDR block of the VPC. |
| Next Hop Type | Select IPsec Connection. |
| Next Hop | Select the IPsec-VPN connection for which you want to create the policy-based route. |

| Parameter | Description |
|----------------|---|
| Publish to VPC | <p>Specify whether to advertise the route to the VPC route table. Valid values:</p> <ul style="list-style-type: none"> ◦ Yes: automatically advertises the route to the route table of the VPC. We recommend that you select this value. ◦ No: does not advertise the route to the VPC route table. <p> Note If you select No, you must manually advertise the route to the VPC route table.</p> |
| Weight | <p>Select a weight. Valid values:</p> <ul style="list-style-type: none"> ◦ 100: specifies a high priority for the policy-based route. ◦ 0: specifies a low priority for the policy-based route. <p> Note If two policy-based routes are configured with the same destination CIDR block, you cannot set the weights of the routes to 100.</p> |

Advertise a policy-based route

1. [Log on to the VPN Gateway console.](#)
 2. In the left-side navigation pane, choose **VPN > VPN Gateways**.
 3. Select the region where the VPN gateway is deployed.
 4. On the **VPN Gateways** page, find the VPN gateway you want to manage and click the gateway ID in the **Instance ID/Name** column.
 5. On the **Policy-based Routing** tab, find the policy-based route that you want to advertise and click **Publish** in the **Actions** column.
 6. In the **Publish Route Entry** message, click **OK**.
- If you want to withdraw the policy-based route, click **Unpublish**.

Modify a policy-based route

You can change the weight of a policy-based route.

1. [Log on to the VPN Gateway console.](#)
2. In the left-side navigation pane, choose **VPN > VPN Gateways**.
3. Select the region where the VPN gateway is deployed.
4. On the **VPN Gateways** page, find the VPN gateway you want to manage and click the gateway ID in the **Instance ID/Name** column.
5. On the **Policy-based Routing** tab, find the policy-based route that you want to modify and click **Edit** in the **Actions** column.
6. In the panel that appears, specify a new weight for the route and click **OK**.

Delete a policy-based route

1. [Log on to the VPN Gateway console.](#)
2. In the left-side navigation pane, choose **VPN > VPN Gateways**.
3. Select the region where the VPN gateway is deployed.
4. On the **VPN Gateways** page, find the VPN gateway you want to manage and click the gateway ID in the **Instance ID/Name** column.

5. On the **Policy-based Routing** tab, find the policy-based route that you want to delete and click **Delete** in the **Actions** column.
6. In the **Delete Route Entry** message, click **OK**.

19.5.3.3. Manage destination-based routes

Destination-based routing is a technique that routes network traffic to specified destination IP addresses. This topic describes how to create, advertise, modify, and delete a destination-based route.

Prerequisites

An IPsec-VPN connection is created. For more information, see [Create an IPsec-VPN connection](#).

Create a destination-based route

After you create an IPsec-VPN connection, you can create a destination-based route for the IPsec-VPN connection.

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > VPN Gateways**.
3. Select the region where the VPN gateway is deployed.
4. On the **VPN Gateways** page, find the VPN gateway you want to manage and click the gateway ID in the **Instance ID/Name** column.
5. On the **Destination-based routing** tab, click **Add Route Entry**.
6. In the **Add Route Entry** dialog box, set the following parameters and click **OK**.

| Parameter | Description |
|-------------------------------|---|
| Destination CIDR block | Enter the CIDR block that you want to access. |
| Next Hop Type | Select IPsec Connection. |
| Next Hop | Select the IPsec-VPN connection for which you want to create a destination-based route. |
| Publish to VPC | <p>Specify whether to advertise the destination-based route to the virtual private cloud (VPC) route table.</p> <ul style="list-style-type: none"> ◦ Yes: automatically advertises the route to the route table of the VPC. We recommend that you select this value. ◦ No: does not advertise the destination-based route to the VPC route table. <p> Note If you select No, you must manually advertise the destination-based route to the VPC route table.</p> |
| Weight | <p>Select a weight. Valid values:</p> <ul style="list-style-type: none"> ◦ 100: specifies a high priority for the destination-based route. ◦ 0: specifies a low priority for the destination-based route. <p> Note If two destination-based routes are configured with the same destination CIDR block, you cannot set the weights of the routes to 100.</p> |

Advertise the destination-based route

1. [Log on to the VPN Gateway console](#).

2. In the left-side navigation pane, choose **VPN > VPN Gateways** .
3. Select the region where the VPN gateway is deployed.
4. On the **VPN Gateways** page, find the VPN gateway you want to manage and click the gateway ID in the **Instance ID/Name** column.
5. On the **Destination-based Routing** tab, find the destination-based route that you want to manage and click **Publish** in the **Actions** column.
6. In the **Publish Route Entry** message, click **OK**.
If you want to withdraw the destination-based route, click **Unpublish**.

Modify the destination-based route

You can change the weight of the destination-based route.

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > VPN Gateways** .
3. Select the region where the VPN gateway is deployed.
4. On the **VPN Gateways** page, find the VPN gateway you want to manage and click the gateway ID in the **Instance ID/Name** column.
5. On the **Destination-based Routing** tab, find the destination-based route that you want to manage and click **Edit** in the **Actions** column.
6. In the panel that appears, specify the weight of the destination-based route and click **OK**.

Delete the destination-based route

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > VPN Gateways** .
3. Select the region where the VPN gateway is deployed.
4. On the **VPN Gateways** page, find the VPN gateway you want to manage and click the gateway ID in the **Instance ID/Name** column.
5. On the **Destination-based Routing** tab, find the destination-based route that you want to manage and click **Delete** in the **Actions** column.
6. In the **Delete Route Entry** message, click **OK**.

19.5.4. Delete a VPN gateway

This topic describes how to delete a VPN gateway. After you delete a VPN gateway, you can no longer use the VPN gateway to establish IPsec-VPN or SSL-VPN connections.

Context

Before you delete a VPN gateway, make sure that the following conditions are met:

- The IPsec-VPN connections on the VPN gateway are deleted. For more information, see [Delete an IPsec-VPN connection](#).
- The SSL server associated with the VPN gateway is deleted. For more information, see [Delete an SSL server](#).

Procedure

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > VPN Gateways**.
3. Select the region where the VPN gateway is deployed.
4. On the **VPN Gateways** page, find the VPN gateway that you want to delete, and then click **Delete** in the **Actions** column.

- In the **Delete VPN Gateway** message, click **OK**.

19.5.5. View the monitoring information about a VPN gateway

This topic describes how to view the monitoring information about a VPN gateway.

Prerequisites

Before you start, make sure that the following requirements are met:

- Log Service is activated.
- CloudMonitor is activated.

Procedure

- Log on to the [VPN Gateway console](#).
- In the left-side navigation pane, choose **VPN > VPN Gateways**.
- On the **VPN Gateway** page, find the VPN gateway you want to manage and click the gateway ID.
- On the details page of the VPN gateway, click the **Monitor** tab.
- On the **Monitor** tab, view the monitoring information.

By default, the system displays the monitoring information within the last hour. You can also view the monitoring information within the last 3, 6, or 12 hours. In addition, you can view the monitoring information in a specified time period. VPN Gateway provides the following monitoring information:

| Metric | Description |
|--------------------------|--|
| Vpn rxPkgs | The rate at which packets are received by the VPN gateway. Unit: packet/s. This is the default unit. You can also select kilopacket/s or megapacket/s from the drop-down list next to the monitored item. |
| Vpn txPkgs | The rate at which packets are sent by the VPN gateway. Unit: packet/s. This is the default unit. You can also select kilopacket/s or megapacket/s from the drop-down list next to the monitored item. |
| Vpn net rate | The rate of inbound traffic. Unit: bit/s. This is the default unit. You can also select Kbit/s, Mbit/s, or Gbit/s from the drop-down list next to the monitored item. |
| Vpn net rate | The rate of outbound traffic. Unit: bit/s. This is the default unit. You can also select Kbit/s, Mbit/s, or Gbit/s from the drop-down list next to the monitored item. |
| ssl connect count(count) | The number of clients that are connected to Alibaba Cloud by using the SSL-VPN connection. |

19.6. Manage a customer gateway

19.6.1. Create a customer gateway

This topic describes how to create a customer gateway. You can use a customer gateway to establish an IPsec-VPN connection between a virtual private cloud (VPC) and a data center or between two VPCs. After you create a customer gateway, you can update the information about a gateway device in the data center to Alibaba Cloud. Then, you can connect the customer gateway to a VPN gateway. A customer gateway can connect to multiple VPN gateways.

Procedure

1. [Log on to the VPN Gateway console.](#)
2. In the left-side navigation pane, choose **VPN > Customer Gateways**.
3. Select the region where you want to deploy the customer gateway.

 **Note** Make sure that the customer gateway and the VPN gateway to be connected belong to the same region.

4. On the **Customer Gateways** page, click **Create Customer Gateway**.
5. On the **Create Customer Gateway** page, set the following parameters and click **Submit**.

| Parameter | Description |
|---------------------|--|
| Organization | Select the organization to which the customer gateway belongs. |
| Resource Set | Select the resource set to which the customer gateway belongs. |
| Region | Select the region where you want to deploy the customer gateway.  Note Make sure that the customer gateway and the VPN gateway to be connected belong to the same region. |
| Zone | Select the zone where you want to deploy the customer gateway. |
| Name | Enter a name for the customer gateway. The name must be 2 to 128 characters in length and can contain digits, hyphens (-), and underscores (_). The name must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . |
| IP Address | Enter the static public IP address of the gateway device in the data center. |
| Description | Enter a description for the customer gateway. The description must be 2 to 256 characters in length and can contain digits, hyphens (-), underscores (_), periods (.), commas (,), and colons (:). The description must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . |

19.6.2. Modify a customer gateway

This topic describes how to modify the name and description of a customer gateway.

Prerequisites

A customer gateway is created. For more information, see [Create a customer gateway](#).

Procedure

1. [Log on to the VPN Gateway console.](#)

2. In the left-side navigation pane, choose **VPN > Customer Gateways**.
3. Select the region where the customer gateway is deployed.
4. On the **Customer Gateways** page, find the target customer gateway, click the  icon in the **Instance ID** column. In the dialog box that appears, enter a name and click **OK**.
The name must be 2 to 128 characters in length, and can contain digits, underscores (_), and hyphens (-). It must start with a letter or Chinese character.
5. Click the  icon in the **Description** column. In the dialog box that appears, enter a new description and click **OK**.
The description must be 2 to 256 characters in length, and cannot start with http:// or https://.

19.6.3. Delete a customer gateway

This topic describes how to delete a customer gateway.

Procedure

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > Customer Gateways**.
3. Select the region where the customer gateway is deployed.
4. On the **Customer Gateways** page, find the customer gateway that you want to delete, and then click **Delete** in the **Actions** column.
5. In the **Delete Customer Gateway** message, click **OK**.

19.7. Configure IPsec-VPN connections

19.7.1. Configuration overview

This topic describes how to connect a VPC to an on-premises data center through IPsec-VPN.

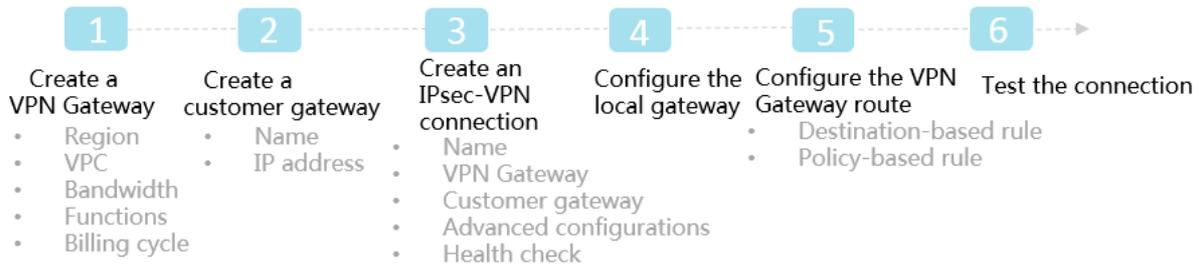
Prerequisites

Before creating a site-to-site VPN connection, make sure the following conditions are met:

- The protocols IKEv1 and IKEv2 are supported by the gateway device of the on-premises data center.
IPsec-VPN supports IKEv1 and IKEv2 protocols. Devices that support these two protocols can connect to Alibaba Cloud VPN Gateway, including devices of Huawei, H3C, Hillstone, SANGFOR, Cisco ASA, Juniper, SonicWall, Nokia, IBM, and Ixia.
- A static public IP address is configured for the local gateway.
- The IP address ranges of the VPC and on-premises data center to be connected do not conflict with each other.

Procedure

The following figure shows the procedure of connecting a VPC to an on-premises data center through IPsec-VPN.



1. Create a VPN Gateway

Enable the IPsec-VPN function. Up to 10 IPsec-VPN connections can be established in a VPN Gateway.

2. Create a customer gateway

By creating a customer gateway, you can register the local gateway to Alibaba Cloud and connect the customer gateway to the VPN Gateway. A customer gateway can be connected to multiple VPN Gateways.

3. Create an IPsec connection

An IPsec connection is a VPN channel established between a VPN Gateway and a customer gateway. The encrypted communication between the VPN Gateway and the on-premises data center can be achieved only after the IPsec connection is established.

4. Configure the local gateway

You need to load the VPN Gateway configurations to the local gateway device.

5. Configure the VPN Gateway route

You need to configure a route in the VPN Gateway and publish it to the VPC route table.

6. Test the connection

Log on to an ECS instance (without a public IP address) in the connected VPC. ping the private IP address of a server in the on-premises data center to check whether the connection is established.

19.7.2. Manage an IPsec-VPN connection

19.7.2.1. Create an IPsec-VPN connection

This topic describes how to create an IPsec-VPN connection. After you create a VPN gateway and a customer gateway, you can create an IPsec-VPN connection between the two gateways for encrypted data transmission.

Procedure

1. Log on to the VPN Gateway console.
2. In the left-side navigation pane, choose VPN > IPsec Connections.
3. In the top navigation bar, select the region where you want to create the IPsec-VPN connection.
4. On the IPsec Connections page, click Create IPsec Connection.
5. On the Create IPsec Connection page, configure the IPsec-VPN connection based on the following information and click Submit.

| Parameter | Description |
|--------------|--|
| Organization | Select the organization to which the IPsec-VPN connection belongs. |
| Resource Set | Select the resource set to which the IPsec-VPN connection belongs. |

| Parameter | Description |
|--|---|
| Region: | Select the region where the IPsec-VPN connection is created. |
| Zone | Select the zone where the IPsec-VPN connection is created. |
| Name | Enter a for the IPsec-VPN connection. The name must be 2 to 128 characters in length, and can contain digits, hyphens (-), and underscores (_). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . |
| VPN Gateway | Select the VPN gateway to be connected through the IPsec-VPN connection. |
| Customer Gateway | Select the customer gateway to be connected through the IPsec-VPN connection. |
| Local Network | Enter the CIDR block of the VPC to be connected to the data center. The CIDR block is used in Phase 2 negotiations. You can add multiple CIDR blocks of the VPC only if IKEv2 is used. |
| Remote Network | Enter the CIDR block of the data center to be connected to the VPC. This CIDR block is used in Phase 2 negotiations. You can add multiple CIDR blocks of the data center only if IKEv2 is used. |
| Effective Immediately | Specify whether to start connection negotiations immediately. <ul style="list-style-type: none"> ◦ Yes: immediately starts negotiations after you complete the configuration. ◦ No: starts negotiations when traffic is detected. |
| Advanced Configurations | Select the type of advanced configurations. <ul style="list-style-type: none"> ◦ Default: Use default settings. ◦ Configure: Use custom settings. |
| Advanced configuration: IKE configuration | |
| Pre-Shared Key | Enter the pre-shared key used for authentication between the VPN gateway and the customer gateway. You can specify a key, or use the default key that is randomly generated by the system. |
| Version | Select an IKE version. <ul style="list-style-type: none"> ◦ ikev1 ◦ ikev2 IKEv1 and IKEv2 are supported. Compared with IKEv1, IKEv2 simplifies the Security Association (SA) negotiation process and provides better support for scenarios where multiple CIDR blocks are used. We recommend that you select IKEv2. |
| Negotiation Mode | Select the negotiation mode of IKEv1. <ul style="list-style-type: none"> ◦ main: This mode offers higher security during negotiations. ◦ aggressive: This mode is faster and has a higher success rate. Connections negotiated in both modes ensure the same security level of data transmission. |

| Parameter | Description |
|--|--|
| Encryption Algorithm | Select the encryption algorithm to be used in Phase 1 negotiations. Supported algorithms are aes, aes192, aes256, des, and 3des. |
| Authentication Algorithm | Select the authentication algorithm to be used in Phase 1 negotiations. Supported algorithms are sha1 and md5. |
| DH Group | Select the Diffie-Hellman key exchange algorithm to be used in Phase 1 negotiations. |
| SA Life Cycle (seconds) | Specify the lifecycle of the SA after Phase 1 negotiations succeed. Default value: 86400 . Unit: seconds. |
| LocalId | Specify the ID of the VPN gateway. The ID is used in Phase 1 negotiations. The default value is the public IP address of the VPN gateway. If you set LocalId to a FQDN, we recommend that you set Negotiation Mode to Aggressive. |
| RemoteId | Specify the ID of the customer gateway. The ID is used in Phase 1 negotiations. The default value is the public IP address of the customer gateway. If you set RemoteId to a FQDN, we recommend that you select set Negotiation Mode to Aggressive. |
| Advanced configuration: IPsec configuration | |
| Encryption Algorithm | Select the encryption algorithm to be used in Phase 2 negotiations. Supported algorithms are aes, aes192, aes256, des, and 3des. |
| Authentication Algorithm | Select the authentication algorithm to be used in Phase 2 negotiations. Supported algorithms are sha1 and md5. |
| DH Group | Select the Diffie-Hellman key exchange algorithm to be used in Phase 2 negotiations. <ul style="list-style-type: none"> ◦ If you select a value other than disabled, the PFS feature is enabled by default, which necessitates key update for every renegotiation. Therefore, you must also enable PFS for the client. ◦ For clients that do not support PFS, select disabled. |
| SA Life Cycle (seconds) | Specify the lifecycle of the SA after Phase 2 negotiations succeed. Default value: 86400 . Unit: seconds. |

19.7.2.2. Modify an IPsec-VPN connection

This topic describes how to modify the name, advanced settings, and health check for an IPsec-VPN connection.

Prerequisites

An IPsec-VPN connection is created. For more information, see [Create an IPsec-VPN connection](#).

Procedure

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > IPsec Connections**.
3. In the top navigation bar, select the region where the IPsec-VPN connection is created.
4. On the **IPsec Connections** page, find the target IPsec-VPN connection and then click **Edit** in the **Actions** column.
5. In the **Modify IPsec Connections** dialog box, modify the name, advanced settings, and health check, and click **Submit**.

19.7.2.3. Download the configuration file of an IPsec-VPN connection

This topic describes how to download the configurations of an IPsec-VPN connection, and load the configurations to the customer gateway device after an IPsec-VPN connection is configured.

Prerequisites

An IPsec-VPN connection is created. For more information, see [Create an IPsec-VPN connection](#).

Procedure

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > IPsec Connections**.
3. In the top navigation bar, select the region where the IPsec-VPN connection is created.
4. On the **IPsec Connections** page, find the target IPsec-VPN connection, and then choose **More > Download Configuration** in the **Actions** column.

Note RemoteSubnet and LocalSubnet in the downloaded configurations are opposite to RemoteSubnet and LocalSubnet that you specify when you create an IPsec-VPN connection. For a VPN gateway, RemoteSubnet refers to the CIDR block of the on-premises data center and LocalSubnet refers to the CIDR block of the VPC network. For a customer gateway, LocalSubnet refers to the CIDR block of the on-premises data center and RemoteSubnet refers to the CIDR block of the VPC network.

19.7.2.4. Configure a security group

This topic describes how to configure a security group to control the inbound and outbound traffic of ECS instances in the security group after an IPsec-VPN connection is created.

Procedure

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > IPsec Connections**.
3. In the top navigation bar, select the region where the IPsec-VPN connection is created.
4. On the **IPsec Connections** page, find the IPsec-VPN connection you want to manage, and then choose **More > Configure Routing Group** in the **Actions** column.
5. In the **Configure Routing Group** panel, set the following parameters, and then click **Submit**.

| Parameter | Description |
|-----------------------|---|
| Security Group | Select the security group to which you want to add the security group rule. |
| Rule Direction | Select the direction in which the rule applies. <ul style="list-style-type: none"> ◦ Outbound: from the ECS instances in the current security group to other ECS instances on Alibaba Cloud or resources on the Internet. ◦ Inbound: from other ECS instances on Alibaba Cloud or resources on the Internet to the ECS instances in the current security group. |

| Parameter | Description |
|---|---|
| Policy | <p>Select an authorization policy.</p> <ul style="list-style-type: none"> ◦ Allow: allows access requests on the specified ports. ◦ Deny: discards requests on the specified ports without returning messages. <p>If you specify different authorization policies for two security group rules but the other settings are the same, the Deny rule prevails over the Allow rule.</p> |
| Protocol Type | Select a protocol type. |
| Port Range | <p>Select a port range for the security group rule. Specify the port range in the format of 1/200 or 80/80.</p> <ul style="list-style-type: none"> ◦ 1/200 specifies a port range from 1 to 200. ◦ 80/80 specifies port 80. <p>Valid values: -1 and 1 to 65535. A value of -1/-1 specifies all ports.</p> |
| Priority | Set a priority for the rule. Valid values: 1 to 100. The default value is 1, which specifies the highest priority. |
| Authorization Type | <p>Select the authorization type of the security group rule.</p> <p>You can select only Address.</p> |
| Network Type | <p>Select a network interface controller (NIC) type.</p> <ul style="list-style-type: none"> ◦ Internal: controls inbound and outbound traffic within Alibaba Cloud. ◦ External: controls inbound and outbound traffic over the Internet. |
| Authorization Object | <p>Select the CIDR blocks to be authorized by the security group rule.</p> <p>You can specify up to 10 CIDR blocks at a time.</p> |
| Enable Automatically Configure Routers | Specify whether to automatically configure routers. The feature is enabled by default. |
| Description | <p>Enter a description for the security group rule.</p> <p>The description must be 2 to 256 characters in length, and cannot start with <code>http://</code> or <code>https://</code>. You can leave this parameter empty.</p> |

19.7.2.5. View IPsec-VPN connection logs

This topic describes how to view IPsec-VPN connection logs that are generated within the last 30 days to troubleshoot connection errors. You can query log data generated within 10 minutes.

Procedure

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > IPsec Connections**.
3. In the top navigation bar, select the region where the IPsec-VPN connection is created.
4. On the **IPsec Connections** page, find the target IPsec-VPN connection, and then choose **More > View Logs**

in the **Actions** column.

5. In the **IPsec Connection Logs** dialog box, set the time range and query the logs.

19.7.2.6. Delete an IPsec-VPN connection

This topic describes how to delete an IPsec-VPN connection.

Procedure

1. [Log on to the VPN Gateway console.](#)
2. In the left-side navigation pane, choose **VPN > IPsec Connections**.
3. In the top navigation bar, select the region where you want to delete the IPsec-VPN connection.
4. On the **IPsec Connections** page, find the IPsec-VPN connection that you want to delete, and click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

19.7.3. View the monitoring information about an IPsec-VPN connection

This topic describes how to view the monitoring information about an IPsec-VPN connection.

Prerequisites

Before you start, make sure that the following requirements are met:

- Log Service is activated.
- CloudMonitor is activated.

Procedure

1. [Log on to the VPN Gateway console.](#)
2. In the left-side navigation pane, choose **VPN > IPsec Connections**.
3. On the **IPsec Connections** page, find the IPsec-VPN connection that you want to manage and click its ID.
4. On the **Monitor** tab of the IPsec-VPN connection, view the monitoring information about the connection.

By default, the system displays monitoring information within the last hour. You can also view monitoring information within the last 3, 6, or 12 hours. In addition, you can view monitoring information within a specified time period. VPN Gateway provides the following monitoring information:

| Metric | Description |
|------------------------------------|---|
| Ipsec receive pkg Counter | The rate at which packets are received by the IPsec-VPN connection. Unit: packet/s. This is the default unit. You can also select kilopacket/s or megapacket/s from the drop-down list next to the monitored item. |
| Ipsec transform pkg Counter | The rate at which packets are sent by the IPsec-VPN connection. Unit: packet/s. This is the default unit. You can also select kilopacket/s or megapacket/s from the drop-down list next to the monitored item. |

| Metric | Description |
|----------------------|---|
| Ipsec receive rate | The rate of inbound traffic. Unit: bit/s. This is the default unit. You can also select Kbit/s, Mbit/s, or Gbit/s from the drop-down list next to the monitored item. |
| Ipsec transform rate | The rate of outbound traffic. Unit: bit/s. This is the default unit. You can also select Kbit/s, Mbit/s, or Gbit/s from the drop-down list next to the monitored item. |

19.7.4. MTU considerations

The maximum transmission unit (MTU) is the size of the largest packet that can be transmitted over a network layer protocol, such as TCP. Packets are measured in bytes. The MTU takes both the sizes of headers and data into account.

Segments transmitted over an IPsec tunnel are encrypted and then encapsulated into packets for routing purpose. The size of a segment must fit the MTU of the packet that carries the segment. Therefore, the MTU of the segment must be smaller than the MTU of the packet.

Gateway MTU

You must set the MTU of the local VPN gateway to a value no greater than 1,360 bytes. We recommend that you set the MTU to 1,360 bytes.

The TCP protocol negotiates the maximum segment length (MSS) of each packet segment between the sender and the receiver. We recommend that you set the TCP MSS of the on-premises VPN gateway to 1,359 bytes to facilitate the encapsulation and transfer of TCP packets.

19.8. Configure SSL-VPN

19.8.1. SSL-VPN configuration overview

This topic describes how to remotely access a VPC through SSL-VPN.

Prerequisites

The following conditions must be met before you deploy a VPN Gateway:

- The client and the VPC are not using the same private CIDR block.
- The client is able to access the Internet.

Procedure

The following figure describes the process for configuring remote access to a VPC through SSL-VPN.

1. Create a VPN gateway
Create a VPN gateway and enable the SSL-VPN function.
2. Create an SSL server
Specify the range of IP addresses to be accessed and the range of IP addresses to be allocated to the client for making connections.
3. Create a client certificate

Create a client certificate according to the SSL server configuration, and then download the client certificate and the configuration file.

4. Configure the client

Download and install the VPN client, and then import the client certificate and the configuration file.

5. Configure the security group

Make sure that the rules in the security group of the ECS instance allow remote access from clients.

19.8.2. Manage an SSL server

19.8.2.1. Create an SSL server

This topic describes how to create an SSL server. Before you can create an SSL-VPN connection, you must create an SSL server.

Prerequisites

A VPN gateway is created and SSL-VPN is enabled for the VPN gateway. For more information, see [Create a VPN gateway](#).

Procedure

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > SSL Servers**.
3. In the top navigation bar, select the region where you want to create the SSL server.
4. On the **SSL Servers** page, click **Create SSL Server**.
5. On the **Create SSL Server** page, set the following parameters and click **Submit**.

| Parameter | Description |
|---------------------|---|
| Organization | Select the organization to which the SSL server belongs. |
| Resource Set | Select the resource set to which the SSL server belongs. |
| Region | Select the region where you want to deploy the SSL server. |
| Zone | Select the zone where you want to deploy the SSL server. |
| Name | Enter a name for the SSL server. The name must be 2 to 128 characters in length, and can contain digits, hyphens (-), and underscores (_). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . |
| VPN Gateway | Select the VPN gateway that you want to associate with the SSL server. Make sure that SSL-VPN is enabled for the VPN gateway. |

| Parameter | Description |
|-------------------------|---|
| Local Network | <p>Enter the CIDR block that the client needs to access through the SSL-VPN connection. It can be the CIDR block of a VPC, a vSwitch, a data center connected to a VPC through an Express Connect circuit, or a cloud service such as ApsaraDB RDS or Object Storage Service (OSS).</p> <p>Click + to add more CIDR blocks.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p>? Note The subnet mask of the specified CIDR block must be 16 to 29 bits in length.</p> </div> |
| Client CIDR Block | <p>Enter the CIDR block from which an IP address is allocated to the virtual NIC of the client. Do not enter the private CIDR block of the client. When the client accesses the destination network through an SSL-VPN connection, the VPN gateway allocates an IP address from the client CIDR block to the client.</p> |
| Advanced Configurations | <p>Select the type of advanced configurations.</p> <ul style="list-style-type: none"> ◦ Default: Use the default advanced configurations. ◦ Custom: Use custom configurations. You can set the following parameters: <ul style="list-style-type: none"> ▪ Protocol: Select the protocol used by the SSL-VPN connection. Valid values: UDP and TCP. ▪ Port: Specify the port used by the SSL-VPN connection. Invalid values: 22, 2222, 22222, 9000, 9001, 9002, 7505, 80, 443, 53, 68, 123, 4510, 4560, 500, and 4500. ▪ Encryption Algorithm: Select the encryption algorithm used by the SSL connection. Valid values: AES-128-CBC, AES-192-CBC, AES-256-CBC, and none. ▪ Enable Compression: Specify whether to compress the data that is transmitted over the SSL-VPN connection. |

19.8.2.2. Modify an SSL server

This topic describes how to modify the name, server CIDR block, client CIDR block, and advanced settings of an SSL server.

Prerequisites

An SSL server is created. For more information, see [Create an SSL server](#).

Procedure

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > SSL Servers**.
3. In the top navigation bar, select the region where you want to create the SSL server.
4. In the top navigation bar, select the region where the SSL server is created.
5. On the **SSL Servers** page, find the SSL server, and then click **Edit** in the **Actions** column.
6. On the **Edit SSL Server** page, modify the name, server CIDR block, client CIDR block, and advanced settings of the SSL server, and then click **OK**.

19.8.2.3. Configure a security group

This topic describes how to configure a security group to control the inbound and outbound traffic of ECS instances in the security group after an IPsec-VPN connection is created.

Procedure

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > SSL Servers**.
3. Select the region where the SSL server is deployed.
4. On the **SSL Servers** page, find the SSL server you want to manage, and then click **Configure Routing Group** in the **Actions** column.
5. In the **Configure Routing Group** panel, set the following parameters, and then click **Submit**.

| Parameter | Description |
|-----------------------------|---|
| Security Group | Select the security group to which you want to add the security group rule. |
| Rule Direction | Select the direction in which the rule applies. <ul style="list-style-type: none"> ◦ Outbound: from the ECS instances in the current security group to other ECS instances on Alibaba Cloud or resources on the Internet. ◦ Inbound: from other ECS instances on Alibaba Cloud or resources on the Internet to the ECS instances in the security group. |
| Policy | Select an authorization policy. <ul style="list-style-type: none"> ◦ Allow: allows access requests on the specified ports. ◦ Deny: discards requests received on the specified ports without returning messages. <p>If you specify different authorization policies for two security group rules but the other settings are the same, the Deny rule prevails over the Allow rule.</p> |
| Protocol Type | Select a protocol type. |
| Port Range | Specify a port range for the security group rule. Specify the port range in the format of 1/200 or 80/80. <ul style="list-style-type: none"> ◦ 1/200 specifies a port range from 1 to 200. ◦ 80/80 specifies port 80. <p>Valid values: -1 and 1 to 65535. A value of -1/-1 specifies all ports.</p> |
| Priority | Set a priority for the rule. Valid values: 1 to 100. The default value is 1, which specifies the highest priority. |
| Authorization Type | Select the authorization type of the security group rule. You can select only Address . |
| Network Type | Select the network interface controller (NIC) type. <ul style="list-style-type: none"> ◦ Internal: controls inbound and outbound traffic within Alibaba Cloud. ◦ External: controls inbound and outbound traffic over the Internet. |
| Authorization Object | Select the CIDR blocks to be authorized by the security group rule. You can specify up to 10 CIDR blocks at a time. |

| Parameter | Description |
|-------------|---|
| Description | The description of the security group rule. The description must be 2 to 256 characters in length, and cannot start with <code>http://</code> or <code>https://</code> . You can leave this parameter empty. |

19.8.2.4. Delete an SSL server

This topic describes how to delete an SSL server.

Procedure

1. [Log on to the VPN Gateway console.](#)
2. In the left-side navigation pane, choose **VPN > SSL Servers**.
3. In the top navigation bar, select the region where you want to delete the SSL server.
4. On the **SSL Servers** page, find the SSL server that you want to delete, and click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

19.8.3. Manage an SSL client certificate

19.8.3.1. Create an SSL client certificate

After you create an SSL server, you must create an SSL client certificate based on the configuration of the SSL server.

Prerequisites

An SSL server is created. For more information, see [Create an SSL server](#).

Procedure

1. [Log on to the VPN Gateway console.](#)
2. In the left-side navigation pane, choose **VPN > SSL Clients**.
3. Select the region where the SSL client is deployed.
4. On the **SSL Clients** page, click **Create Client Certificate**.
5. On the **Create SSL Client Certificate** page, configure the client certificate based on the following information, and then click **Submit**.

| Parameter | Description |
|---------------------|--|
| Organization | Select the organization to which the SSL client belongs. |
| Resource Set | Select the resource set to which the SSL client belongs. |
| Region: | Select the region where the SSL client is deployed. |
| Zone | Select the zone where the SSL client is deployed. |

| Parameter | Description |
|-------------|---|
| Name | Enter a name for the SSL client certificate. The name must be 2 to 128 characters in length, and can contain digits, hyphens (-), and underscores (_). It must start with a letter and cannot start with <code>http://</code> or <code>https://</code> . |
| VPN Gateway | Select the VPN gateway that you want to associate with the SSL client certificate. |
| SSL Server | Select the SSL server that you want to associate with the SSL client certificate. |

19.8.3.2. Download an SSL client certificate

This topic describes how to download an SSL client certificate. Before you use an SSL client to initiate an SSL-VPN connection, you must import the SSL client certificate to the SSL client.

Prerequisites

An SSL client certificate is created. For information, see [Create an SSL client certificate](#).

Procedure

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > SSL Clients**.
3. Select the region where the SSL client is deployed.
4. On the **SSL Clients** page, find the target SSL client certificate, and then click **Download** in the **Actions** column.

19.8.3.3. Delete an SSL client certificate

This topic describes how to delete an SSL client certificate.

Procedure

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > SSL Clients**.
3. Select the region where the SSL client is deployed.
4. On the **SSL Clients** page, find the SSL client certificate that you want to delete, and click **Delete** in the **Actions** column.
5. In the message that appears, click **OK**.

19.8.4. Query SSL-VPN connection logs

This topic describes how to query the logs of an SSL server and an SSL client. To troubleshoot the issues in an SSL-VPN connection, you can query the logs of the SSL server and the SSL client to which the connection is established.

Context

You can query logs that are created within the most recent month. The time range that you can specify for a query cannot exceed 10 minutes.

Query the logs of an SSL server.

1. [Log on to the VPN Gateway console](#).

2. In the left-side navigation pane, choose **VPN > SSL Servers**.
3. In the top navigation bar, select the region where you want to create the SSL server.
4. On the **SSL Servers** page, find the SSL server that you want to manage and click **View Logs** in the **Actions** column.
5. On the **SSL VPN Connection Logs** page, specify a time range. The system queries logs that are created during the specified time period.

Query the logs of an SSL client

1. [Log on to the VPN Gateway console](#).
2. In the left-side navigation pane, choose **VPN > SSL Clients**.
3. Select the region where the SSL client is deployed.
4. On the **SSL Clients** page, find the SSL client certificate that you want to manage and click **View Logs** in the **Actions** column.
5. On the **SSL-VPN Client Logs** page, specify a time range. The system queries logs that are created during the specified time period.

20. Elastic IP Address

20.1. What is Elastic IP Address?

This topic provides an overview of Elastic IP Address. An elastic IP address (EIP) is a public IP address that you can purchase and use as an independent resource. You can associate an EIP with an Elastic Compute Service (ECS) instance, an internal-facing Server Load Balancer (SLB) instance, or a secondary elastic network interface (ENI) deployed in a virtual private cloud (VPC). You can also associate an EIP with a NAT gateway, or a high-availability virtual IP address (HAVIP).

An EIP is a NAT IP address provisioned in the Internet-facing gateway of Alibaba Cloud and is mapped to the associated cloud resource by using NAT. After an EIP is associated with a cloud resource, the cloud resource can use the EIP to communicate with the Internet.

Differences between an EIP and the static public IP address of an ECS instance

The following table describes the differences between an EIP and the static public IP address of an ECS instance.

| Item | EIP | Static public IP address |
|---|----------------|--------------------------|
| Supported network | VPC | VPC |
| Used as an independent resource | Supported | Not supported |
| Associated with and disassociated from an ECS instance at any time | Supported | Not supported |
| Displayed in the ENI information of the operating system of the associated ECS instance | Not displayed. | Not displayed |

Benefits

EIPs have the following benefits:

- Purchase and use as independent resources
You can purchase and use an EIP as an independent resource. EIPs are not bundled with other computing or storage resources.
- Associate with resources at any time
You can associate an EIP with a cloud resource as needed. You can also disassociate and release the EIP at any time.
- Modify bandwidth limits on demand
You can modify the bandwidth limit of an EIP at any time to meet your business requirements. The modification immediately takes effect.

20.2. Log on to the EIP console

This topic describes how to log on to the Apsara Uni-manager Management Console to manage your elastic IP addresses (EIPs). The Google Chrome browser is used as an example.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.

- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

 **Note** When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top menu bar, choose **Products > Networking > Elastic IP Address**.

20.3. Quick start

20.3.1. Tutorial overview

This topic provides an overview of the tutorial that guides you in creating an EIP and associating an EIP with an ECS instance to allow the ECS instance to access the Internet.

This tutorial walks you through the following tasks:

1. [Apply for new EIPs](#)

An EIP is a public IP address that you can purchase and hold as an independent resource. To get started, you must create an EIP.

2. [Associate an EIP with an ECS instance](#)

You can associate an EIP with an ECS instance deployed in a VPC to enable the ECS instance to connect to the Internet.

3. [Disassociate an EIP from a cloud resource](#)

You can disassociate an ECS instance from an EIP when the ECS instance no longer requires access to the Internet.

4. [Release an EIP](#)

You can release an EIP if it is no longer needed.

20.3.2. Apply for EIPs

This topic describes how to apply for elastic IP addresses (EIPs). An EIP is a public IP address that you can purchase and hold as an independent resource.

Procedure

1. [Log on to the EIP console](#).
2. On the **Elastic IP Addresses** page, click **Create EIP**.
3. On the **Create Elastic IP Address** page, set the following parameters and click **Submit**.

| Parameter | Description |
|--------------------------|--|
| Organization | Select the organization to which the EIP belongs. |
| Resource Set | Select the resource set to which the EIP belongs. |
| Region | Select the region where you want to create the EIP. Make sure that the EIP and the cloud resources to be associated with the EIP are deployed in the same region. |
| Quantity | Enter the number of EIPs that you want to create. |
| EIP Name | Enter the name of the EIP that you want to create. |
| Connection Type | Select a line type for the EIP. |
| Network Type | Select a network type for the EIP. Valid values: <ul style="list-style-type: none"> ◦ Internet: The EIP is used for communication over the Internet. ◦ Hybrid Cloud: The EIP is used for communication within a hybrid cloud. If you want to allow a data center to access the Internet by using SNAT and DNAT, you must select this type. |
| IP Address | Enter the EIP that you want to request. Make sure that you enter an idle IPv4 address. Otherwise, you cannot request the specified EIP. <div style="background-color: #e0f2f1; padding: 5px; border: 1px solid #ccc;"> ? Note If you do not specify an EIP, the system automatically assigns one. </div> |
| Maximum Bandwidth | Specify the maximum bandwidth value for the EIP. Unit: Mbit/s. |

20.3.3. Associate an EIP with an ECS instance

This topic describes how to associate an elastic IP address (EIP) with an Elastic Compute Service (ECS) instance that is deployed in a virtual private cloud (VPC). ECS instances can communicate with the Internet after they are associated with EIPs.

Prerequisites

An ECS instance is created. For more information, see the [Create an instance](#) topic in the [Quick start](#) chapter of *Elastic Compute Service User Guide*.

Procedure

1. [Log on to the EIP console](#).
2. In the top navigation bar, select the region where the EIP is deployed.
3. On the [Elastic IP Addresses](#) page, find the EIP that you want to manage and click **Bind Resource** in the **Actions** column.
4. In the **Bind Elastic IP Address to Resources** dialog box, set the following parameters and click **OK**.

| Parameter | Description |
|----------------------|------------------------------|
| Instance Type | Select ECS Instance . |

| Parameter | Description |
|-----------------------------------|---|
| Binding mode | <p>Select the mode in which you want to associate the EIP.</p> <p>You can select only Normal. In NAT Mode:</p> <ul style="list-style-type: none"> ◦ The EIP is associated with the ECS instance in NAT mode. Both the private IP address and public IP address of the ECS instance are available. ◦ The EIP is not displayed in the operating system. To query the public IP address of the ECS instance, call the DescribeInstances operation. ◦ The EIP does not support NAT application layer gateway (ALG) protocols such as H.323, Session Initiation Protocol (SIP), Domain Name System (DNS), Real-Time Streaming Protocol (RTSP), and Trivial File Transfer Protocol (TFTP). |
| Select an instance to bind | <p>Select the ECS instance to be associated with the EIP.</p> <p>Make sure that the following requirements are met:</p> <ul style="list-style-type: none"> ◦ The ECS instance is deployed in a VPC. ◦ The ECS instance is in the Running or Stopped state. ◦ Each ECS instance can be associated with only one EIP. ◦ The ECS instance and the EIP reside in the same region. ◦ The ECS instance is not assigned a static public IP address. In addition, the ECS instance is not associated with another EIP. |

20.3.4. Disassociate an EIP from a cloud resource

This topic describes how to disassociate an elastic IP address (EIP) from a cloud resource. After an EIP is disassociated from a cloud resource, the cloud resource can no longer communicate with the Internet by using the EIP.

Procedure

1. [Log on to the EIP console](#).
2. In the top navigation bar, select the region where the EIP is deployed.
3. On the **Elastic IP Addresses** page, find the EIP that you want to manage and click **Unbind** in the **Actions** column.
4. In the message that appears, click **OK**.

20.3.5. Release an EIP

This topic describes how to release an elastic IP address (EIP) that you no longer need.

Prerequisites

The EIP is not associated with a cloud resource. For more information, see [Disassociate an EIP from a cloud resource](#).

Procedure

1. [Log on to the EIP console](#).
2. In the top navigation bar, select the region where the EIP is deployed.
3. On the **Elastic IP Addresses** page, find the EIP that you want to manage and choose  **> Release** in the **Actions** column.
4. In the **Release EIP** dialog box, click **OK**.

20.4. Manage EIPs

20.4.1. Apply for EIPs

This topic describes how to apply for elastic IP addresses (EIPs). An EIP is a public IP address that you can purchase and hold as an independent resource.

Procedure

1. [Log on to the EIP console](#).
2. On the **Elastic IP Addresses** page, click **Create EIP**.
3. On the **Create Elastic IP Address** page, set the following parameters and click **Submit**.

| Parameter | Description |
|--------------------------|--|
| Organization | Select the organization to which the EIP belongs. |
| Resource Set | Select the resource set to which the EIP belongs. |
| Region | Select the region where you want to create the EIP. Make sure that the EIP and the cloud resources to be associated with the EIP are deployed in the same region. |
| Quantity | Enter the number of EIPs that you want to create. |
| EIP Name | Enter the name of the EIP that you want to create. |
| Connection Type | Select a line type for the EIP. |
| Network Type | Select a network type for the EIP. Valid values: <ul style="list-style-type: none"> ◦ Internet: The EIP is used for communication over the Internet. ◦ Hybrid Cloud: The EIP is used for communication within a hybrid cloud. If you want to allow a data center to access the Internet by using SNAT and DNAT, you must select this type. |
| IP Address | Enter the EIP that you want to request. Make sure that you enter an idle IPv4 address. Otherwise, you cannot request the specified EIP. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> ? Note If you do not specify an EIP, the system automatically assigns one. </div> |
| Maximum Bandwidth | Specify the maximum bandwidth value for the EIP. Unit: Mbit/s. |

20.4.2. Bind an EIP to a cloud instance

20.4.2.1. Associate an EIP with an ECS instance

This topic describes how to associate an elastic IP address (EIP) with an Elastic Compute Service (ECS) instance that is deployed in a virtual private cloud (VPC). ECS instances can communicate with the Internet after they are associated with EIPs.

Prerequisites

An ECS instance is created. For more information, see the **Create an instance** topic in the **Quick start** chapter of *Elastic Compute Service User Guide*.

Procedure

1. [Log on to the EIP console](#).
2. In the top navigation bar, select the region where the EIP is deployed.
3. On the **Elastic IP Addresses** page, find the EIP that you want to manage and click **Bind Resource** in the **Actions** column.
4. In the **Bind Elastic IP Address to Resources** dialog box, set the following parameters and click **OK**.

| Parameter | Description |
|-----------------------------------|---|
| Instance Type | Select ECS Instance . |
| Binding mode | <p>Select the mode in which you want to associate the EIP.</p> <p>You can select only Normal. In NAT Mode:</p> <ul style="list-style-type: none"> ◦ The EIP is associated with the ECS instance in NAT mode. Both the private IP address and public IP address of the ECS instance are available. ◦ The EIP is not displayed in the operating system. To query the public IP address of the ECS instance, call the DescribeInstances operation. ◦ The EIP does not support NAT application layer gateway (ALG) protocols such as H.323, Session Initiation Protocol (SIP), Domain Name System (DNS), Real-Time Streaming Protocol (RTSP), and Trivial File Transfer Protocol (TFTP). |
| Select an instance to bind | <p>Select the ECS instance to be associated with the EIP.</p> <p>Make sure that the following requirements are met:</p> <ul style="list-style-type: none"> ◦ The ECS instance is deployed in a VPC. ◦ The ECS instance is in the Running or Stopped state. ◦ Each ECS instance can be associated with only one EIP. ◦ The ECS instance and the EIP reside in the same region. ◦ The ECS instance is not assigned a static public IP address. In addition, the ECS instance is not associated with another EIP. |

20.4.2.2. Associate an EIP with an SLB instance

This topic describes how to associate an elastic IP address (EIP) with a Server Load Balancer (SLB) instance. After you associate an EIP with an SLB instance, the SLB instance can forward requests from the Internet.

Prerequisites

An SLB instance is created. For more information, see the **Create an SLB instance** topic in the **Quick start** chapter of *Server Load Balancer User Guide*.

Procedure

1. [Log on to the EIP console](#).
2. In the top navigation bar, select the region where the EIP is deployed.
3. On the **Elastic IP Addresses** page, find the EIP that you want to manage and click **Bind Resource** in the **Actions** column.

4. In the **Bind Elastic IP Address to Resources** dialog box, set the following parameters and click **OK**.

| Parameter | Description |
|----------------------------|---|
| Instance Type | Select SLB Instance . |
| Select an instance to bind | <p>Select the SLB instance to be associated with the EIP.</p> <p>Make sure that the following requirements are met:</p> <ul style="list-style-type: none"> ◦ The SLB instance is deployed in a virtual private cloud (VPC). ◦ The SLB instance and the EIP reside in the same region. ◦ Each SLB instance can be associated with only one EIP. <p>Note If you associate the EIP with an internal-facing SLB instance and the SLB instance has traffic of private network workloads, transient connections may occur. We recommend that you perform the association during off-peak hours or switch the workloads to another SLB instance.</p> |

20.4.2.3. Associate an EIP with a NAT gateway

This topic describes how to associate an elastic IP address (EIP) with a NAT gateway. After you associate an EIP with a NAT gateway, you can create DNAT and SNAT entries for the EIP.

Prerequisites

A NAT gateway is created. For more information, see the **Create a NAT gateway** topic in the **Quick start** chapter of *NAT Gateway User Guide*.

Procedure

1. [Log on to the EIP console](#).
2. In the top navigation bar, select the region where the EIP is deployed.
3. On the **Elastic IP Addresses** page, find the EIP that you want to manage and click **Bind Resource** in the **Actions** column.
4. In the **Bind Elastic IP Address to Resources** dialog box, set the following parameters and click **OK**.

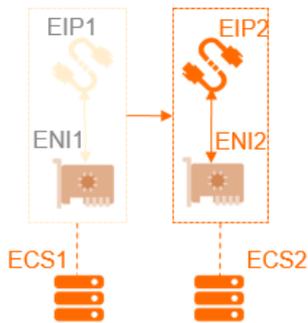
| Parameter | Description |
|----------------------------|---|
| Instance Type | Select NAT Gateway . |
| Select an instance to bind | <p>Select the NAT gateway to be associated with the EIP.</p> <p>Make sure that the following requirements are met:</p> <ul style="list-style-type: none"> ◦ The NAT gateway and the EIP reside in the same region. ◦ Each NAT gateway can be associated with at most 20 EIPs, among which at most 10 pay-by-data-transfer EIPs can be associated. |

20.4.2.4. Bind an EIP to a secondary ENI

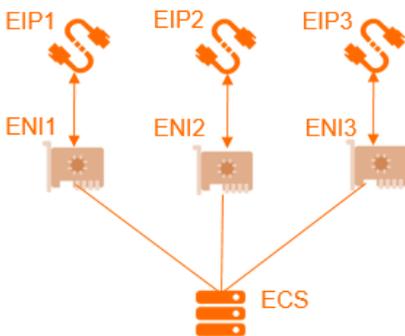
20.4.2.4.1. Overview

You can associate elastic IP addresses (EIPs) with elastic network interfaces (ENIs). Then, you can associate the ENIs with an Elastic Compute Service (ECS) instance. This way, the ECS instance can use multiple EIPs. You can use EIPs to improve the service availability, flexibility, and scalability.

Each ENI is assigned a private IP address. After you associate an EIP with an ENI, both the private IP address and the EIP are available for the ENI. You can change the private IP address and public IP address of an ECS instance by replacing the secondary ENI that is associated with the ECS instance. When you replace the secondary ENI of an ECS instance, the reliability and availability of your service are not affected.



You can associate multiple ENIs with an ECS instance. Make sure that an EIP is associated with each ENI. This way, the ECS instance can use multiple EIPs. The ECS instance can use the EIPs to provide Internet-facing services. You can configure security group rules for the ECS instance to control access from the Internet.



Association modes

You can associate an EIP with an ENI in NAT mode.

The following table describes the features in NAT mode.

| Item | NAT mode |
|---|---|
| Whether the EIP is displayed in the ENI information of the operating system | No |
| Types of ENIs that can be associated with EIPs | Primary and secondary ENIs |
| Number of EIPs that can be associated with a primary ENI | 1 |
| Number of EIPs that can be associated with a secondary ENI | Depends on the number of private IP addresses of the secondary ENI <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p>Note Each EIP is mapped to a private IP address of a secondary ENI. If a secondary ENI is assigned 10 private IP addresses, at most 10 EIPs can be associated with the secondary ENI.</p> </div> |
| Whether the private network features of a secondary ENI are available after an EIP is associated with the secondary ENI | Yes |

| Item | NAT mode |
|---------------------|--|
| Supported protocols | EIPs deployed as NAT application layer gateways (ALGs) do not support protocols such as H.323, Session Initiation Protocol (SIP), Domain Name System (DNS), and Real-Time Streaming Protocol (RTSP). |

20.4.2.4.2. Associate an EIP with a secondary ENI in NAT mode

This topic describes how to associate an elastic IP address (EIP) with a secondary elastic network interface (ENI) in NAT mode. After you associate an EIP with a secondary ENI, both the private IP address and public IP address that are assigned to the secondary ENI are available. In this case, the EIP is not displayed in the secondary ENI information.

Prerequisites

Before you associate an EIP with a secondary ENI in NAT mode, make sure that the following requirements are met:

- A secondary ENI that is deployed in a virtual private cloud (VPC) is created. The secondary ENI and the EIP reside in the same region. For more information, see the **Create an ENI** topic in the **Elastic network interfaces** chapter of *Elastic Compute Service User Guide*.
- The secondary ENI is not associated with an Elastic Compute Service (ECS) instance. If the secondary ENI is associated with an ECS instance, disassociate the secondary ENI from the ECS instance first. Then, associate the EIP with the secondary ENI in NAT mode and associate the secondary ENI with the ECS instance. For more information, see the **Unbind a secondary ENI from an instance** topic in the **Elastic network interfaces** chapter of *Elastic Compute Service User Guide*.

Procedure

1. [Log on to the EIP console](#).
2. In the top navigation bar, select the region where the EIP is deployed.
3. On the **Elastic IP Addresses** page, find the EIP that you want to manage and click **Bind Resource** in the **Actions** column.
4. In the **Bind Elastic IP Address to Resources** dialog box, set the following parameters and click **OK**.

| Parameter | Description |
|----------------------|---|
| Instance Type | Select Secondary ENI . |
| Binding mode | <p>Select Normal.</p> <p>In NAT Mode:</p> <ul style="list-style-type: none"> ◦ The number of EIPs that can be associated with a secondary ENI is based on the number of private IP addresses that are assigned to the secondary ENI. ◦ When an EIP is associated with a secondary ENI in NAT mode, both the private IP address and public IP address that are assigned to the secondary ENI are available. ◦ The EIP is not displayed in the operating system. To query the EIP, call the DescribeEipAddresses operation. ◦ The EIP does not support NAT application layer gateway (ALG) protocols such as H.323, Session Initiation Protocol (SIP), Domain Name System (DNS), Real-Time Streaming Protocol (RTSP), and Trivial File Transfer Protocol (TFTP). |

| Parameter | Description |
|----------------------------|---|
| Select an instance to bind | <p>Select the secondary ENI to be associated with the EIP.</p> <p>Make sure that the following requirements are met:</p> <ul style="list-style-type: none">○ The secondary ENI is deployed in a VPC.○ The secondary ENI and the EIP reside in the same region. |

20.4.3. Resize the maximum bandwidth

This topic describes how to resize the maximum bandwidth for elastic IP addresses (EIPs). After you change the maximum bandwidth value for an EIP, the new value immediately takes effect.

Procedure

1. [Log on to the EIP console](#).
2. In the top navigation bar, select the region where the EIP is deployed.
3. On the **Elastic IP Addresses** page, find the EIP that you want to manage and choose  > **Modify Configuration** in the **Actions** column.
4. On the **Change Specifications** page, specify a new bandwidth value and click **Submit**.

20.4.4. Disassociate an EIP from a cloud resource

This topic describes how to disassociate an elastic IP address (EIP) from a cloud resource. After an EIP is disassociated from a cloud resource, the cloud resource can no longer communicate with the Internet by using the EIP.

Procedure

1. [Log on to the EIP console](#).
2. In the top navigation bar, select the region where the EIP is deployed.
3. On the **Elastic IP Addresses** page, find the EIP that you want to manage and click **Unbind** in the **Actions** column.
4. In the message that appears, click **OK**.

20.4.5. Release an EIP

This topic describes how to release an elastic IP address (EIP) that you no longer need.

Prerequisites

The EIP is not associated with a cloud resource. For more information, see [Disassociate an EIP from a cloud resource](#).

Procedure

1. [Log on to the EIP console](#).
2. In the top navigation bar, select the region where the EIP is deployed.
3. On the **Elastic IP Addresses** page, find the EIP that you want to manage and choose  > **Release** in the **Actions** column.
4. In the **Release EIP** dialog box, click **OK**.

21. Apsara Stack Security

21.1. What is Apsara Stack Security

Apsara Stack Security is a solution that protects Apsara Stack assets with a full suite of security features, such as network, server, application, data, and security management.

Background information

Traditional security solutions for IT services detect attacks on network perimeters. These solutions use hardware products such as firewalls and intrusion prevention systems (IPSs) to protect networks against attacks.

With the development of cloud computing, an increasing number of enterprises and organizations use cloud computing services instead of traditional IT services. Cloud computing features low costs, on-demand flexible configuration, and high resource utilization. Cloud computing environments do not have definite network perimeters. As a result, traditional security solutions cannot effectively safeguard cloud assets.

With the powerful data analysis capabilities and professional security operations team of Alibaba Cloud, Apsara Stack Security provides integrated security protection services for networks, applications, and servers.

Complete security solution

Apsara Stack Security consists of Apsara Stack Security Standard Edition and optional security services to provide a comprehensive security solution.

| Security domain | Service name | Description |
|----------------------|--|---|
| Security management | Threat Detection Service (TDS) | Monitors traffic and overall security status to audit and centrally manage assets. |
| Server security | Server Guard | Protects Elastic Compute Service (ECS) instances against intrusions and malicious code. |
| | Server Security | Protects physical servers against intrusions. |
| Application security | Web Application Firewall (WAF) | Protects web applications against attacks and ensures that mobile and PC users can securely access web applications over the Internet. |
| Network security | Anti-DDoS | Ensures the availability of network links and improves business continuity. |
| | Cloud Firewall | Allows you to centrally manage access control policies for traffic transferred within your business system (east-west) and between the Internet and your business system (north-south). |
| Data security | Sensitive Data Discovery and Protection (SDDP) | Prevents data leaks and helps your business system meet compliance requirements. |
| Security O&M service | On-premises security service | Helps you establish and optimize the cloud security system to protect your business system against attacks by using security features of Apsara Stack Security and other Apsara Stack services. |

21.2. Usage notes

Before you log on to Apsara Stack Security Center, you must verify that your computer meets the configuration requirements.

For more information about the configuration requirements, see [Configuration requirements](#).

Configuration requirements

| Item | Requirement |
|------------------|---|
| Browser | <ul style="list-style-type: none"> • Internet Explorer: V11 or later • Google Chrome (recommended): V42.0.0 or later • Mozilla Firefox: V30 or later • Safari: V9.0.2 or later • GmSSL browser that runs the Chrome kernel: V69.0.0 or later |
| Operating system | <ul style="list-style-type: none"> • Windows XP • Windows 7 or later • macOS |

21.3. Quick start

21.3.1. User roles and permissions

This topic describes the user roles involved in Apsara Stack Security.

All roles in Apsara Stack Security Center are provided by default. You cannot add custom roles. Before you log on to Apsara Stack Security Center, make sure that your account is assigned the required role. For more information, see [Default roles in Apsara Stack Security](#).

Default roles in Apsara Stack Security

| Role | Permission |
|--|--|
| System administrator of Apsara Stack Security Center | Manages and configures system settings for Apsara Stack Security Center. The system administrator has permissions to manage Apsara Stack accounts, synchronize data, configure alerts, and configure global settings. |
| Security administrator of Apsara Stack Security Center | <p>Monitors the security status across Apsara Stack and configures security policies for each functional module of Apsara Stack Security. The security administrator has permissions on all features under Threat Detection, Network Security, Application Security, Server Security, Physical Server Security, and Asset Management.</p> <p> Note The permissions on Web Application Firewall (WAF) and Cloud Firewall must be separately assigned.</p> |
| Department security administrator | <p>Monitors the security status of cloud resources in a specific department and configures security policies for each functional module of Apsara Stack Security for this department. The department security administrator has permissions on all features under Threat Detection, Network Security, Application Security, Server Security, Physical Server Security, and Asset Management. In addition, the department security administrator can specify alert notification methods and alert contacts in the department.</p> <p> Note The permissions on WAF and Cloud Firewall must be separately assigned.</p> |
| Auditor of Apsara Stack Security Center | Conducts security audits across Apsara Stack. The auditor can view audit events and raw logs, configure audit policies, and access all features under Security Audit. |

If you do not have an account that assumes the required role, contact the administrator to create an account and assign the role to the account. For more information, see the [Create a user](#) topic in *Apsara Uni-manager Management Console User Guide*.

21.3.2. Log on to Apsara Stack Security Center

This topic describes how to log on to Apsara Stack Security Center.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

 **Note** When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top navigation bar, choose **Security > Apsara Stack Security**.
5. On the **Apsara Stack Security Center** page, select **Region**.
6. Click **YD** to go to Apsara Stack Security Center.

21.4. Threat Detection Service

21.4.1. Overview

This topic introduces the basic concepts related to Threat Detection Service (TDS).

TDS provides comprehensive protection for enterprises. It can monitor vulnerabilities, intrusions, web attacks, DDoS attacks, threat intelligence, and public opinions. TDS uses modeling and analysis to obtain key information based on traffic characteristics, host behavior, and host operation logs. In addition, TDS identifies intrusions that cannot be detected by traffic inspection or file scan. You can use the input of cloud analysis models and intelligence data to discover sources and behavior of attacks and assess threats.

TDS provides the following features:

- **Overview**: provides a security situation overview and information about security screens.
- **Security Alerts**: displays security alerts that occur in your business system.
- **Attack Analysis**: displays application attacks and brute-force attacks that occur in your system.
- **Cloud Service Check**: checks whether risks exist in the configurations of Apsara Stack services.
- **Application Whitelist**: allows you to create and apply application whitelist policies to your servers that require

special protection. After you create the policies, Apsara Stack Security identifies trusted, suspicious, and malicious programs based on intelligent learning. This prevents unauthorized programs from running.

- **Assets:** manages servers and cloud services on Apsara Stack.
- **Security Reports:** allows you to configure security report tasks on Apsara Stack.

21.4.2. Security overview

21.4.2.1. View security overview information

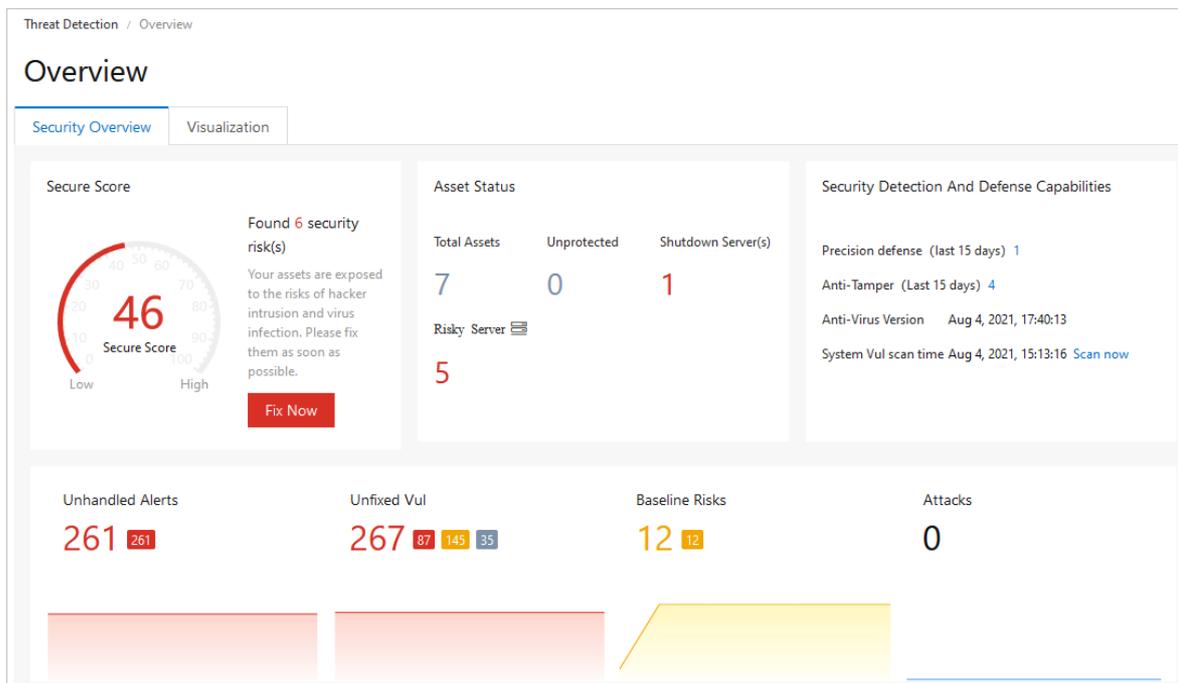
This topic describes how to view security statistics, attack trends, and network traffic information on the Apsara Stack platform.

Context

The **Security Overview** tab provides an overview of detected security events, the latest threats, and inherent vulnerabilities of the system. A security administrator can view information on the **Security Overview** tab to better understand the security posture of the system.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Threat Detection > Overview**. On the page that appears, click the **Security Overview** tab.
3. View the security posture of the Apsara Stack platform.



Sections on the Security Overview tab

| Section | Description |
|--------------|---|
| Secure Score | The security score of assets and the number of detected security risks. |
| Asset Status | The total number of assets and the numbers of servers that are not protected, servers that are stopped, and servers that are at risk. |

| Section | Description |
|---|---|
| Security Detection And Defense Capabilities | The numbers of precise defense events and anti-tampering events over the last 15 days, the time when the antivirus database was last updated, and the time when vulnerability scanning was last performed. This allows you to obtain the defense situation and security status of your assets in real time. |
| Threat statistics | The numbers of alerts that are not handled, vulnerabilities that are not fixed, baseline risks, and attacks. |
| Config Assessment Risks | Risks in the baseline configurations of cloud services. |
| Issue Resolved | Statistics on alerts, vulnerabilities, and baseline risks that have been processed over the last 15 days. The statistics are displayed in a bar and trend chart. |

21.4.3. Security alerts

21.4.3.1. View security alerts

This topic describes how to view security alerts on the Security Alerts page.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Threat Detection > Security Alerts**.
3. (Optional)Set filter conditions for security alerts.

? **Note** If you want to view all alerts, do not set the conditions.

Ur... X
No... X
War... X
▼
Unhandle... ▼
All ▼
Asset Group ▼
Alert/Asset
Q

| Filter condition | Description |
|------------------------------------|--|
| Severity | The severity level. You can select one or more levels. Valid values: <ul style="list-style-type: none"> ○ Urgent ○ Warning ○ Notice |
| Alert status | The status of alerts. Valid values: <ul style="list-style-type: none"> ○ Unhandled Alerts ○ Handled |
| Alert type | The type of alerts. Select All or a specific type. |
| Affected asset group | The affected asset group. Select Asset Group or a specific group. |
| Search for alerts by name or asset | Enter an alert name or a keyword of affected assets to search for alerts. |

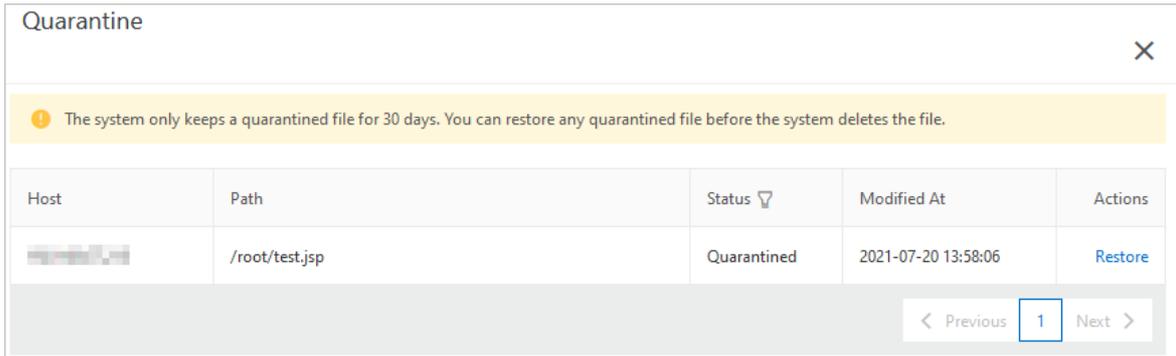
4. View security alerts and their details in the list.

21.4.3.2. Manage quarantined files

This topic describes how to manage threat files that are quarantined by the system. The system deletes a quarantined file 30 days after the file is quarantined. You can restore the file before it is deleted.

Procedure

1. Log on to Apsara Stack Security Center.
2. In the left-side navigation pane, choose Threat Detection > Security Alerts.
3. In the upper-right corner of the Alerts page, click Quarantine.
4. In the Quarantine panel, view the information about a quarantined file, such as the IP address of the host, path, status, and operation time.



5. (Optional) If a file is incorrectly quarantined, click Restore in the Actions column to restore the file.

 **Notice** Before you restore a quarantined file, make sure that the file is normal and does not bring risks.

The restored file is removed from the Quarantine panel and is displayed in the security alert list again.

21.4.3.3. Configure security alerts

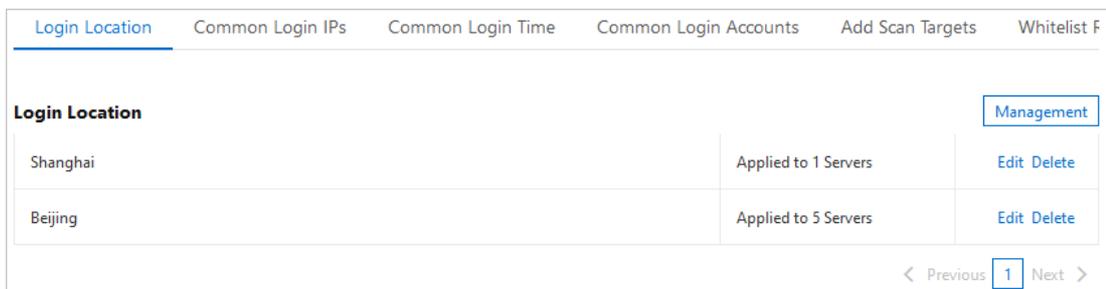
This topic describes how to configure logon settings and web directories that are scanned.

Procedure

1. Log on to Apsara Stack Security Center.
2. In the left-side navigation pane, choose Threat Detection > Security Alerts.
3. In the upper-right corner of the page that appears, click Settings.

You can perform the following operations:

- o Add an approved logon location
 - a. In the Login Location section, click Management on the right.



- b. Select the logon location that you want to add, and select the servers that allow logons from the added location.

- c. Click **Ok**.

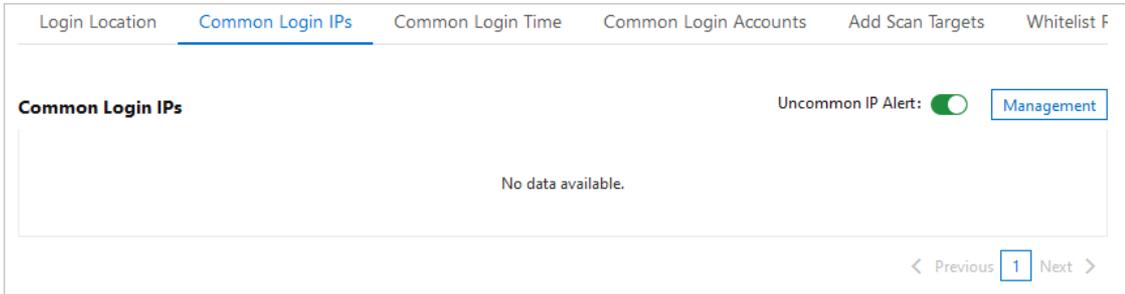
Threat Detection Service (TDS) allows you to edit and delete added logon locations.

- Find the required logon location and click **Edit** on the right to change the servers that are allowed to be logged on to from this location.
- Find the required logon location and click **Delete** on the right to delete the logon location.

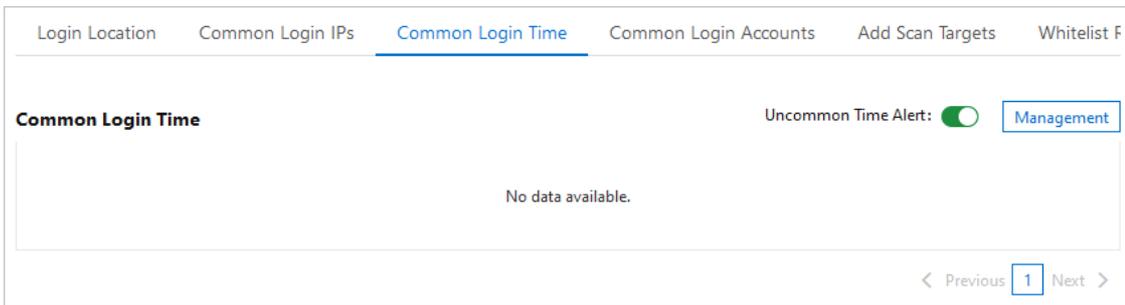
o **Configure advanced logon settings**

Note You can specify the IP addresses, accounts, and time periods that are allowed to log on to your assets. After you configure these settings, alerts are triggered if your assets receive logon requests that do not meet the requirements. The procedure to configure advanced logon settings is similar to that to configure **common logon locations**. You can follow the preceding procedure to **add**, **edit**, and **delete** advanced logon settings.

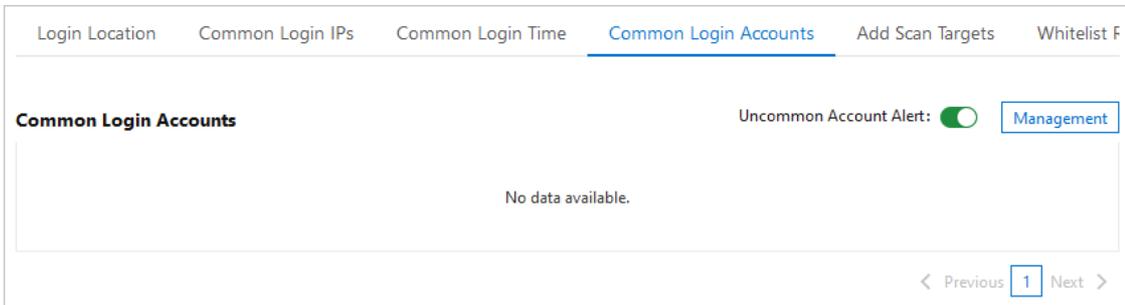
- On the right of **Common Login IPs**, turn on or off Uncommon IP Alert. After the switch is turned on, alerts are triggered if your assets receive logon requests from unapproved IP addresses.



- On the right of **Common Login Time**, turn on or off Uncommon Time Alert. After the switch is turned on, alerts are triggered if your assets receive logon requests at an unapproved time.

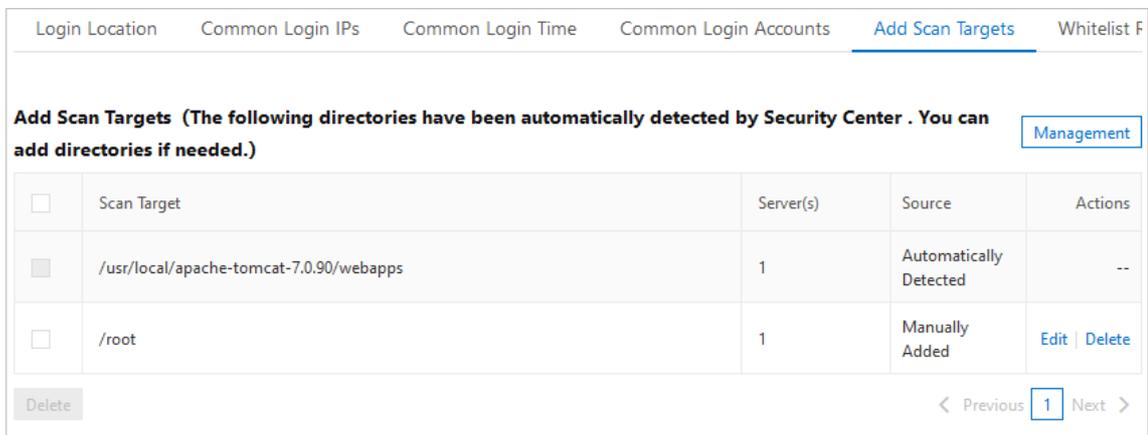


- On the right of **Common Login Accounts**, turn on or off Uncommon Account Alert. After the switch is turned on, alerts are triggered if your assets receive logon requests from unapproved accounts.



○ **Add Scan Targets**

Apsara Stack Security automatically scans web directories of data assets in your servers and runs dynamic and static scan tasks. You can also manually add other web directories of your servers.



- On the right of **Add Scan Targets**, click **Add**.

- b. Enter a valid web directory and select the servers on which the directory is scanned. The web directory is added to the scan list.

 **Note** Root directories are not allowed. This ensures performance and efficiency.

- c. Click **Ok**.

21.4.4. Attack analysis

This topic describes the statistics provided by the attack analysis feature. The statistics include the total number of attacks, distribution of attack types, top five attack sources, top five attacked assets, and an attack list.

Background information

The attack analysis feature provides basic attack detection and prevention capabilities in Apsara Stack Security Center. We recommend that you optimize firewalls and enhance business security to develop a more fine-grained and in-depth defense system.

On the **Attack Awareness** page, you can specify a time range to view these attack details. You can view the attack analysis statistics of the current day, last 7 days, or last 15 days. You can also set Time Range to **Custom** to view the statistics of a time range within the last 30 days.

- **Attacks:** the total number of attacks detected in your assets within a specific time range.
- **Attack Type Distribution:** the attack types and the number of attacks for each type.
- **Top 5 Attack Sources:** the top five IP addresses from which the most attacks are launched.
- **Top 5 Attack Assets:** the top five assets that are attacked the most frequently.
- **Attack list:** the details about each attack. The details include the attack time, source IP address, attacked asset, attack type, and total number of attacks.

 **Note** The attack list displays a maximum of 10,000 attacks. You can specify **Time Range** to view details about the attacks that occur over the specified time range.

Parameters in the attack list

| Parameter | Description |
|----------------|---|
| Attacked At | The time at which an attack occurs. |
| Attack Source | The source IP address of an attack. |
| Attacked Asset | The name, public IP address, and private IP address of an attacked asset. |
| Attacks | The total number of attacks. |
| Attack Type | The type of an attack. The types of attacks that can be detected include SSH brute-force attacks and remote code execution attacks. |

- Search for an attack.
To view the details about a specific attack, specify search conditions in the search box above the attack list. Search conditions include the attack type, attacked asset, and source IP address.
- View the details of an attacked asset.
To view the details about an attacked asset, move the pointer over the name of the attacked asset.
- Export the attack list.

To export and save the attack list to your computer, click the  icon in the upper-left corner above the attack list. The attack list is exported to an Excel file.

21.4.5. Cloud service check

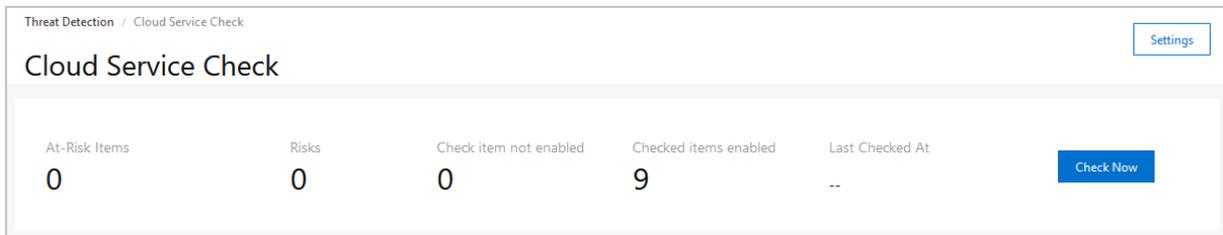
21.4.5.1. Overview

Threat Detection Service (TDS) provides the cloud service check feature. This feature allows you to check for the configuration risks of your Apsara Stack services. This topic describes the features and check items that are supported by the cloud service check feature.

Background information

The cloud service check feature allows you to perform network access control and data security checks. The checks help you detect configuration risks of your Apsara Stack services and provide repair solutions.

You can view the number of **Checked items enabled** on the **Cloud Service Check** page.



Cloud service check list

The following table describes the check items.

| Type | Supported item | Description |
|------|---|---|
| | PolarDB - Backup configurations | Checks whether the automatic backup feature is enabled for PolarDB. Regular backups help you improve database security. You can restore data if an error occurs in your database. PolarDB supports automatic backup. We recommend that you enable automatic backup to create a backup on a daily basis. |
| | Container Registry - Repository permission configurations | Checks whether a Container Registry repository is set to private. Container Registry supports public and private repositories. Public repositories allow users to anonymously download images over the Internet. If images in a repository contain sensitive information, we recommend that you set the repository to private. If images in a repository do not contain sensitive information, ignore related alerts. |
| | OSS - Server-side encryption | Checks whether the data encryption feature is enabled for Object Storage Service (OSS) buckets. OSS supports server-side encryption to secure data that is persistently stored in OSS. We recommend that you enable server-side encryption to protect sensitive data. |
| | OSS - Sensitive information leakage scans | Checks whether access permissions on sensitive files in OSS buckets are required. |

| Type | Supported item | Description |
|---|--|--|
| Data security | ApsaraDB RDS - Cross-region backup configurations | Checks whether the cross-region backup feature is enabled for ApsaraDB RDS instances. ApsaraDB RDS for MySQL provides the cross-region backup feature that automatically synchronizes local backup files to OSS buckets in another region. This implements geo-disaster recovery. We recommend that you enable the cross-region backup feature. |
| | KVStore for Redis - Backup configurations | Checks whether the data backup feature is enabled for KVStore for Redis instances. |
| | ApsaraDB for MongoDB - SSL encryption | Checks whether SSL encryption is enabled for ApsaraDB for MongoDB databases. We recommend that you enable the SSL encryption feature to improve the security of data links in ApsaraDB for MongoDB databases. |
| | ApsaraDB for MongoDB - Backup configurations | Checks whether the automatic backup feature is enabled for ApsaraDB for MongoDB databases. Regular backups help you improve database security. You can restore data if an error occurs in your database. ApsaraDB for MongoDB provides automatic backup policies. We recommend that you enable automatic backup to create a backup on a daily basis. |
| | ECS - Disk encryption | Checks whether encryption is enabled for disks on Elastic Compute Service (ECS) instances. |
| | ECS - Automatic snapshot policies | Checks whether the automatic snapshot feature is enabled for the disks on ECS instances. The automatic snapshot feature improves the security of ECS instances and supports disaster recovery. |
| | OSS - Bucket permissions | Checks whether the OSS bucket ACL is set to <i>private</i> . |
| | OSS - Logging | Checks whether the logging feature is enabled for OSS. |
| | OSS - Cross-region replication | Checks whether the cross-region replication feature is enabled for OSS. |
| | ApsaraDB RDS - Database security policies | Checks whether the SQL audit, SSL encrypted transmission, and transparent database encryption features are enabled for ApsaraDB RDS databases. |
| | ApsaraDB RDS - Backup configurations | Checks whether the data backup feature is enabled for ApsaraDB RDS instances. |
| SSL Certificates Service - Expiration check | Checks whether your SSL certificate expires. If your SSL certificate expires, you are not allowed to use SSL Certificates Service. | |
| | ECS - Security group policies | Checks the security group policies of ECS. We recommend that you grant permissions to users based on the principle of least privilege. We also recommend that you specify 0.0.0.0/0 only for the ports that must be open to all services, such as port 80, 443, 22, or 3389. |
| | OSS - Bucket hotlink protection | Checks whether the hotlink protection feature is enabled for OSS buckets. The OSS hotlink protection feature checks the Referer header to deny access from unauthorized users. We recommend that you enable this feature. |

| Type | Supported item | Description |
|------------------------|--|---|
| Network access control | VPC - DNAT rules | Checks whether a port is open to the Internet. When you create a DNAT rule for a NAT gateway that is deployed in a virtual private cloud (VPC), we recommend that you do not open internal management ports to the Internet. Do not open all ports or an important port, for example, ports 22, 3389, 1433, or 3306, to the Internet. |
| | Apsara Stack Security - Back-to-origin configuration for Anti-DDoS | Checks whether Anti-DDoS is configured to allow the requests from only Web Application Firewall (WAF) back-to-origin IP addresses. After you set up Anti-DDoS or WAF, you must hide the IP addresses of the backend servers to prevent attacks on the cloud assets. |
| | Apsara Stack Security - WAF back-to-origin configurations | Checks whether WAF allows requests from only WAF back-to-origin IP addresses. After you set up Anti-DDoS or WAF, you must hide the IP addresses of the backend servers to prevent attacks on the cloud assets. |
| | SLB - IP address whitelist configurations | Checks the access control configurations of Server Load Balancer (SLB) instances. Checks whether access control is enabled for HTTP and HTTPS services and checks whether 0.0.0.0/0 is added to the IP address whitelist. |
| | SLB - Open ports | Checks whether SLB opens ports to the Internet for forwarding unnecessary public services. |
| | ApsaraDB RDS - IP address whitelist configurations | Checks whether a whitelist is configured for ApsaraDB RDS and whether the whitelist contains 0.0.0.0/0. If the whitelist contains 0.0.0.0/0, requests from all IP addresses are allowed. We recommend that you configure the whitelist to allow requests only from specific IP addresses. |
| | KVStore for Redis - IP address whitelist configurations | Checks whether a whitelist is configured for KVStore for Redis and whether the whitelist contains 0.0.0.0/0. If the whitelist contains 0.0.0.0/0, requests from all IP addresses are allowed. We recommend that you configure the whitelist to allow requests only from specific IP addresses. |
| | AnalyticDB for PostgreSQL - IP address whitelist configurations | Checks whether a whitelist is configured for AnalyticDB for PostgreSQL and whether the whitelist contains 0.0.0.0/0. If the whitelist contains 0.0.0.0/0, requests from all IP addresses are allowed. We recommend that you configure the whitelist to allow requests only from specific IP addresses. |
| | PolarDB - IP address whitelist configurations | Checks whether a whitelist is configured for PolarDB and whether the whitelist contains 0.0.0.0/0. If the whitelist contains 0.0.0.0/0, requests from all IP addresses are allowed. We recommend that you configure the whitelist to allow requests only from specific IP addresses. |
| | ApsaraDB for MongoDB - IP address whitelist configurations | Checks whether a whitelist is configured for ApsaraDB for MongoDB and whether the whitelist contains 0.0.0.0/0. If the whitelist contains 0.0.0.0/0, requests from all IP addresses are allowed. We recommend that you configure the whitelist to allow requests only from specific IP addresses. |

21.4.5.2. Run cloud service checks

Threat Detection Service (TDS) provides the cloud service check feature. This feature allows you to check for security risks in the configurations of your cloud services. This topic describes how to manually run cloud service checks on your cloud services. This topic also describes how to specify a detection interval for periodic automatic checks.

Context

Apsara Stack Security supports manual checks and periodic automatic checks to scan for security risks in the configurations of cloud services.

- **Manual checks:** On the **Cloud Service Check** page, you can click **Check Now** to check for security risks in the configurations of your cloud services.
- **Periodic automatic checks:** By default, Apsara Stack Security automatically runs checks during the time range 00:00 - 06:00 every other day. You can also customize a time range for periodic automatic checks. This way, you can detect and handle the security risks in the configurations of your cloud services at the earliest opportunity.

Manual checks

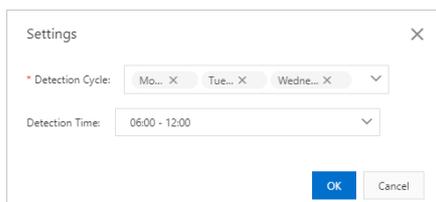
1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Threat Detection > Cloud Service Check**.
3. On the **Cloud Service Check** page, click **Check Now** to check whether the configurations of all your cloud services contain risks and the number of affected assets.

 **Note** Do not perform other operations until the check is complete.

After the check is complete, the detected risks are listed based on risk severities in descending order.

Automatic checks

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Threat Detection > Cloud Service Check**.
3. In the upper-right corner of the **Cloud Platform Configuration Assessment** page, click **Settings**.
4. In the **Settings** dialog box, configure the **Detection Cycle** and **Detection Time** parameters.



- **Detection Cycle:** Monday to Sunday. You can select multiple values.
 - **Detection Time:** Valid values are 24:00 - 06:00 , 06:00 - 12:00 , 12:00 - 18:00 , and 18:00 - 24:00 . You can select only one time range.
5. Click **Ok**.
During the selected time range, Apsara Stack Security automatically runs checks on all check items.

21.4.5.3. View the check results of configuration assessment for your cloud services and handle the detected risks

This topic describes how to view the check results of configuration assessment for your cloud services and handle the detected configuration risks in Apsara Stack Security. You can view the check items, details of check items, potential impacts caused by the detected configuration risks, and suggestions on how to handle the detected configuration risks. You can handle the detected configuration risks on the Cloud Service Check page in a centralized manner.

View check results

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Threat Detection > Cloud Service Check**.
3. On the **Cloud Service Check** page, view the details of check results.

Threat Detection / Cloud Service Check Settings

Cloud Service Check

At-Risk Items: 0 Risks: 0 Check item not enabled: 0 Checked items enabled: 9 Last Checked At: -- Check Now

All Risks All Types Enter a check item name to:

| <input type="checkbox"/> Checked Item | Severity/Affected Assets | Type | Last Checked | Actions |
|--|--------------------------|------------------------|--------------|----------------------------------|
| <input type="checkbox"/> RDS - Whitelist Configuration | Unchecked | Network access control | -- | Verify Whitelist |
| <input type="checkbox"/> OSS - Bucket Access Permissions | Unchecked | Data Security | -- | Verify Whitelist |
| <input type="checkbox"/> MongoDB - Whitelist Configuration | Unchecked | Network access control | -- | Verify Whitelist |
| <input type="checkbox"/> Redis - Whitelist Configuration | Unchecked | Network access control | -- | Verify Whitelist |
| <input type="checkbox"/> RDS - Database Security Policy | Unchecked | Data Security | -- | Verify Whitelist |
| <input type="checkbox"/> OSS - Logging Configuration | Unchecked | Data Security | -- | Verify Whitelist |

Verify Items per Page: 20 < Previous 1 Next >

- o **View the statistics of the last check**

You can view the total number of at-risk items and the numbers of risks at different levels in the **At-Risk Items** section, and the number of assets on which risks are detected in the **Risks** section. You can also view the number of disabled check items in the **Check item not enabled** section, the number of enabled check items in the **Checked items enabled** section, and the time when the check was last performed in the **Last Checked At** section.

- o **View check items**

You can view the information about the check items in the check item list. The information includes the risk severities of check items, the number of affected assets, the types of affected assets, the types of check items, and the time when the check was last performed.

- o **View the details of check results**

You can click the name of a check item in the **Checked Item** column to go to the panel that displays the details of the check item. In the panel, you can view the description of the check item, potential impacts caused by the detected risks, and suggestions on how to handle the risks.

Handle the detected configuration risks of your cloud services

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Threat Detection > Cloud Service Check**.
3. On the **Cloud Service Check** page, handle the configuration risks detected on your cloud services.
 - o **Verify the configurations after modification**

If you have modified the configurations for which risks are detected, find the check item in the check item list and click **Verify** in the Actions column to check whether the new configurations are at risk.

The screenshot displays the 'Cloud Service Check' page. At the top, there are summary statistics: At-Risk Items (0), Risks (0), Check item not enabled (0), and Checked items enabled (9). A 'Check Now' button is located to the right of these statistics. Below the summary is a table of check items. The table has columns for 'Checked Item', 'Severity/Affected Assets', 'Type', 'Last Checked', and 'Actions'. All items in the table are marked as 'Unchecked'. The 'Actions' column for each item contains a 'Verify' link and a 'Whitelist' link. At the bottom of the table, there is a 'Verify' button and a pagination control showing 'Items per Page' set to 20, with 'Previous' and 'Next' navigation options.

o Add check items to a whitelist

If you trust a check item for which risks are detected, find the check item in the check item list and click **Whitelist** in the Actions column to add the check item to a whitelist. Then, the state of the check item is displayed as **Ignored** in the Severity/Affected Assets column. **Ignored** check items are not counted in the total number of at-risk items in the **At-Risk Items** section.

In the check item list, you can click **Remove** to remove the ignored check items from the whitelist.

Note After you add a check item to the whitelist, the risk that is detected for the check item is ignored only for this time. If the risk is detected again, Apsara Stack Security still displays the check result of this check item.

21.4.6. Application whitelist

The application whitelist feature prevents unauthorized programs from running on your servers and provides a trusted running environment for your servers.

Context

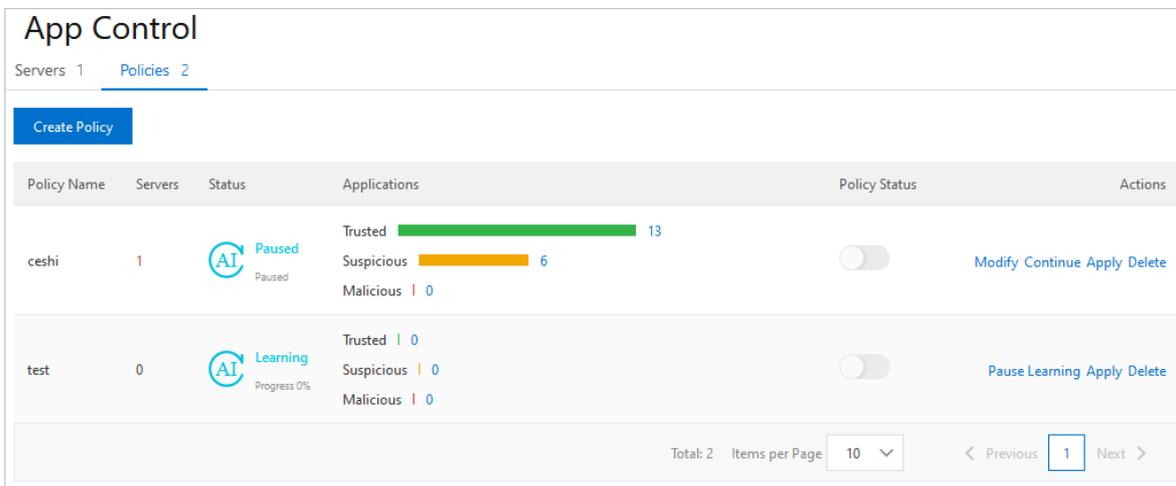
The application whitelist feature allows you to apply application whitelist policies to your servers that require special protection. Apsara Stack Security identifies trusted, suspicious, and malicious programs based on the policies. Then, you can add the identified programs to an application whitelist based on your business requirements. This prevents unauthorized programs from running. This feature protects your servers from untrusted and malicious programs and improves resource usage.

After you create an application whitelist policy, you can apply it to a server that requires special protection. Then, Apsara Stack Security scans for suspicious or malicious programs on the server and generates alerts for the programs that are not in the application whitelist.

Note If a program that is not in the application whitelist starts, an alert is generated. The program may be a normal program that is newly started or a malicious program that is inserted into your compromised server. If the program is a normal program, a frequently used program, or a third-party program installed by you, we recommend that you add the program to the application whitelist. After you add the program to the application whitelist, Apsara Stack Security no longer generates alerts for this program the next time the program starts. If the program is malicious, we recommend that you immediately delete this program and check whether the configuration files such as cron tasks are tampered with.

Step 1: Create an application whitelist policy

1. Log on to Apsara Stack Security Center.
2. In the left-side navigation pane, choose **Threat Detection > Application Whitelists**.
3. On the **Application Whitelist** page, click the **Policies** tab.
4. On the App Control page, click the **Policies** tab. Then, click **Create Policy**.
5. In the Create Policy step of the **Create Whitelist Policy** panel, configure the following parameters:
 - o **Policy Name**: the name of the application whitelist policy.
 - o **Intelligent Learning Duration**: the duration for Apsara Stack Security to perform intelligent learning. Valid values: 1 Day, 3 Days, 7 Days, and 15 Days. The intelligent learning feature uses machine learning to automatically collect and categorize large amounts of alert data. Apsara Stack Security can identify suspicious or malicious processes based on the collected data.
 - o **Servers for Intelligent Learning**: the servers to which you want to apply the application whitelist policy.
6. Click **Next** to create the application whitelist policy.
 After the application whitelist policy is created, the policy details are displayed in the policy list on the Policies tab.



The following table describes the parameters in the list of application whitelist policies.

| Parameter | Description |
|--------------------|---|
| Policy Name | The name of the application whitelist policy. |
| Servers | The number of servers to which the application whitelist policy is applied. |

| Parameter | Description |
|---------------------|---|
| Status | <p>The status of the policy. Valid values:</p> <ul style="list-style-type: none"> ◦ Applied: Intelligent learning is complete. The policy has been applied to servers. ◦ Pending Confirmation: Intelligent learning is complete. The policy must be confirmed and enabled. <p>After intelligent learning is complete, you must turn on the switch in the Policy Status column to enable this policy. The policy takes effect only after it is enabled. Apsara Stack Security automatically identifies the programs on your servers as trusted, suspicious, or malicious programs.</p> <ul style="list-style-type: none"> ◦ Paused: Intelligent learning is manually paused. You can click Continue in the Actions column to resume intelligent learning. ◦ Learning: Intelligent learning is in progress. <p>After an application a whitelist policy is created, Apsara Stack Security automatically performs intelligent learning based on the policy. The status of a new application whitelist policy is Learning.</p> |
| Applications | <p>The number of programs of each type on all servers to which the policy is applied. The program types include trusted, suspicious, and malicious.</p> |
| Actions | <p>The operations that you can perform on a policy. You can perform the following operations:</p> <ul style="list-style-type: none"> ◦ Apply: Add or remove servers to which the policy is applied in the Apply Whitelist Policy panel. ◦ Modify: Modify the policy in the Modify Whitelist Policy panel. You can change the values of Policy Name and Intelligent Learning Duration, and add or remove the servers on which intelligent learning is automatically performed. ◦ Pause Learning: Pause intelligent learning. ◦ Continue: Resume intelligent learning. <p>After you click Continue, the status of the policy changes to Learning. You can view the learning progress of the policy in the Status column.</p> <ul style="list-style-type: none"> ◦ Delete: Delete the policy. <p>After the policy is deleted, the servers to which the policy is applied are no longer protected by the policy.</p> |

Step 2: Apply the created application whitelist policy to servers

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Threat Detection > Application Whitelists**.
3. On the **Servers** tab of the App Control page, click **Add Server**.
4. In the **Add Server** panel, configure the following parameters:
 - **Whitelist Policy**: Select the created application whitelist policy from the drop-down list.
 - **Event Handling**: The default value is **Alert**, which indicates that Apsara Stack Security generates an alert when a suspicious program is detected.

If a program that is not in the application whitelist starts, Apsara Stack Security automatically generates an alert. You can click the number in the **Suspicious Events** column to go to the **Alerts** tab of the server details page and view the alert details.

 - **Servers**: Select the server to which you want to apply the application whitelist policy. You can select multiple servers.

To search for a server, enter the server name in the **Servers** search box and click the search icon. **Fuzzy match is supported.**

5. Click **OK**. The application whitelist policy is applied to the selected servers. After the application whitelist is created, you can view the protected servers and the name of the application whitelist policy in the server list on the **Servers** tab.

The **Servers** tab displays the following information of a protected server:

- **Server Name/IP**: the name and IP address of the server to which the application whitelist policy is applied.
- **Whitelist Policy**: the name of the application whitelist policy that is applied to the server.
- **Suspicious Events**: the number of programs that are not in the application whitelist and have started. If a suspicious program starts on the server, Apsara Stack Security detects the program and generates an alert.
- **Event Handling**: The default value is **Alert**, which indicates that Apsara Stack Security generates an alert when a suspicious program is detected.

If a program that is not in the application whitelist starts, Apsara Stack Security automatically generates an alert. You can click the number in the **Suspicious Events** column to go to the **Alerts** tab of the server details page and view the alert details.

- **Actions**: After you click **Delete** in the **Actions** column, the application whitelist policy no longer applies to the server.

After you click **Delete** in the **Actions** column, the application whitelist policy becomes invalid for the server. In this case, if a program that is added to the application whitelist starts on this server, Apsara Stack Security generates an alert.

Add a program to or remove a program from an application whitelist

After you configure an application whitelist policy for your server, you can view the detailed information in the server list on the **Servers** tab. The information includes the details of the protected server and the name of the application whitelist policy that is applied to the server. You can click a policy name in the **Whitelist Policy** column to view the programs running on the server. You can also view the number of trusted, suspicious, and malicious programs and their detailed information.

The following information about each program on the server is displayed:

- **Type**: the type of the program. Programs are classified into trusted, suspicious, and malicious programs.
- **Process Name**: the name of the program.
- **Hash**: the hash function of the program. A hash function is used to identify whether a program is unique. This helps protect servers against malicious programs.
- **Path**: the file path of the program on the server.
- **Degree of Trustability**: the degree of trustability for the program. The value of this parameter is determined by Apsara Stack Security. Valid values: 0%, 60%, and 100%. The value 0% indicates malicious programs, 60% indicates suspicious programs, and 100% indicates trusted programs.

 **Note** We recommend that you handle the program whose Degree of Trustability is 0% at the earliest opportunity.

- **Actions**: the operations that can be performed on the program. You can determine whether to add the program to the whitelist based on the services deployed on your server. You can perform the following operations:
 - **Add to Whitelist**: If you trust the program, add it to the whitelist.
 - **Remove from Whitelist**: After you remove the program from the whitelist, Apsara Stack Security identifies the program as untrusted. If this program starts, Apsara Stack Security generates an alert.

21.4.7. Assets

21.4.7.1. View the security status of a server

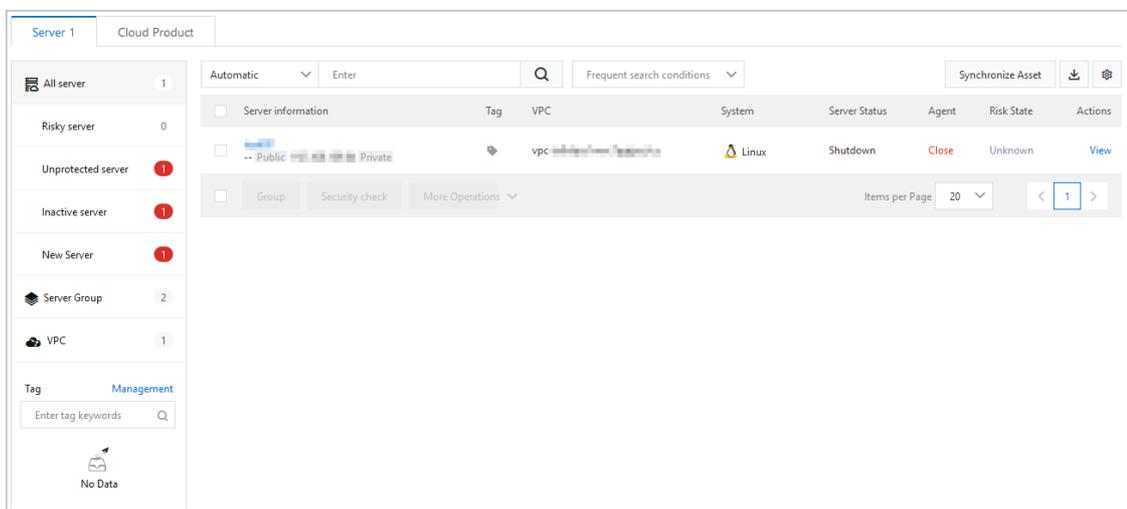
The Assets page displays security information about each protected server. The information includes the virtual private cloud (VPC) where each server resides, server status, and risk status. This topic describes how to search for specific servers and view the security status of these servers. This topic also describes how to specify search conditions and select the items that you want to display on the Assets page.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Assets**.
3. On the **Server(s)** tab of the Assets page, view the security status of each server.

You can perform the following operations:

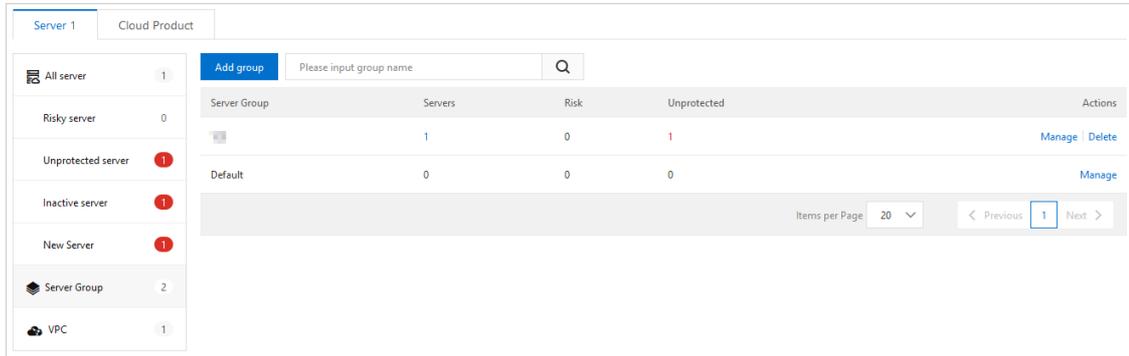
- **Filter servers by status**
 - In **All server**, you can view the numbers of all servers, risky servers, unprotected servers, new servers, and servers that are shut down.



To view the security information about a server, you can click the name of the server or click **Fix** in the **Actions** column. For more information, see [View the details of a single asset](#).

- You can click **Risky server**, **Unprotected server**, **Shutdown Server(s)**, or **New Server(s)** to view security information about specific servers.
- **Filter servers by group**

- You can click **Server Group** to view the numbers of all servers, servers that are at risk, and unprotected servers in each server group. You can also view the total number of server groups.

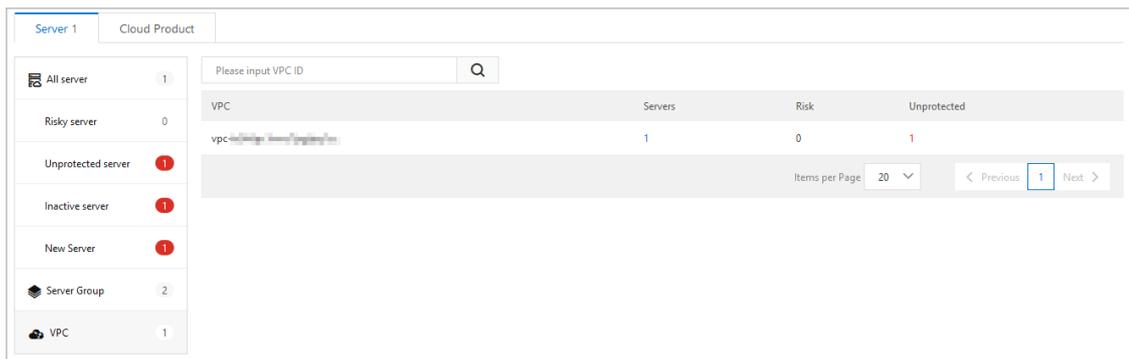


To manage server groups, you can click **Manage** or **Delete** in the Actions column. For more information, see [Manage asset groups](#).

- You can select a server group and click the number in the **Servers**, **Risk**, or **Unprotected** column to view the security information about specific servers in this group.

○ **Filter servers by VPC ID**

- You can click **VPC** to view the numbers of all servers, servers that are at risk, and unprotected servers in each VPC. You can also view the total number of VPCs.



- You can find a VPC and click the number in the **Servers**, **Risk**, or **Unprotected** column to view the security information about specific servers in this VPC.

○ **Filter servers by tag**

In the navigation tree, you can click a tag to view the security information about servers to which the tag is added.

○ **Filter servers by condition**

If you click **All Servers**, **Server Group**, **VPC**, or a tag in the navigation tree, you can specify filter conditions above the right-side list to search for specific servers.

- Use one filter condition to search for specific servers:

You can select a filter condition and select or enter keywords to search for specific servers. The filter conditions include **Internet IP**, **Private IP**, **Instance name**, **System**, **Baseline problems**, **Vul problems**, **Alert problems**, **Risk Status**, **Online or Offline**, **Tag**, **Group name**, **OS**, and **Is there a snapshot risk**.

 **Note**

You can specify multiple filter conditions at a time and specify a Boolean operator for the conditions. The following list describes the Boolean operators:

- Boolean operators:
 - **AND**: specifies the **AND** logical relation for the conditions.
 - **OR**: specifies the **OR** logical relation for the conditions.
- If you want to search for servers that meet at least one of the filter conditions, you must set the Boolean operator to **OR**.
- If a filter condition requires you to enter a keyword, you must enter the keyword and click the **Search** icon. Results are displayed only after you click the **Search** icon.

- Use multiple filter conditions to search for specific servers:

If you select multiple filter conditions, they are all applied to search for specific servers.

You can also click **Server Group**, **VPC**, or a tag, and use the search box above the asset list to search for specific servers.

- **Save frequently used filter conditions**

You can save the filter conditions that are applied as frequently used search conditions. To save the conditions, click **Save** above the right-side list, and enter a name in the **Save condition** dialog box. Then, you can select the saved conditions from the **Frequent search conditions** drop-down list on the right of the **Search** icon.

- **Customize displayed items**

On the **Assets** page, you can click the  icon in the upper-right corner. Then, you can select the items that you want to display on the **Assets** page.

21.4.7.2. View the security status of cloud services

The **Assets** page displays the security information about each protected cloud service. The information includes the at-risk services and the types of services such as SLB and NAT Gateway. This topic describes how to configure search conditions to view the security status of cloud services.

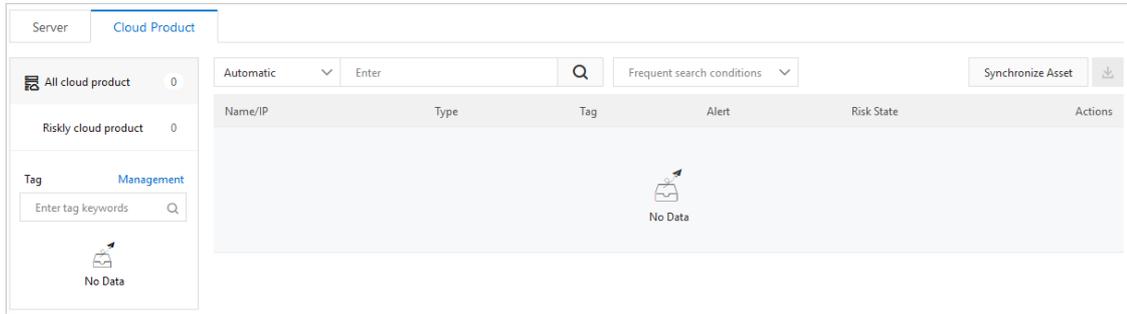
Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Assets**.
3. On the **Cloud Product** tab of the **Assets** page, view the security status of cloud services.

You can perform the following operations based on your business requirements:

- **Search by asset status**

- In the left-side pane of the **Cloud Product** tab, you can view the numbers of **All cloud product** and **Risky cloud product**. You can also view the security status of all cloud services.



- Click **Risky cloud product** to view the cloud services that are at risk.

You can click the name of the required cloud service or click **View** in the **Actions** column that corresponds to a service to view detailed information. For more information, see [View the details of a single asset](#).

○ **Search by asset type**

Cloud services are classified into two asset types:

- **SLB**
- **NAT Gateway**

In the left-side pane of the **Cloud Product** tab, you can view the number of cloud services of each type. You can click **SLB** or **NAT** to view the security status of the required cloud service.

○ **Search by tag**

In the **Tag** section in the left-side pane of the **Cloud Product** tab, you can view the number of cloud services to which each tag is added. You can click a tag to view the security status of cloud services to which the tag is added.

○ **Filter by search condition**

You can click **All cloud product**, **SLB**, or **NAT** in the left-side pane of the **Cloud Product** tab and configure search conditions in the search box to search for specific assets.

For example, you can click **All cloud product** and configure search conditions to search for specific assets.

- Use multiple subconditions to search for specific assets:

Select a condition from the drop-down list of the search box on the **Cloud Product** tab, and select a subcondition or enter a keyword into the search box to search for specific assets. Supported search conditions are **Internet IP**, **Instance name**, **Alert problems**, **Risk Status**, **Tag**, and **Group name**.

Note

You can configure multiple search conditions at a time and specify a Boolean operator for search conditions. The following list describes the Boolean operators:

- Boolean operators:
 - **AND**: specifies the **AND** logical relation for the search conditions.
 - **OR**: specifies the **OR** logical relation for the search conditions.
- If you need to use one search condition and multiple keywords to search for specific servers, set the Boolean operator to **OR**.
- If the search condition requires you to enter a keyword, enter a keyword and click the **Search** icon. Results appear after you click the **Search** icon.

- Use multiple filter conditions to search for specific servers:
Apply multiple search conditions.
 - You can click **SLB**, **NAT**, or a tag specified in the **Tag** section and configure conditions in the search box on the **Cloud Product** tab to search for specific assets.
 - You can also click **All cloud product**, **SLB**, or **NAT** and select a tag specified in the **Tag** section to search for specific assets.
- **Set frequently used search conditions**
You can save the filter conditions that are applied as frequently used search conditions. Click **Save** below the search box and enter a name in the **Save condition** dialog box. Then, you can select the saved search condition from the **Frequent search conditions** drop-down list on the right of the search box.

21.4.7.3. View the details of a single asset

The **Assets** page provides details about all assets. These details include basic information, alert management status, baseline check analysis, and asset fingerprints. This topic describes how to view the details of a server or a cloud service.

Context

The basic information about assets is displayed on the **Assets** page. Different types of assets, such as servers and cloud services, are managed in different ways.

The following table lists the features that are supported by servers and cloud services on the **Assets** page. The following list describes the marks that are used to indicate whether a feature is supported by servers or cloud services:

- Cross (×): not supported.
- Tick (√): supported.

| Feature | Description | Server | Cloud service |
|-------------------|---|--------|--------------------------------------|
| Basic Information | Risk state: displays the number of risks detected on an asset. The following types of risks can be detected: <ul style="list-style-type: none"> ● Vulnerability ● Alert ● Baseline risk | √ | √ (Only alerts can be processed.) |
| | Detail: displays the configuration and protection status of an asset. You can specify a group and a tag for the asset. | √ | √ (Asset grouping is not supported.) |
| | Asset investigation: displays asset fingerprints, including ports, software, processes, and accounts. | √ | X |
| | Vulnerability check: displays the types of vulnerabilities that can be detected. You can specify the types of vulnerabilities that you want to detect for an asset. | √ | X |
| | Login security setting: displays the configured logon locations, IP addresses, time periods, and accounts. You can manage relevant alerts for an asset. | √ | X |
| Vulnerabilities | Displays the results of vulnerability detection on an asset. | √ | X |

| Feature | Description | Server | Cloud service |
|--------------------|--|--------|---------------|
| Alerts | Displays the alerts that are generated for an asset. | √ | √ |
| Baseline Risks | Displays the results of the baseline check on an asset. | √ | X |
| Asset Fingerprints | Displays the details of asset fingerprints for an asset. | √ | X |

Procedure

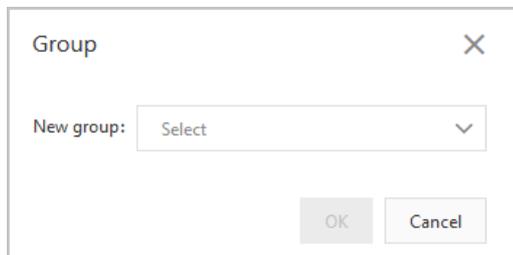
1. Log on to Apsara Stack Security Center.
2. In the left-side navigation pane, choose **Threat Detection > Assets**.
3. On the **Assets** page, click the **Server(s)** or **Cloud Product** tab.
4. On the **Server(s)** or **Cloud Product** tab, find the required asset and click its name.
5. View the details of the asset.

On the asset details page, click the **Basic Information**, **Vulnerabilities**, **Alerts**, **Baseline Risks**, or **Asset Fingerprints** tab to view relevant details.

The following list describes the details of the asset:

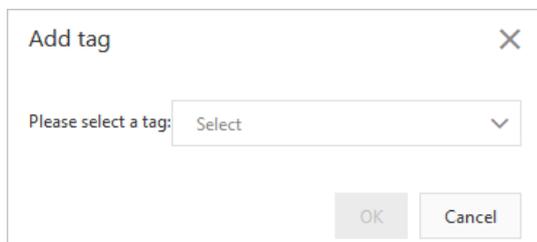
- **Basic Information:** This tab consists of sections in which you can view asset details and manage an asset.
 - **Risk State:** This section displays the numbers of vulnerabilities, alerts, and baseline risks on the asset. You can click the number under Vulnerabilities, Alerts, or Baseline Risks to view the details.
 - **Detail:** This section displays information about the asset configuration and security protection settings, and allows you to manage asset tags and groups.
 - **Change asset groups**

Click **Group**. In the **Group** dialog box, select a new group and click **OK**.



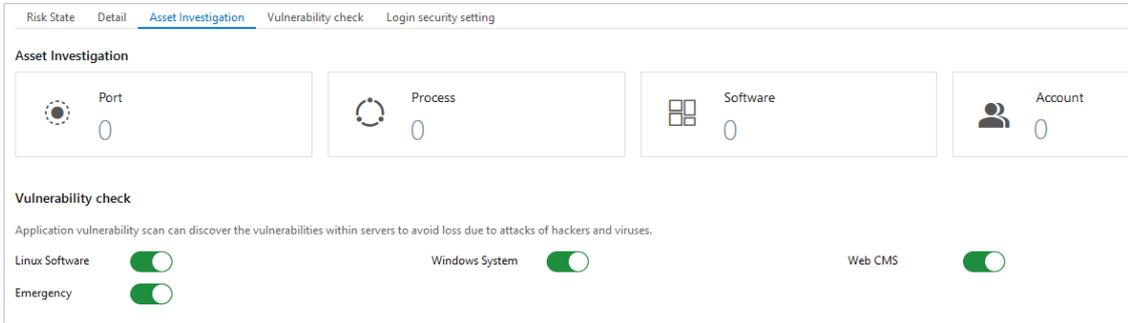
- **Modify tags**

Click the  icon. In the **Add tag** dialog box, select a tag and click **OK**.

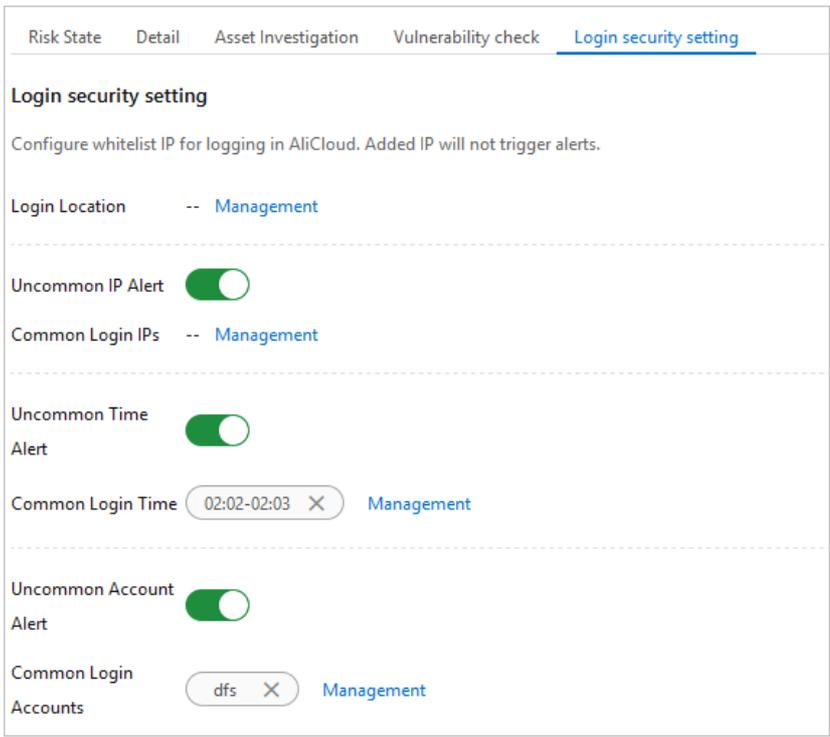


You can click the  icon in the right of a tag to delete the tag.

- **Asset Investigation:** This section displays the fingerprints of an asset. You can click the number under an item to go to the **Asset Fingerprints** tab to view the details.



- **Vulnerability check:** This section displays vulnerability check items that are enabled or disabled for an asset. You can enable or disable different types of vulnerability checks for the asset. The vulnerabilities include Linux software vulnerabilities, Windows system vulnerabilities, Web CMS vulnerabilities, and urgent vulnerabilities.
- **Login security setting:** This section allows you to specify approved logon locations, configure advanced logon settings, and turn on or turn off alerting for unapproved IP addresses, time, and accounts. The advanced logon settings include approved IP addresses, time periods, and accounts. You can also specify approved IP addresses, time periods, and accounts for a specific asset.



- **Vulnerabilities:** This tab displays vulnerabilities detected on an asset.

| Priority | Disclosure Time | Vulnerability | Related process | Vul (cve) | Status | Actions |
|----------|-----------------|--|-----------------|------------------------|---------|------------------------|
| High | Aug 10, 2020 | RHSA-2018:1062-Important: kernel security, bug fix, and enhancement update | | CVE-2016-3672 Total 30 | Unfixed | Fix Verify Details |
| High | Aug 10, 2020 | RHSA-2018:1453-Critical: dhcp security update | | CVE-2018-1111 | Unfixed | Fix Verify Details |
| High | Aug 10, 2020 | RHSA-2018:3665-Important: NetworkManager security update | | CVE-2018-15688 | Unfixed | Fix Verify Details |
| High | Aug 10, 2020 | RHSA-2017:3263-Moderate: curl security update | | CVE-2017-1000257 | Unfixed | Fix Verify Details |

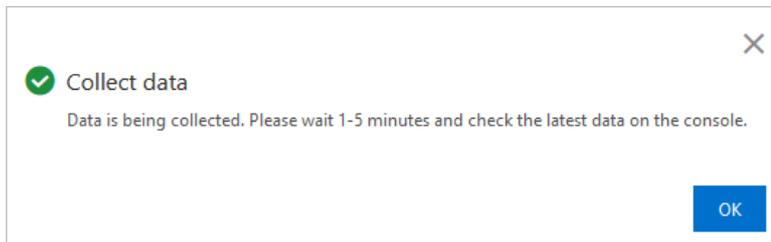
- **Alerts:** This tab displays alerts generated for an asset.
- **Baseline Risks:** This tab displays baseline risks of an asset.

| Severity | Baseline | Checked Item | Failed Items/Affected Servers | Category | Last Check |
|----------|---|--------------|-------------------------------|-------------------------|------------------------|
| High | Alibaba Cloud Standard - CentOS Linux 7/8 Security Baseline Check | 15 | 5 / 1 | Best security practices | Aug 13, 2020, 00:35:11 |
| High | Weak password - Linux system login weak password baseline | 1 | Risk free | Weak password | Aug 13, 2020, 00:35:11 |

- **Asset Fingerprints:** This tab displays the fingerprints, including ports, processes, software, and accounts of an asset.

You can manually collect the latest fingerprints of an asset.

- a. You can click the **Port**, **Software**, **Process**, **Account**, or **Scheduled Tasks** tab. In the upper-right corner, click **Collect data now**.
- b. In the **Collect data** message, click **OK**.



After the data collection task is submitted, it takes one to five minutes to collect the fingerprints of the required asset. After the data collection task is complete, you can view the latest fingerprints of the asset.

21.4.7.4. Enable and disable server protection

This topic describes how to enable and disable server protection.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Assets**.
3. On the **Server(s)** tab of the page that appears, enable or disable server protection for specified servers.

- **Enable server protection**

Select one or more servers where the agent is in the **Close** state, and choose **More operations > Turn on protection**.

After server protection is enabled, the status in the **Agent** column changes to **Enable**.

- **Disable server protection**

If you confirm that a server does not require protection from Apsara Stack Security, you can disable protection for the server. Select one or more servers where the agent is in the **Enable** state, and choose **More operations > Suspend Protection**.

Note After server protection is disabled, Apsara Stack Security stops protecting your servers. For example, Apsara Stack Security no longer detects vulnerabilities or generates alerts for detected risks. We recommend that you proceed with caution.

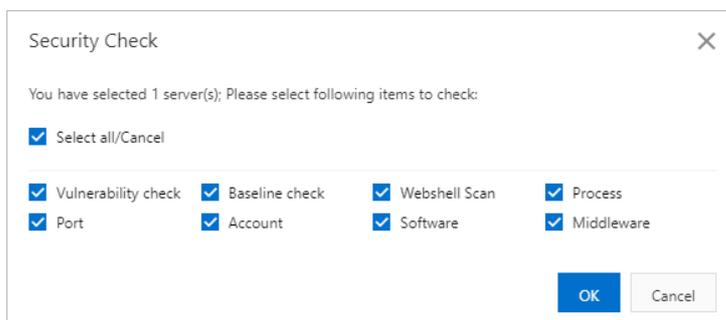
After server protection is disabled, the status of the agent on your servers changes to **Close**.

21.4.7.5. Perform a quick security check

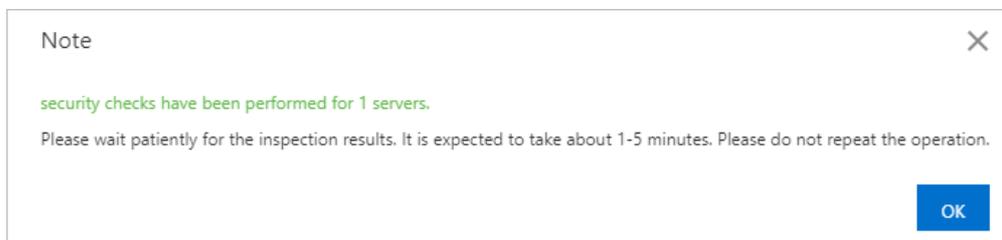
The Server tab of the Assets page allows you to run security checks. You can dispatch security check tasks to scan for vulnerabilities, baseline risks, or webshells, and collect asset fingerprints on a specific server. The asset fingerprints are ports, software, processes, and accounts. This topic describes how to perform a security check on servers.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Assets**.
3. On the **Server** tab, select one or more servers on which you want to perform a security check.
4. In the lower part of the page, click **Security check**.
5. In the **Security Check** dialog box, select check items.



6. Click **OK** to start the check.
7. In the message that appears, click **OK**.



After the security check is complete, the check results are automatically displayed on the details pages of the selected servers.

21.4.7.6. Manage server groups

This topic describes how to create, modify, delete, and replace server groups.

Create a server group

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Assets**.
3. On the **Server(s)** tab of the page that appears, click **Server Group** in the navigation tree.

Note By default, the assets that are not grouped are in the **Default** group.

4. Click **Add group**.
5. In the **Add group** dialog box, configure parameters for the new group.

To configure the parameters, perform the following steps:

- i. Enter a name for the new group in the **Group name** field.
- ii. Add servers to the new group.

You can add servers in the **Default** group to the new group. You can also move servers from another group to the new group. To add or move servers, select **Default** or other groups in the **Asset Group** section, and select or clear the required check boxes in the asset list in the right area of the section.

6. Click **OK**.
In the server group list, you can view the new group.

Modify or delete a server group

The following procedure describes how to modify or delete a server group. When you modify a server group, you can rename the group or adjust the servers in the group.

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Threat Detection > Assets**.
3. On the **Server(s)** tab of the page that appears, click **Server Group** in the navigation tree.
4. Find the server group that you want to modify or delete. In the Actions column, click **Manage** or **Delete**.

You can perform the following operations based on your business requirements:

- **Modify the group**
 - a. In the Actions column, click **Manage**. The Group dialog box appears.
 - b. In the **Group** dialog box, select the group in the **Asset Group** section.
 - c. In the right area of the section, clear the check boxes that correspond to the required servers in the asset list.
 - d. Click **OK**. The server group is modified.

- **Delete the group**

In the Actions column, click **Delete**. In the message that appears, click **OK**.

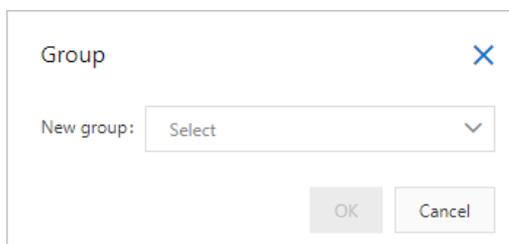
 **Note** After you delete a group, servers in this group are moved to the **Default** group.

Replace a server group

You can add servers to a server group to manage multiple servers at a time. We recommend that you add the same types of servers to a server group. For example, if you configure a baseline check template, you can specify a server group and apply the template to all servers in the group. You can also filter and view servers based on server groups.

To add servers to a specific server group, perform the following steps:

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Threat Detection > Assets**.
3. On the **Server(s)** tab of the page that appears, select one or more servers and click **Group** in the lower part of the page.
4. In the **Group** dialog box, select a new server group.



5. Click **OK**.

21.4.7.7. Manage asset tags

This topic describes how to add asset importance tags to your assets and how to create, modify, and delete custom tags.

Context

Apsara Stack Security provides the asset importance tags described in the following table to classify assets. You

can select appropriate importance tags for your assets.

An asset importance tag is transformed to an **asset importance score**. An **asset importance score** is used to calculate a vulnerability priority score. You can determine whether to preferentially fix a vulnerability based on the vulnerability priority score. We recommend that you add importance asset tags to core assets. Apsara Stack Security prompts you to fix vulnerabilities based on the importance of each asset. The following table describes the relationships between asset importance tags and asset importance scores.

| Asset importance tag | Asset importance score | Recommendation |
|----------------------|------------------------|--|
| Important Assets | 1.5 | Assets that are related to crucial business or store core business data. Virus intrusion into the assets adversely affects the system and causes major loss. |
| General Assets | 1 | Assets that are related to non-crucial business and are highly replaceable. Virus intrusion into the assets causes less impact on the system. |
| Test Assets | 0.5 | Assets for functional or performance tests, or assets that can cause less impact on the system. |

 **Note** If you do not add asset importance tags, the **General Assets** tag is automatically added to each asset. This tag indicates that the asset importance score is 1.

Create a custom tag

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Assets**.
3. On the page that appears, click the **Server(s)** or **Cloud Product** tab.
4. In the navigation tree of the **Server(s)** or **Cloud Product** tab, click **Management** on the right of **Tag**.
5. In the **Add tag** dialog box, enter the tag name in the **Tag** field.

6. In the **Asset Group** section, select a server group. Then, select the required servers to add the new tag to the selected servers in the right area of the section.
7. Click **OK**.

In the asset list of the **Server(s)** or **Cloud Product** tab, you can click the  icon in the **Tag** column to add the new tag to an asset.

 **Note** You can add multiple tags to one asset. All tags of an asset are displayed in the **Tag** column.

Modify or delete a custom tag

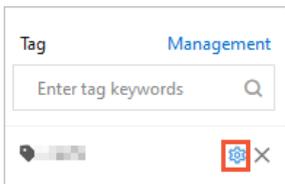
The following procedure describes how to modify or delete a custom tag. When you modify a tag, you can rename the tag or adjust the servers to which the tag is added.

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Assets**.
3. On the page that appears, click the **Server(s)** or **Cloud Product** tab.
4. On the **Server(s)** or **Cloud Product** tab, modify or delete a tag.

Perform the following operations to modify or delete a tag:

- **Modify a tag**

- a. Find the tag that you want to modify and move the pointer over the  icon on the right of the tag.

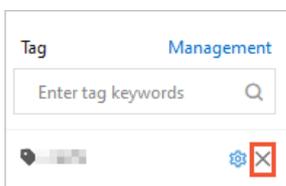


- b. In the **Tag** dialog box, enter a new name in the **Tag** field, add the tag to more servers, or remove the tag from specific servers.

- c. Click **OK**.

o **Delete a tag**

Find the tag that you want to delete and click the  icon in the **Tag** column. In the message that appears, click **OK**.



21.4.8. Vulnerability scan

21.4.8.1. Quick start

This topic describes how to get started with the vulnerability scan feature.

The following procedure shows how to use the vulnerability scan feature:

1. Configure the following detection items and the required cycles based on your environment requirements:
 - o Overall Monitoring: Configure detection features and the monitoring cycle of each detection feature. For more information, see [Configure overall monitoring](#).
 - o Basic Monitoring: Configure Weak Password Vulnerability Monitoring, Operation Security Vulnerability Monitoring, CMS Application Vulnerability Monitoring, and Baseline Monitoring. For more information, see [Configure basic monitoring](#).
 - o Web Monitoring: Configure the monitoring cycle and the types of web vulnerabilities that you want to monitor. For more information, see [Configure web monitoring](#).
 - o Whitelist: Add the assets that do not require detection to the whitelist. For more information, see [Configure a whitelist](#).
2. Import assets that require vulnerability scans.
 - o Import internal assets: Configure a scan engine to import your internal assets in virtual private clouds (VPCs). For more information, see [Configure a scan engine for internal assets](#).
 - o Import Internet assets: Import your Internet assets. For more information, see [Import assets](#).

 **Note** The number of imported assets cannot exceed the specified upper limit.

3. View and confirm the results of vulnerability scans.
 - o View the overall information to obtain the results of vulnerability scans. For more information, see [View the information on the Overview page](#).
 - o View and confirm vulnerability risks. For more information, see [Manage security vulnerabilities](#).
 - o View and confirm host compliance risks. For more information, see [Manage host compliance risks](#).

- View and confirm external risks, such as code leak risks. For more information, see [Manage external risks](#).
4. Generate vulnerability scan reports.
Generate reports to audit the vulnerabilities and baseline risks on assets on a regular basis. For more information, see [Manage external risks](#).

21.4.8.2. View the information on the Overview page

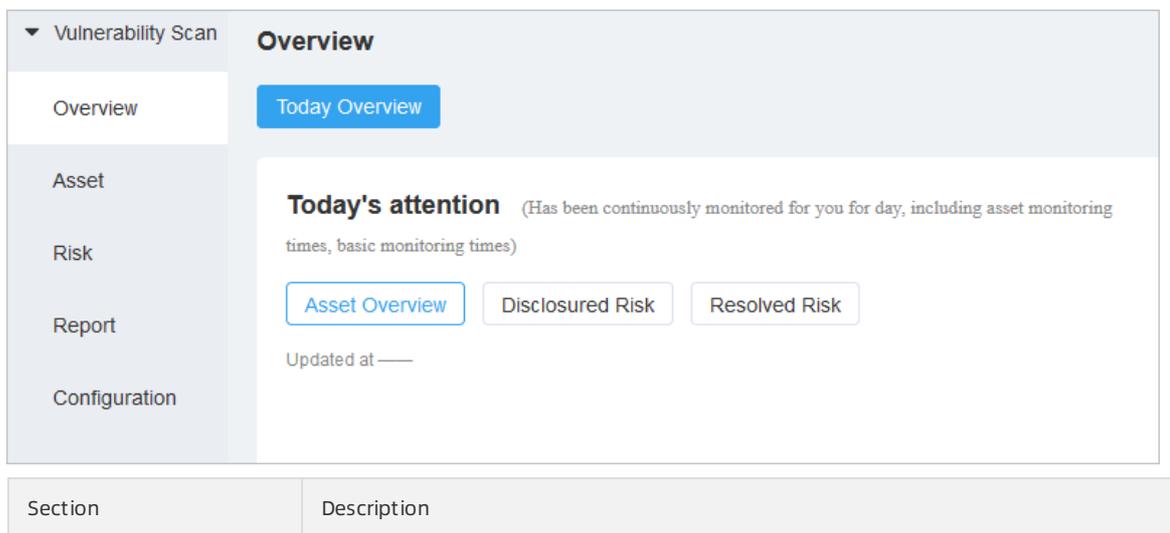
This topic describes the overall results of vulnerability scans. Security administrators can understand the vulnerability situation based on the overall results.

Context

The vulnerability scan feature can identify the following vulnerabilities: web security vulnerabilities, content management system (CMS) application vulnerabilities, weak password vulnerabilities, O&M security vulnerabilities, and baseline security vulnerabilities.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan > Overview**.
3. View the overall results of vulnerability scans.



| Section | Description |
|-------------------------|---|
| Today's attention | <p>View Asset Overview, Disclosed Risk, and Resolved Risk of the current day.</p> <ul style="list-style-type: none"> ◦ Asset Overview: displays the numbers of hosts, websites, and domain names for the current day and provides a security score for the current assets. The radar chart on the right shows the distribution of web security vulnerabilities, CMS application vulnerabilities, weak password vulnerabilities, O&M security vulnerabilities, and baseline security vulnerabilities. ◦ Disclosed Risk: displays the numbers of high-risk vulnerabilities, medium-risk vulnerabilities, and low-risk vulnerabilities, and the total number of these vulnerabilities for the current day. These vulnerabilities are not fixed. The Disclosed Risk Distribution section on the right displays the distribution of unfixed vulnerabilities. ◦ Resolved Risk: displays the numbers of high-risk vulnerabilities, medium-risk vulnerabilities, and low-risk vulnerabilities, and the total number of these vulnerabilities for the current day. These vulnerabilities are fixed. The Resolved Risk Distribution section on the right displays the distribution of fixed vulnerabilities. |
| Asset Risk Top 5 | <p>View the top five assets that are at risk on the Security Vulnerabilities and Host Compliance tabs.</p> <p>These assets are displayed by asset or group.</p> |
| Risk Monitoring Trend | <p>View the trend charts of vulnerabilities on the Security Vulnerabilities and Host Compliance tabs.</p> <p>Fixed and unfixed vulnerabilities are identified by lines in different colors. You can move the pointer over a line to view the numbers of unfixed vulnerabilities and fixed vulnerabilities for the specific day.</p> |
| Asset Monitoring Trend | <p>View the trends in the numbers of protected hosts and websites.</p> <p>Hosts and websites are identified by lines in different colors. You can move the pointer over a line to view the number of protected hosts and websites for the specific day.</p> |
| Risk Asset Ranking List | <p>View the rankings of assets that are at risk on the Latest Risk and High Risk tabs.</p> |
| Port Service Statistics | <p>View the statistics on the Port and Host Service tabs.</p> |

21.4.8.3. Asset management

21.4.8.3.1. View the results of asset analysis

This topic describes how to view the analysis results of websites and hosts.

Context

The asset analysis feature allows you to view the analysis results of websites and hosts. For the websites, you can view Web Service, Open Source Framework, and Device Type. For the hosts, you can view Host Port, Host Service, and Operation System.

Procedure

1. [Log on to Apsara Stack Security Center.](#)

2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan**.
3. In the **Vulnerability Scan** pane, click **Asset**.
4. On the **Asset** page, click the **Asset Analysis** tab to view websites.

Asset

Asset Analysis Asset List Asset Import Availability Monitoring Custom Update Detection

Web Asset Updated at Aug 05, 2021, 08:00:00

Web Service

No Date

Open Source Framework

No Date

Device Type

No Date

5. View hosts.

Host Asset Updated at Aug 05, 2021, 08:00:00

Host Port

No Date

Host Service

No Date

Operation System

No Date

21.4.8.3.2. Import assets

This topic describes how to import Internet assets.

Context

The vulnerability scan feature works only on imported assets. If you want to scan the vulnerabilities of your assets,

you must import your assets.

The assets that the feature supports include Internet assets and internal assets. The internal assets refer to the assets in a virtual private cloud (VPC).

- To import internal assets, you must add a scan engine. For more information, see [Configure a scan engine for internal assets](#).
- To import Internet assets, perform the operations provided in this topic.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan**.
3. In the **Vulnerability Scan** pane, click **Asset**.
4. On the **Asset** page, click the **Asset Import** tab to view imported assets.

| Asset import task | Asset Type | Import Progress | Start Time | End Time | Operation |
|-------------------|----------------------|-----------------|------------------------|------------------------|-----------|
| | Private Asset | Finish | Dec 22, 2020, 09:46:29 | Dec 22, 2020, 09:47:02 | |
| | Public network asset | Finish | Dec 22, 2020, 00:15:00 | Dec 22, 2020, 00:17:09 | |
| | Public network asset | Finish | Dec 21, 2020, 21:59:38 | Dec 21, 2020, 22:06:51 | |
| | Public network asset | Finish | Dec 21, 2020, 21:56:46 | Dec 21, 2020, 22:06:38 | |
| | Public network asset | Finish | Dec 21, 2020, 21:53:24 | Dec 21, 2020, 21:51:41 | |

5. Click **Asset Import**. On the **Public network asset Import** page, create an asset import task.
 - i. In the **Import Asset** section, select **Manual Import** and enter the required assets in the field. Then, read and select the disclaimer.
 - You can enter domain names, URLs, IP addresses, and CIDR blocks.
 - You can enter the information about multiple assets at a time. Press Enter after you enter the information about one asset.
 - You cannot enter the information about the assets in VPCs.
 - The number of imported assets must be less than the number of remaining assets supported by the platform.

Note For example, if the number of remaining assets supported by the platform is 100 and 90 assets are entered, all the assets can be scanned. If 110 assets are entered, only 100 assets can be scanned, and the 10 assets that remain cannot be scanned.

- ii. In the **Asset Info** section, group the imported assets and configure an owner and a tag for the assets.
 - **Asset Group**: Select a group from the drop-down list. You can click the icon to create, edit, or delete a group.
 - **Person in charge**: Select an owner from the drop-down list. You can click the icon to create, edit, or delete an owner.
 - **Asset Tag**: Click **Add Tag** to add a tag to the imported assets.

iii. In the **Import Set** section, select the operations that you want to perform after the assets are imported.

| Operation | | Description |
|-----------------|---|---|
| Asset Discovery | <i>Auto Import subdomains</i> | Automatically queries the subdomain assets of the imported domain names. |
| | <i>Auto import associated IP</i> | Automatically adds IP address assets that are mapped to the domain names. |
| | <i>Auto synchronize tags and groups</i> | Applies the group and tag of the imported assets to the assets that are discovered by the system. |
| Web Asset | <i>Open WEB Monitoring</i> | Enables the web monitoring feature on the imported website assets. If you want to select the web monitoring rules to use, click the  icon. In the dialog box that appears, select the required web monitoring rules. For more information about how to configure web monitoring rules, see Configure web monitoring . |

iv. In the **Whitelist** section, add the assets that do not need to be scanned.

You can enter IP addresses and URLs. If you add more than one asset, you must press Enter after you enter the information about an asset.

v. Click **Save**.

6. Manage the created asset import task.

After the asset import task is created, you can view the task in the task list. You can also perform the following operations on the new asset import task.

| Icon | Description |
|---|--|
|  | View the details, results, and process of the asset import task. |
|  | Delete the asset import task. |

21.4.8.3.3. Manage assets

This topic describes how to view and manage assets.

Context

You can view the information about assets in the asset list. If the purpose or owner of an asset changes, security administrators can move the asset to another group or change the owner.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan**.
3. In the Vulnerability Scan pane, click **Asset**.
4. On the **Asset** page, click the **Asset List** tab to view assets.
5. Click the **Web Asset** tab and manage websites.

- i. Specify filter conditions to search for specific websites.

The filter feature allows you to search for the required websites in a more efficient manner.

| Filter condition | Description |
|------------------------------|---|
| Asset Type | The asset type, such as Internet assets or a specific VPC. |
| Asset Source | The source from which the asset is imported. Valid values: Manual and System Find . |
| Asset Group | The group to which the asset belongs. |
| Web Status | The status of the website. |
| Person in charge | The owner of the website. |
| Asset Change | The change status of the asset. Valid values: All , New , Update , No Update , and Offline . |
| Web Monitoring | The monitoring status of the website. |
| Risk Level | The risk level of the asset. |
| Web Service | The service type and version of the website. |
| WAF Recognition | Specifies whether the asset is identified by WAF. |
| Open Source Framework | The open source framework type of the asset. |
| Device Type | The type of the device. |
| Time range | The time range during which assets are imported. |
| Key information | The crucial information of the asset. The crucial information includes the website, domain name, IP address, title, and tag. |

- ii. Click **Export** to export the asset list to an Excel file.
- iii. Select one or more websites that you want to manage. The following table describes the operations that you can perform on websites.

| Action | Description |
|--------------------------------|---|
| Batch Web Monitoring | <p>Allows you to enable or disable web monitoring for multiple websites.</p> <ul style="list-style-type: none"> ▪ Batch Open Monitoring: To enable web monitoring for multiple websites, select Batch Open Monitoring from the drop-down list of Batch Web Monitoring. ▪ Batch Stop Monitoring: To disable web monitoring for multiple websites, select Batch Stop Monitoring from the drop-down list of Batch Web Monitoring |
| Change Group | Allows you to change the asset group of multiple assets at a time. |
| Change Person in charge | Allows you to change the owner of multiple assets at a time. |
| Batch Delete | Allows you to delete multiple assets at a time. After the assets are deleted, the assets are not scanned by the vulnerability scan feature. |

- 6. Click the **Host Asset** tab and manage the hosts.

- i. Specify filter conditions to search for specific hosts.

The filter feature allows you to search for the required hosts in a more efficient manner.

| Filter condition | Description |
|-------------------------|---|
| Asset Type | The asset type, such as Internet assets or a specific VPC. |
| Asset Source | The source from which the asset is imported. Valid values: Manual and System Find . |
| Asset Group | The group to which the asset belongs. |
| Person in charge | The owner of the asset. |
| Asset Change | The change status of the asset. Valid values: All , New , Update , No Update , and Offline . |
| Risk Level | The risk level of the asset. Valid values: All , High , Middle , Low , and Security . |
| SurviveStatus | The status of the asset. Valid values: Alive and Close . |
| Operation System | The operating system of the host. |
| Host Port | The port of the host. |
| CDN Recognition | Specifies whether Content Delivery Network (CDN) is configured for the asset. |
| Host Service | The service of the host. |
| Time range | The time range during which assets are imported. |
| Key information | The crucial information of the asset. The crucial information includes the IP address, host, tag, and domain name. |

- ii. Click **Export** to export the asset list to an Excel file.
- iii. Select one or more hosts that you want to manage. The following table describes the operations that you can perform on hosts.

| Action | Description |
|--------------------------------|---|
| Change Group | Allows you to change the asset group of multiple assets at a time. |
| Change Person in charge | Allows you to change the owner of multiple assets at a time. |
| Batch Delete | Allows you to delete multiple assets at a time. After the assets are deleted, the assets are not scanned by the vulnerability scan feature. |

21.4.8.3.4. Manage asset availability

This topic describes how to manage the availability of assets.

Context

If hosts or websites are used for a long period of time, they may become unavailable due to errors. Availability monitoring allows a security administrator to discover unavailable assets. Then, the security administrator can troubleshoot the issues that cause the assets to become unavailable.

Availability monitoring supports the following methods:

- HTTP monitoring: This method is used to monitor websites.
- PING monitoring: This method is used to monitor hosts.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan**.
3. In the **Vulnerability Scan** pane, click **Asset**.
4. On the **Asset** page, click the **Availability Monitoring** tab to view availability monitoring tasks.
5. Create an availability monitoring task.

Availability monitoring supports HTTP monitoring and PING monitoring.

- To create an HTTP monitoring task, perform the following steps:
 - a. Click **Add Monitoring**.
 - b. Click the **HTTP Monitoring** tab.

Asset

Asset Analysis Asset List Asset Import **Availability Monitoring** Custom Update Detection

[Back](#) **Add Monitoring**

HTTP Monitoring **PING Monitoring**

Monitoring Name

Monitoring Target

Monitoring Frequency

Request Method HEAD GET POST PUT

Alert Setting Response Time The response time over or equal to ms is regarded as anomaly.

Response Status When the status code is not it is regarded as an anomaly.

c. Configure the following parameters.

| Parameter | Description |
|-----------------------------|---|
| Monitoring Name | The name of the availability monitoring task. |
| Monitoring Target | The website that you want to monitor. |
| Monitoring Frequency | The interval at which you want to monitor the website. Valid values: 1 Minute , 5 Minute , 15 Minute , and 30 Minute . |
| Request Method | The request method that is used to send HTTP request packets. Valid values: HEAD , GET , POST , and PUT . |
| Alert Setting | <p>The policy based on which Apsara Stack Security reports alerts. If one of the following conditions is met, the website is unavailable:</p> <ul style="list-style-type: none"> Response Time: If the actual response time is greater than the specified value, an exception occurs. Response Status: If an unexpected status code is returned, an exception occurs. |

d. Click **Save**.

o To create a PING monitoring task, perform the following steps:

a. Click **Add Monitoring**.

b. Click the **PING Monitoring** tab.

Asset

HTTP Monitoring
PING Monitoring

Monitoring Name

Monitoring Target

Monitoring Frequency

Alert Setting
Response Time
The response time over or equal to

ms is regarded as anomaly.

Packet loss rate
Packet loss rate exceeding

% is regarded as anomaly.

c. Configure the following parameters.

| Parameter | Description |
|-----------------------------|--|
| Monitoring Name | The name of the availability monitoring task. |
| Monitoring Target | The host that you want to monitor. |
| Monitoring Frequency | The interval at which you want to monitor the host. Valid values: 1 Minute , 5 Minute , 15 Minute , and 30 Minute . |
| Alert Setting | The policy based on which Apsara Stack Security reports alerts. If one of the following conditions is met, the host is unavailable: <ul style="list-style-type: none">■ Response Time: If the actual response time is greater than the specified value, an exception occurs.■ Response Status: If an unexpected status code is returned, an exception occurs. |

d. Click **Save**.

6. Manage more than one availability monitoring task at a time.

You can manage more than one availability monitoring task in the monitoring task list at a time.

- Start more than one availability monitoring task at a time

Select more than one availability monitoring task and choose **Batch Monitoring Manage > Batch Open Monitoring**.

- Stop more than one availability monitoring task at a time

Select more than one availability monitoring task and choose **Batch Monitoring Manage > Batch Stop Monitoring**.

- Delete more than one availability monitoring task at a time

Select more than one availability monitoring task and choose **Batch Monitoring Manage > Batch Delete Monitoring**.

21.4.8.3.5. Manage custom update detection tasks

This topic describes how to manage custom update detection tasks.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan**.
3. In the **Vulnerability Scan** pane, click **Asset**.
4. On the **Asset** page, click **Custom Update Detection** to view custom update detection tasks.
5. Create a custom update detection task.
 - i. Click **Add Detection**.
 - ii. On the **Add Custom Update Detection** page, configure the parameters.

Asset

Asset Analysis
Asset List
Asset Import
Availability Monitoring
Custom Update Detection

Back
Add Custom Update Detection

Detection Name

Detection Target Public network asset ▼

Host Asset | Web Asset
Group Filter ▼
NATIP

NATIP
 Asset Group

No Data

Port Range Customize ▼

Multiple ports are separated by commas, and consecutive ports are represented by such as: 80, 8000-9000

Save
Cancel

| Parameter | Description |
|-------------------------|--|
| Detection Name | The name of the custom update detection task. |
| Detection Target | The asset that you want to detect. The value is fixed as Public network asset . <ol style="list-style-type: none"> a. Click Host Asset or Website Asset based on your business requirements. b. Select the assets that you want to detect from the asset list. <div style="background-color: #e1f5fe; padding: 5px; margin-top: 10px; border: 1px solid #cfe2f3;"> ? Note You can search for assets by Asset Group or IP. </div> |
| Port Range | The range of ports that you want to detect. Valid values: Customize , Full Port , Top100 , and Top1000 . <div style="background-color: #e1f5fe; padding: 5px; margin-top: 10px; border: 1px solid #cfe2f3;"> ? Note The Port Range parameter appears only after you set the Detection Target parameter to Host Asset. </div> <ul style="list-style-type: none"> ▪ Customize: You can specify custom ports to detect. ▪ Full Port: All ports are detected. ▪ Top 100: Top 100 ports are detected. ▪ Top 1000: Top 1000 ports are detected. |

iii. Click **Save**.

21.4.8.4. Risk management

21.4.8.4.1. Manage vulnerabilities

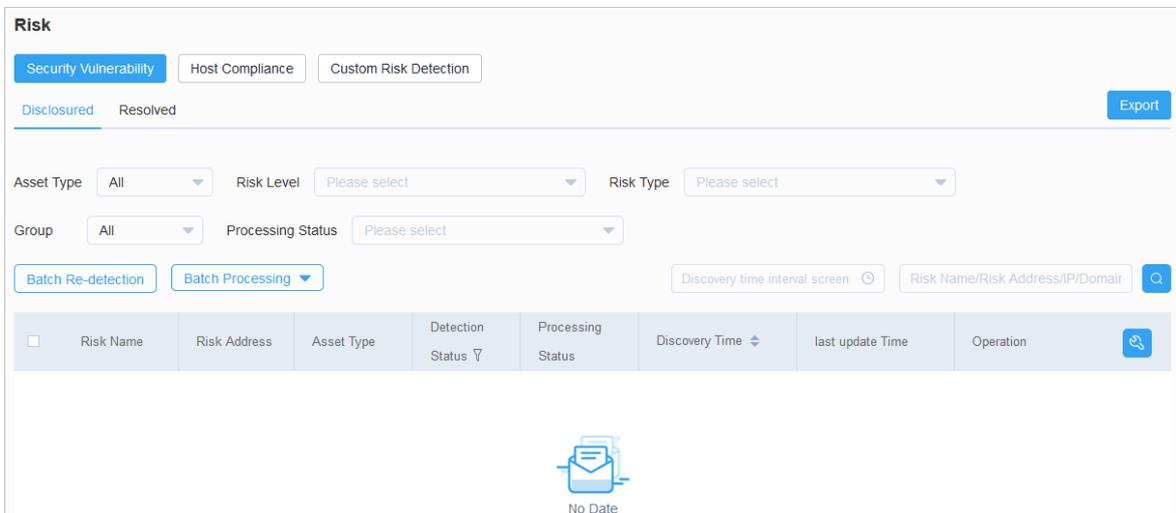
This topic describes how to view and handle the vulnerabilities that are detected by the vulnerability scan feature.

Context

On the **Security Vulnerability** tab, security administrators can view the vulnerabilities that are detected by the vulnerability scan feature.

Procedure

1. Log on to **Apsara Stack Security Center**.
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan**.
3. In the **Vulnerability Scan** pane, click **Risk**.
4. On the **Risk** page, click the **Security Vulnerability** tab to view vulnerabilities.



5. Click the **Disclosed** or **Resolved** tab to view unfixed vulnerabilities or fixed vulnerabilities.
6. View risk statistics. The statistics include **Unfixed Risks**, **Recovered Risks**, **Unconfirmed Risks**, **Confirmed Risks**, and **Ignored Risks**.
5. Specify search conditions to view specific vulnerabilities. The conditions include **Asset Type** and **Risk Level**.
6. Handle vulnerabilities.

Security administrators can analyze and confirm whether the vulnerabilities affect the security of assets based on the vulnerability information.

- o Confirm risks
 - if a vulnerability affects the security of assets, confirm the risk after the security vulnerability is fixed.
 - a. Find the vulnerability and click the  icon in the **Operation** column.
 - b. In the drop-down list, select **Confirm Risk**.
 - c. In the dialog box that appears, click **OK**.
- o Ignore risks
 - if a vulnerability is a false positive or does not affect the security of assets, ignore the risk.

- a. Find the vulnerability and click the  icon in the **Operation** column.
 - b. In the drop-down list, select **Ignore Risk**.
 - c. In the dialog box that appears, click **OK**.
7. Click **Export** to export the list of vulnerabilities to your computer.

21.4.8.4.2. Manage host compliance risks

This topic describes how to view and confirm host compliance risks.

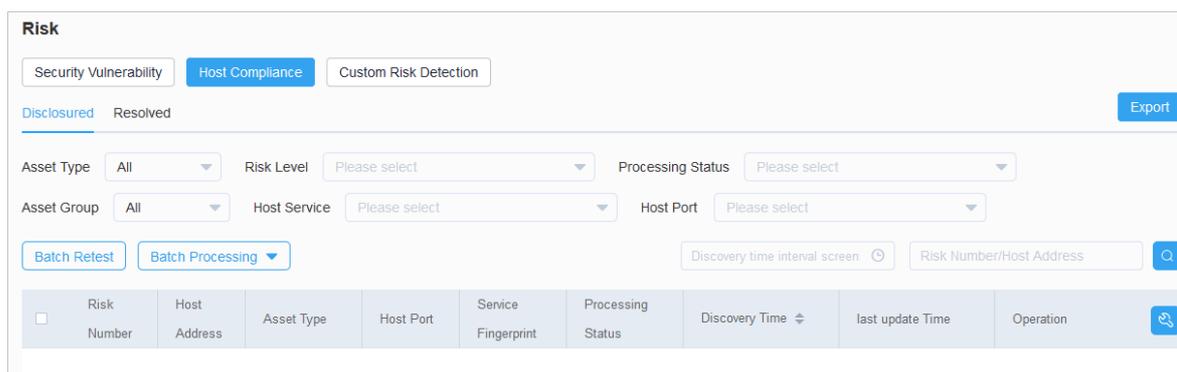
Context

On the **Host Compliance** tab, security administrators can view the host compliance issues that are detected by the vulnerability scan feature.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan**.
3. In the **Vulnerability Scan** pane, click **Risk**.
4. On the **Risk** page, click the **Host Compliance** tab to view host compliance risks.

You can click the **Disclosed** or **Resolved** tab to view **unfixed vulnerabilities** or **fixed vulnerabilities**.



5. Specify conditions to search for host compliance risks based on the conditions. The conditions include **Asset Type** and **Risk Level**.
6. Handle host compliance risks.

Security administrators can analyze and confirm whether host compliance risks affect the security of assets based on the risk information.

- o Confirm risks

If a host compliance risk affects the security of assets, harden the security of hosts and confirm the risk.

- a. Find the required risk and click the  icon in the **Actions** column.
- b. In the drop-down list, select **Confirm Risk**.
- c. In the message that appears, click **OK**.

- o Ignore risks

If a host compliance risk proves to be a false positive or does not affect the security of assets, ignore the risk.

- a. Find the required risk and click the  icon in the **Actions** column.

- b. In the drop-down list, select **Ignore Risk**.
 - c. In the message that appears, click **OK**.
7. Click **Export** to export the list of host compliance risks to your computer.

21.4.8.4.3. Manage external risks

This topic describes how to view and identify external risks, such as code leaks.

Prerequisites

A GitHub account is available, and the tokens of the account are obtained.

Context

On the **External Risk** page, security administrators can use the vulnerability scan feature to check whether the GitHub library has risks of code leaks.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan > Risk > External Risk**.
3. Click the **Code Disclosure** tab.
4. Add the assets that you want to monitor.

- i. In the **Asset Monitoring List** section, click **Token setting**.
 - ii. In the **Enter GitHub Token** dialog box, enter the tokens.

You can enter multiple tokens. Press Enter each time you enter a token.

- iii. Click **Add Asset** and select the external assets that you want to monitor.
5. View the assets that have risks of code leaks.
You can view the unfixed and fixed risks on the **Resolved** and **Disclosed** tabs.
6. Handle the risks of code leaks.

Security administrators can analyze and check whether the security of the assets is affected based on the risk information.

- o Confirm risks

If the risks of code leaks affect the security of your assets, harden the hosts and confirm the risks.

- o Ignore risks

If the risks of code leaks prove to be false positives or do not affect the security of your assets, ignore the risks.

21.4.8.4.4. Create a custom risk detection task

This topic describes how to create a custom risk detection task.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan**.
3. In the **Vulnerability Scan** pane, click **Risk**. On the Risk page, click the **Custom Risk Detection** tab.
4. On the Custom Risk Detection tab, click **Add Detection**.
5. On the **Add Custom Risk Detection** page, configure the parameters.

| Parameter | Description |
|-----------------------------|--|
| Detection Name | The name of the custom risk detection task. |
| Detection Target | The asset on which you want to perform risk detection. The value is fixed as Public network asset . |
| Emergency Detection | The switch that is used to enable or disable the emergency detection feature. If you enable this feature, you can select emergency detection items from the detection item list. |
| Basic Risk Detection | The switch that is used to enable or disable the basic risk detection feature. For more information about how to configure this feature, see Configure basic monitoring . |
| WEB Risk Detection | The switch that is used to enable or disable the web risk detection feature. For more information about how to configure this feature, see Configure web monitoring . |

6. Click **Save**.

21.4.8.5. Report management

21.4.8.5.1. Create a report

This topic describes how to create a report.

Context

A security administrator can create a report to view the security statuses of specific assets during a period of time and implement security measures as required.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan > Report**.
3. On the Report page, click the **Risk Report** tab and then the **Add Report** tab.
4. Configure the following parameters.

Report

Risk Report
Excel Report

Back
Add Report

Report Report

Content Name

Report Range Asset Info Single asset Asset Group All assets

Discovery

Time

Risk Setting Risk Range Resolved Risk Disclosed Risk

Risk Type Security Vulnerability Host Compliance

Create

| Parameter | | Description |
|----------------|----------------|---|
| Report Content | Report Name | The name of the report that you want to create. |
| Report Range | Asset Info | The scope of assets that you want to include in the report. Valid values: Single asset , Asset Group , and All assets . <ul style="list-style-type: none"> ◦ Single asset: Select an asset. ◦ Asset Group: Select an asset group and a tag. <div style="background-color: #e1f5fe; padding: 5px; margin: 5px 0;"> ? Note After you select an asset group and a tag, the assets in the group that have the selected tag are included in the report. </div> <ul style="list-style-type: none"> ◦ All assets: Select assets by tag. |
| | Discovery Time | The time range in which you want to perform risk detection. |
| Risk Setting | Risk Range | The scope of risks that you want to include in the report. Valid values: Resolved Risk and Disclosed Risk . |
| | Risk Type | The types of risks that you want to include in the report. Valid values: Security Vulnerability and Host Compliance . |

5. Click **Create**.

Result

After the report is created, it appears in the report list on the **Report** page.

21.4.8.5.2. Delete multiple reports at a time

This topic describes how to delete multiple reports at a time.

Context

You can delete multiple reports that are no longer required at a time to save storage space.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan > Report**.
3. Click **Risk Report**.
4. In the report list, select the reports that you want to delete.
5. Click **Batch Delete**.

21.4.8.6. Configuration management

21.4.8.6.1. Configure overall monitoring

This topic describes how to configure overall monitoring for the vulnerability scan feature. Overall monitoring includes Asset Monitoring Configuration, Base Risk Monitoring Configuration, External Risk Monitoring Configuration, and Scan Configuration.

Context

Overall monitoring allows you to configure detection features and the monitoring cycle for each detection feature.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan > Configuration > Monitoring Configuration > Overall Monitoring**.
3. In the **Monitoring Status** section, view the status of overall monitoring.
4. Configure detection features.

Detection features include **Asset Monitoring Configuration**, **Base Risk Monitoring Configuration**, **External Risk Monitoring Configuration**, and **Scan Configuration**.

In this step, the **Asset Monitoring Configuration** detection feature is used as an example.

- i. Turn on Asset Monitoring Configuration to enable the asset monitoring feature.
 - After the switch is turned on, the switch is in the **On** state. When in the On state, the switch is blue. After the switch is turned off, the switch is in the **Off** state. When in the Off state, the switch is gray.
 - You must turn on **Asset Monitoring Configuration** and **Base Risk Monitoring Configuration** to enable the two features. External Risk Monitoring Configuration and Scan Configuration are automatically enabled.
- ii. Configure the following parameters.

Asset Monitoring Configuration

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|--------------------------------------|--|
| Monitoring Item | <p>The item that you want to monitor. Valid value: Subdomain Discovery.</p> <p>If you want to import assets, you can set the Import Set parameter to Auto Import subdomains. Then, subdomains are automatically imported.</p> <p>If you select Subdomain Discovery, Apsara Stack Security regularly discovers subdomains for assets whose Import Set parameter is set to Auto Import subdomains.</p> |
| Monitoring Cycle | <p>The cycle based on which you want to perform detection. Valid values: customization, per week, and per month.</p> <ul style="list-style-type: none"> ▪ customization: Specify the interval at which you want to perform detection. Unit: days. ▪ per week: Specify the days of each week on which you want to perform detection. ▪ per month: Specify the days of each month on which you want to perform detection. |
| Detection Time | <p>The time when you want to perform detection. The time varies based on the value of the Monitoring Cycle parameter.</p> <ul style="list-style-type: none"> ▪ If you set the Monitoring Cycle parameter to customization, select a time range of the day in which you want to perform detection. ▪ If you set the Monitoring Cycle parameter to per week, select the days of each week and the time range in which you want to perform detection. ▪ If you set the Monitoring Cycle parameter to per month, select the days of each month and the time range in which you want to perform detection. <div style="background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note For example, if you set the Monitoring Cycle parameter to per week and select Monday to Sunday and 00:00:00 to 24:00:00 for the Detection Time parameter, Apsara Stack Security performs detection 24 hours a day, 7 days a week.</p> </div> |
| Port Range | <p>The ports on which you want to perform detection. Valid values: customization, Full port, TOP100, and TOP1000.</p> <ul style="list-style-type: none"> ▪ customization: Specify the ports to scan. ▪ Full port: Scan all ports. ▪ TOP100: Scan top 100 ports. You can click Add to add more ports. ▪ TOP1000: Scan top 1,000 ports. You can click Add to add more ports. |
| Host Alive Detection Settings | <p>The option that is used to check whether a host is running.</p> <p>By default, the ping feature is used to check whether a host is running. If the host has the ping feature disabled, the status of the host is detected based on top 20 ports and custom ports.</p> <p>To specify custom ports, click Settings. In the Host Alive Detection Settings dialog box, specify the ports in the Custom Port field.</p> |

Base Risk Monitoring Configuration

| Parameter | Description |
|-------------------------|--|
| Monitoring Item | <p>The item that you want to monitor. Valid values: Weak Password, Common Vulnerabilities, Baseline Monitoring, and Host Compliance.</p> <ul style="list-style-type: none"> ▪ Weak Password: Attackers can guess passwords or launch brute-force attacks to crack passwords. Then, the attackers can obtain relevant permissions. If you select this item, weak password vulnerabilities can be identified. ▪ Common Vulnerabilities: Web security vulnerabilities and CMS application vulnerabilities are included. If you select this item, common vulnerabilities can be identified. Then, you can install patches at the earliest opportunity. ▪ Baseline Monitoring: Risks in host configuration and account configuration are detected. ▪ Host Compliance: Host compliance risks are detected. |
| Monitoring Cycle | <p>The cycle based on which you want to perform detection. Valid values: customization, per week, and per month.</p> <ul style="list-style-type: none"> ▪ customization: Specify the interval at which you want to perform detection. Unit: days. ▪ per week: Specify the days of each week on which you want to perform detection. ▪ per month: Specify the days of each month on which you want to perform detection. |
| Detection Time | <p>The time when you want to perform detection. The time varies based on the value of the Monitoring Cycle parameter.</p> <ul style="list-style-type: none"> ▪ If you set the Monitoring Cycle parameter to customization, select a time range of the day in which you want to perform detection. ▪ If you set the Monitoring Cycle parameter to per week, select the days of each week and the time range in which you want to perform detection. ▪ If you set the Monitoring Cycle parameter to per month, select the days of each month and the time range in which you want to perform detection. <p> Note For example, if you set the Monitoring Cycle parameter to per week and select Monday to Sunday and 00:00:00 to 24:00:00 for the Detection Time parameter, Apsara Stack Security performs detection 24 hours a day, 7 days a week.</p> |

External Risk Monitoring Configuration

| Parameter | Description |
|------------------------|---|
| Monitoring Item | <p>The item that you want to monitor. Valid value: Code Disclosure.</p> <p>If you select Code Disclosure, Apsara Stack Security detects leaked source code of your assets.</p> |

| Parameter | Description |
|------------------|--|
| Monitoring Cycle | <p>The cycle based on which you want to perform detection. Valid values: customization, per week, and per month.</p> <ul style="list-style-type: none"> ▪ customization: Specify the interval at which you want to perform detection. Unit: days. ▪ per week: Specify the days of each week on which you want to perform detection. ▪ per month: Specify the days of each month on which you want to perform detection. |
| Detection Time | <p>The time when you want to perform detection. The time varies based on the value of the Monitoring Cycle parameter.</p> <ul style="list-style-type: none"> ▪ If you set the Monitoring Cycle parameter to customization, select a time range of the day in which you want to perform detection. ▪ If you set the Monitoring Cycle parameter to per week, select the days of each week and the time range in which you want to perform detection. ▪ If you set the Monitoring Cycle parameter to per month, select the days of each month and the time range in which you want to perform detection. <p> Note For example, if you set the Monitoring Cycle parameter to per week and select Monday to Sunday and 00:00:00 to 24:00:00 for the Detection Time parameter, Apsara Stack Security performs detection 24 hours a day, 7 days a week.</p> |

Scan Configuration

| Parameter | Description |
|---------------------|---|
| Risk Re-detection | The time at which you want to perform detection again. If risks are detected in specific assets, Apsara Stack Security scans the assets again each day at the time you specify. |
| Asset Scanning Rate | The scan rate. Valid values: Slow Mode , General Mode , Fast Mode , and Turbo Mode . |
| Risk Scanning Rate | The scan rate. Valid values: Slow Mode , General Mode , Fast Mode , and Turbo Mode . |
| UserAgent Setting | The User-Agent property. |

iii. Click Save.

21.4.8.6.2. Configure basic monitoring

This topic describes how to configure basic monitoring.

Context

Basic monitoring includes Weak Password Vulnerability Monitoring, Operation Security Vulnerability Monitoring, CMS Application Vulnerability Monitoring, and Baseline Monitoring.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan.**
3. In the **Vulnerability Scan** pane, click **Configuration.**
4. On the **Configuration** page, click the **Monitoring Configuration** tab.
5. On the **Basic Monitoring** tab of the Monitoring Configuration tab, click **Weak Password Vulnerability Monitoring** to configure rules to monitor weak passwords.
 - o By default, all monitoring items on weak passwords use the default weak password library.
 - o To disable a monitoring item, perform the following step:
 - Find the required metric and click the  icon in the **Operation** column.
 - o To enable a monitoring item, perform the following step:
 - Find the required monitoring item and click the  icon in the **Operation** column.
 - o To specify custom weak passwords for a monitoring item, perform the following steps. In this example, **MySQL Weak Password Vulnerability** is used.
 - a. In the **Default Weak Password** column, turn off the switch. The switch status changes to .
 - b. In the **Operation** column, click the .
 - c. In the **Customize MySQL Weak Password** dialog box, specify custom weak passwords.
 - d. Click **Yes.**
 - o To apply the same custom weak passwords to multiple monitoring items, perform the following steps:
 - a. In the **Default Weak Password** column, turn off the switches for the monitoring items that you want to apply the same custom weak passwords. The switch status changes to .

 **Note** If you want to apply a custom weak password to a monitoring item, you must turn off the switch in the **Default Weak Password** column of the monitoring item.

- b. Click **Tailored Overall Weak Password.**
 - c. In the **Tailored Overall Weak Password** dialog box, specify custom weak passwords.
 - d. Click **Yes.**
6. Click **Operation Security Vulnerability Monitoring** and configure O&M security vulnerability monitoring.

| No. | Description |
|-----|---|
| 1 | The switch that is used to enable or disable the Operation Security Vulnerability Monitoring feature. We recommend that you enable this feature to enhance system security. |
| 2 | The switch that is used to enable or disable a monitoring item. You can disable monitoring items based on your business requirements. |

7. Click **CMS Application Vulnerability Monitoring** and configure monitoring on content management system (CMS) application vulnerabilities.

| No. | Description |
|-----|-------------|
|-----|-------------|

| No. | Description |
|-----|--|
| 1 | The switch that is used to enable or disable the CMS Application Vulnerability Monitoring feature. We recommend that you enable this feature to enhance system security. |
| 2 | The switch that is used to enable or disable a monitoring item. You can disable monitoring items based on your business requirements. |

8. Click **Baseline Monitoring** and configure baseline monitoring.

To add a baseline monitoring item, perform the following steps:

- i. Click **Add**.
- ii. In the **Add Baseline** dialog box, configure the baseline monitoring item.

In this example, a baseline monitoring item is added to block Telnet-based access.

| Parameter | Description |
|-----------------------|---|
| Baseline Name | The name of the baseline monitoring item. Example: Block Telnet-based access. |
| Baseline Rule | The detection rule that is used by the baseline monitoring item. This rule checks whether hosts use disabled ports or run disabled services. Valid values: <ul style="list-style-type: none"> ▪ Port Disabled: ports that you want to disable. Example: 23. ▪ Service Disabled: services that you want to disable. Example: Telnet. |
| Baseline Range | The scope of assets to which the baseline monitoring item can be applied. Valid values: Private IP and NatIP . You must specify this parameter and select specific assets. |

- iii. Click **OK**.

21.4.8.6.3. Configure web monitoring

This topic describes how to configure web monitoring.

Context

Web monitoring allows you to configure monitoring items for monitoring web vulnerabilities. You can also configure conditions to block website crawlers.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan**.
3. In the **Vulnerability Scan** pane, click **Configuration**.
4. On the Configuration page, click the **Monitoring Configuration** tab and then the **Web Monitoring** tab to view existing rules.

 **Note** Default rules are created by the system. You can only view details of the default rules, but cannot modify or delete them.

5. Create a web monitoring rule.

- i. Click **Add Rule**.
- ii. On the **Add Web Monitoring Rule** page, configure the following parameters.

| Parameter | Description |
|---------------------------|--|
| Rule Name | The name of the web monitoring rule. |
| Monitoring Cycle | <p>The monitoring cycle. Valid values: Customization, Per Week, and Per Month.</p> <ul style="list-style-type: none"> ▪ Customization: Specify the interval at which you want to perform detection. ▪ Per Week: Specify the days of each week on which you want to perform detection. ▪ Per Month: Specify the days of each month on which you want to perform detection. |
| Detection Time | <p>The time when you want to perform detection. The time varies based on the value of the Monitoring Cycle parameter.</p> <ul style="list-style-type: none"> ▪ If you set the Monitoring Cycle parameter to Customization, select a time range of the day in which you want to perform detection. ▪ If you set the Monitoring Cycle parameter to Per Week, select the days of each week and the time range in which you want to perform detection. ▪ If you set the Monitoring Cycle parameter to Per Month, select the days of each month and the time range in which you want to perform detection. <div style="background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note For example, if you set the Monitoring Cycle parameter to Per Week and select Monday to Sunday and 00:00:00 to 24:00:00 for the Detection Time parameter, Apsara Stack Security performs detection 24 hours a day, 7 days a week.</p> </div> |
| Monitoring Options | The type of web vulnerabilities that you want to monitor. Supported operations: Select All , Inverse , and Clear . |
| UserAgent | <p>The User-Agent field of the HTTP request packet.</p> <p>The User-Agent field identifies the application type, operating system, software developer, and version number of the proxy software that initiates requests.</p> |
| Cookies | The cookie parameters. |
| Key Page | The web directories or pages that you want to monitor. |
| Excluded Page | The web directories or pages that you do not want to monitor. |
| Crawler Depth | The capturing depth of crawlers. Valid values: 10 , 15 , and 30 . |
| URL Numbers | The number of URLs that are used for crawling. Valid values: 500 , 1000 , and 2000 . |
| scanning Frequency | The scan frequency of web monitoring. Valid values: Request 10 Times Per Second and Request 15 Times Per Second . |

- iii. Click **Yes**.

6. Manage the web monitoring rule.

| Icon | Description |
|---|---|
|  | Modify the rule. |
|  | Delete the rule. |
| Batch Delete | If you want to delete multiple rules, select the rules you want to delete and click Batch Delete . |

21.4.8.6.4. Configure a whitelist

This topic describes how to configure a whitelist.

Context

Apsara Stack Security does not scan the assets that are added to a whitelist. Before you add assets to a whitelist, make sure that the assets are secure.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan**.
3. In the **Vulnerability Scan** pane, click **Configuration**.
4. On the Configuration page, click the **Monitoring Configuration** tab and then the **Whitelist** tab to view the assets that are added to a whitelist.

 **Note** By default, the whitelist feature is enabled. If you do not want to use the whitelist feature, turn off the switch in the upper-right corner.

5. Add assets to a whitelist.
 - i. Click **Add**.
 - ii. In the **Add Whitelist** dialog box, configure the parameters.
 - If you select **Asset Group** for the **Whitelist** parameter, select a group from the second drop-down list. The assets in this group are added to the whitelist.
 - If you select **Customization** for the **Whitelist** parameter, enter the IP addresses or URLs that you want to add to the whitelist in the field that appears.
 - iii. Click **Yes**.
6. Manage the assets that are added to a whitelist.
 - Delete an asset from a whitelist

Find the required asset and click the  icon in the **Operation** column.
 - Remove multiple assets from a whitelist at a time

Select the required assets and click **Batch Delete**.

21.4.8.6.5. Configure a scan engine for internal assets

This topic describes how to configure a scan engine for internal assets, such as the assets of a virtual private cloud (VPC).

Context

You must add a scan engine for a VPC before you can scan for vulnerabilities on the assets of the VPC.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Threat Detection > Vulnerability Scan > Configuration**. On the page that appears, click the **Scan Engine Manage** tab. Then, click the **Private-sector assets** tab.
3. Click the name of the VPC whose assets you want to scan.
4. Click **Add Scan Engine**.
5. In the **Add Scan Engine** dialog box, select a vSwitch for the VPC from the **vSwitch** drop-down list.
6. Click **OK**.

21.4.9. Create a security report

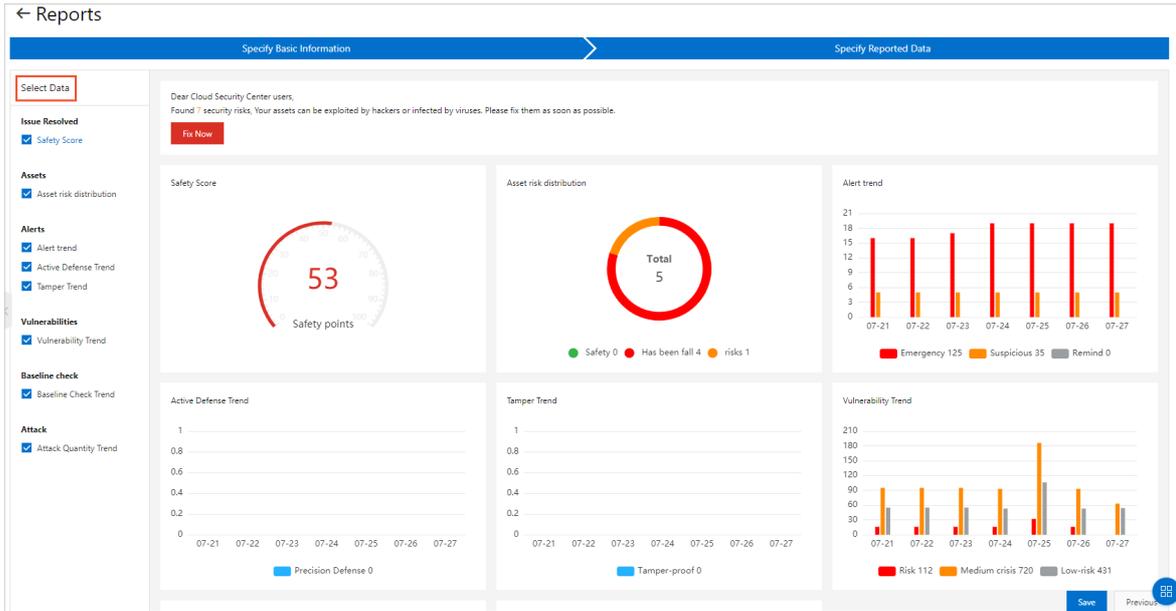
Security reports help monitor the security status of your assets. You can specify the content, types of statistics, and email addresses of recipients to create a security report. This topic describes how to create a security report.

Procedure

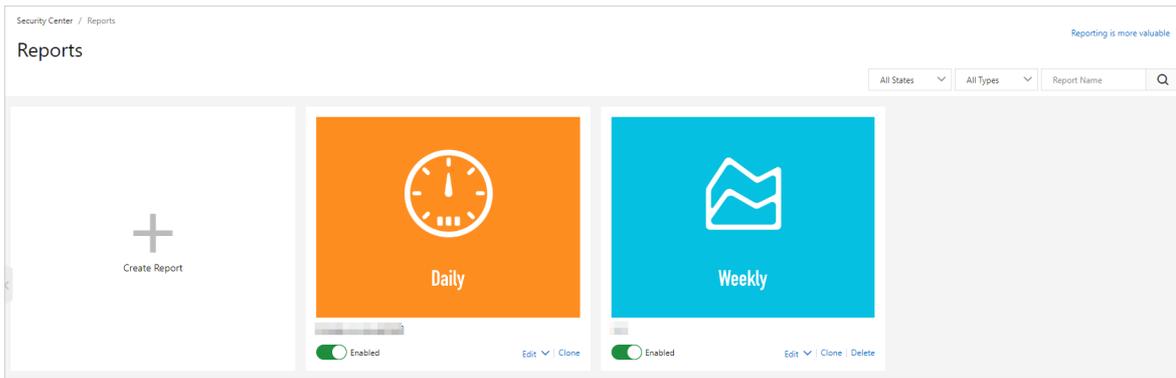
1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Threat Detection > Security Reports**.
3. On the **Reports** page, click **Create Report**.

 **Notice** In addition to the default security report created by Apsara Stack Security, you can create a maximum of nine security reports.

4. In the **Specify Basic Information** step, configure the parameters.
Configure the following parameters:
 - **Report Name**: Enter a name for the security report.
 - **Report Type**: Select a report type from the drop-down list. Valid values: *Daily*, *Weekly*, *Monthly*, and *Custom*.
If you select *Custom*, you must also set the **Data Collection Period** parameter to specify the cycle on which data is collected.
 - **Language**: Select a natural language for the report. Valid values: **简体中文** and **English**.
5. Click **Next**.
6. In the **Specify Reported Data** step, select the types of data that you want to view in the security report. You can select assets, alerts, vulnerabilities, baselines, attacks, and other data related to security operations.



- 7. Click **Save**. The security report is created. You can view the newly created security report on the **Reports** page.

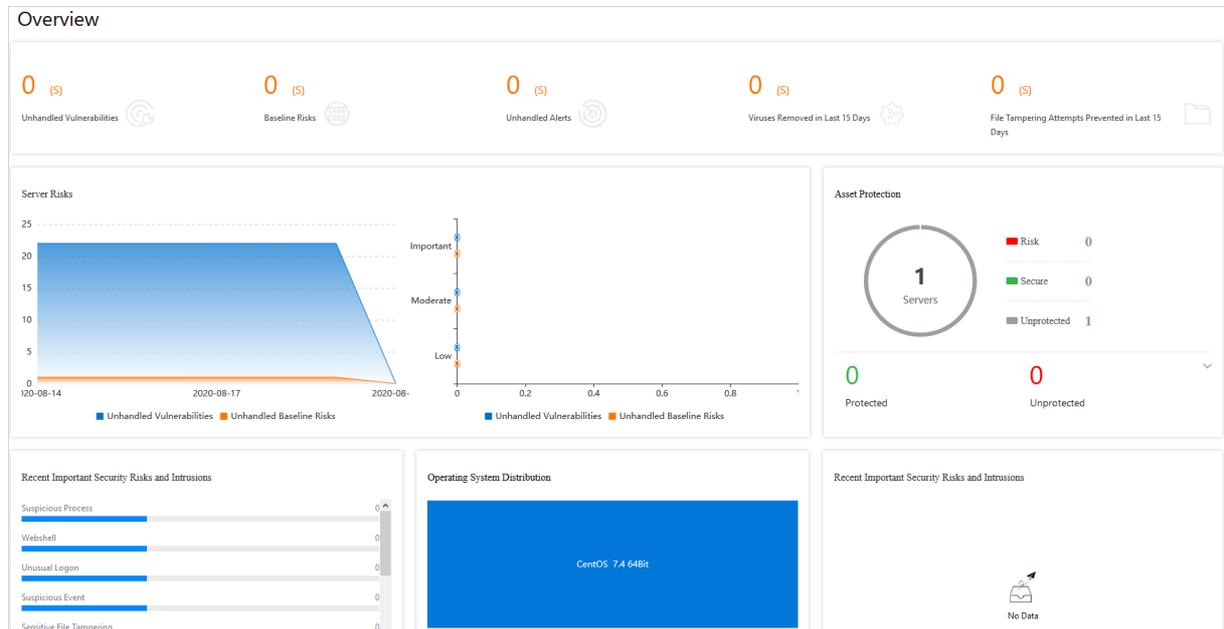


21.5. Server security

21.5.1. Server security overview

This topic describes how to view the details about the security of servers on the server security overview page of Apsara Stack Security Center. This helps security administrators understand the security status of the servers. The servers refer to servers on the cloud.

To view the details about the security of servers, log on to Apsara Stack Security Center and choose **Server Security > Overview**. On the page that appears, you can view detailed information on the following sections: overall statistics, Server Risks, Asset Protection, Operating System Distribution, and Recent Important Security Risks and Intrusions.



- **Overall statistics:** This section displays the numbers of security vulnerabilities and security events on servers. For security vulnerabilities, you can view **Unhandled Vulnerabilities** and **Baseline Risks**. For security events, you can view **Unhandled Alerts**, **Viruses Removed in Last 15 Days**, and **File Tampering Attempts Prevented in Last 15 Days**.
- **Server Risks:** This section displays the number of unhandled vulnerabilities, the number of baseline risks, and the distribution of risk levels.
- **Asset Protection:** This section displays the number of protected servers and the number of offline servers.
- **Recent Important Security Risks and Intrusions:** This section displays the recent important risks and events on your servers. You can click a risk or an event to view the details.
- **Operating System Distribution:** This section displays your servers by operating system.

21.5.2. Server fingerprints

21.5.2.1. Manage listening ports

This topic describes how to view the information about the listening port of a server. The information helps you identify suspicious listening behavior.

Context

This topic is suitable for the following scenarios:

- Check for servers that listen on a specific port.
- Check for ports that a specific server listens.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Servers > Server Fingerprints**.
3. On the **Asset Fingerprints** page, click the **Port** tab to view listening ports, network protocols, and server information.

You can search for a port by using the port number, server process name, server name, or server IP address.

In the server information list, you can view the **process**, **IP address**, and **latest scan time** of a server.

21.5.2.2. Manage software versions

This topic describes how to regularly view and collect the software version information about a server. This helps you check your software assets.

Context

This topic covers the following scenarios:

- Check for software assets that are installed without authorization.
- Check for outdated versions of software assets.
- Locate affected assets if vulnerabilities are detected.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Servers > Server Fingerprints**. On the page that appears, click the **Software** tab.
3. View all the **software assets** that are in use and the **numbers of servers** that use the software assets.
You can search for specific software by using its name, version, installation directory, server name, or IP address.
4. Click software to view the details, such as the software versions and the servers that use the software.

You can click the  icon in the upper-right corner to download a software version table to your computer for subsequent asset check.

21.5.2.3. Manage processes

This topic describes how to regularly collect the process information on a server and record changes. This way, you can view process information and historical process changes.

Context

This task is suitable for the following scenarios:

- Check for servers on which a specific process runs.
- Check for processes that run on a specific server.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Servers > Server Fingerprints**. On the page that appears, click the **Process** tab.
3. View all running processes and the numbers of servers on which these processes run.
You can search for a process by using the **process name**, **running user**, **startup parameter**, or **server name or IP address**.
4. Click the name of a process to view the details of the process, such as the servers, paths, and startup parameters.

21.5.2.4. Manage account information

This topic describes how to regularly collect the account information on a server and record changes to the accounts. This way, you can check your accounts and view historical changes to your accounts.

Context

You can use the information collected in this topic for the following scenarios:

- Check for servers on which a specific account is created.
- Check for accounts that are created on a server.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Servers > Server Fingerprints**. On the page that appears, click the **Account** tab.
3. View all the logged-on accounts and the numbers of servers on which the accounts are created.
You can search for an account by using the account name, root permissions, server name, or IP address.
4. Click an account name to view the details, such as the server information, root permissions, and user group.

21.5.2.5. Manage scheduled tasks

This topic describes how to view scheduled tasks on servers.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Servers > Server Fingerprints**. On the page that appears, click the **Scheduled Tasks** tab.
3. View the paths of all tasks and the numbers of servers that run these tasks.
You can search for a task by using the path, server name, or IP address.
4. Click a task path to view the details, such as the servers, executed commands, and task cycles.

21.5.2.6. Set the fingerprint collection frequency

You can set the frequency at which the data of running processes, system accounts, listening ports, and software versions is collected.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Servers > Server Fingerprints**. In the upper-right corner of the page that appears, click **Settings**.
3. Select the collection frequency from each drop-down list.
4. Click **OK**.

21.5.3. Threat protection

21.5.3.1. Vulnerability management

21.5.3.1.1. Manage Linux vulnerabilities

This topic describes how to manage Linux vulnerabilities.

Context

Apsara Stack Security automatically scans the software that is installed on your servers to detect the vulnerabilities provided in the Common Vulnerabilities and Exposures (CVE) list. Apsara Stack Security also sends you alerts about the detected vulnerabilities. In addition, Apsara Stack Security provides commands that you can use to fix vulnerabilities and allows you to verify vulnerability fixes.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Threat Prevention > Vulnerabilities**. On the page that appears, click the **Linux Software** tab.
3. View the detected Linux vulnerabilities.

 **Note** You can search for specific vulnerabilities by using the search and filter features.

4. Click a vulnerability. In the panel that appears, you can view details about the vulnerability and the servers that are affected by the vulnerability.

 **Note** You can find affected servers by using the search and filter features.

- **Detail:** This tab displays the basic information about the vulnerability, including the name, Common Vulnerability Scoring System (CVSS) score, and description.
 - **Pending vulnerability:** This tab displays the servers that are affected by the vulnerability.
5. Handle the vulnerability based on its impact.

Actions on vulnerabilities

| Action | Impact |
|--------|--|
| Fix | Select this action to fix the vulnerability. |
| Ignore | Select this action to ignore the vulnerability. The system no longer reports alerts for ignored vulnerabilities. |
| Verify | Click Verify to verify the vulnerability fix. If you do not manually verify a fix, the system automatically verifies the fix within 48 hours after the vulnerability is fixed. |

You can manage a vulnerability for one or multiple affected servers at a time.

- To manage a vulnerability for one affected server, find the server and select an action from the **Actions** column of the server.
- To manage a vulnerability for multiple affected servers, select the servers and select an action in the lower-left corner.

21.5.3.1.2. Manage Windows vulnerabilities

This topic describes how to manage Windows vulnerabilities.

Context

Apsara Stack Security automatically checks whether the latest Microsoft updates are installed on your servers, and notifies you of the detected vulnerabilities. Apsara Stack Security also automatically detects and fixes major vulnerabilities on your servers.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Threat Prevention > Vulnerabilities**. On the page that appears, click the **Windows System** tab.
3. View the detected Windows vulnerabilities.

 **Note** You can search for specific vulnerabilities by using the search and filter features.

4. Click a vulnerability. In the panel that appears, you can view details about the vulnerability and the servers that are affected by the vulnerability.

 **Note** You can find affected servers by using the search and filter features.

- **Detail:** This tab displays the basic information about the vulnerability.
 - **Pending vulnerability:** This tab displays the servers that are affected by the vulnerability.
5. Handle the vulnerability based on its impact. [Actions on vulnerabilities](#) describes the actions.

Actions on vulnerabilities

| Action | Impact |
|--------|--|
| Fix | Select this action to fix the vulnerability. The system caches an official Windows patch in the cloud. Your server can automatically download the patch for updates. |
| Ignore | Select this action to ignore the vulnerability. The system no longer reports alerts for ignored vulnerabilities. |
| Verify | Click Verify to verify the vulnerability fix. |

You can manage a vulnerability for one or multiple affected servers at a time.

- To manage a vulnerability for one affected server, select an action from the **Actions** column of the server.
- To manage a vulnerability for multiple affected servers, select the servers and select an action in the lower-left corner.

21.5.3.1.3. Manage Web CMS vulnerabilities

This topic describes how to manage Web CMS vulnerabilities.

Context

The Web CMS vulnerability detection feature obtains information about the latest vulnerabilities and provides patches in the cloud. This helps you detect and fix vulnerabilities.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Threat Prevention > Vulnerabilities**. On the page that appears, click the **Web CMS** tab.
3. View vulnerabilities.

 **Note** You can search for specific vulnerabilities by using the search and filter features.

4. Click a vulnerability. In the panel that appears, you can view details about the vulnerability and servers that are affected by the vulnerability.

 **Note** You can find affected servers by using the search and filter features.

- Handle the vulnerability based on its impact. [Actions on vulnerabilities](#) describes the actions.

Actions on vulnerabilities

| Action | Description |
|--------|---|
| Fix | If you select this action, the system replaces the web files affected by the vulnerability on your server to fix the Web CMS vulnerability. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note Before you fix the vulnerability, we recommend that you back up the web files affected by the vulnerability. For more information about the paths of the web files, click Details in the Actions column.</p> </div> |
| Verify | After a vulnerability is fixed, you can click Verify to verify the fix. If you do not manually verify the fix of a vulnerability, the system automatically verifies the fix within 48 hours after the vulnerability is fixed. |

You can manage a vulnerability for one or more affected servers at a time.

- To manage a vulnerability for one affected server, find the server and select an action in the **Actions** column of the asset.
- To manage a vulnerability for multiple affected servers, select the servers and select an action in the lower-left corner.

21.5.3.1.4. Manage urgent vulnerabilities

This topic describes how to manage urgent vulnerabilities.

Context

Apsara Stack Security automatically detects vulnerabilities on servers, such as the unauthorized Redis access vulnerability and Struts S2-052 vulnerability, and generates vulnerability alerts. After you fix a vulnerability, you can also check whether the fix is successful.

Procedure

- Log on to [Apsara Stack Security Center](#).
- In the left-side navigation pane, choose **Server Security > Threat Prevention > Vulnerabilities**. On the page that appears, click the **Emergency** tab.

- View all vulnerabilities.

You can search for a specific vulnerability by using the search and filter features.

- Click a vulnerability. In the panel that appears, view the details in the following sections: **Details**, **Suggestions**, and **Affected Assets**.

You can find affected servers by using the search and filter features.

- Handle the vulnerability based on the impact of the vulnerability. [Actions on vulnerabilities](#) describes the actions that you can use to handle the vulnerability.

Follow the instructions to manually fix the vulnerabilities on the **Emergency** tab.

Actions on vulnerabilities

| Action | Description |
|--------|--|
| Ignore | Select this action to ignore the vulnerability. The system no longer reports alerts for ignored vulnerabilities. |
| Verify | Click Verify to verify the vulnerability fix. If you do not manually verify a fix, the system automatically verifies the fix within 48 hours after the vulnerability is fixed. |

You can manage a vulnerability for one or more affected servers at a time.

- To manage a vulnerability for one affected server, find the server and select an action from the **Actions** column of the server.
- To manage a vulnerability for multiple affected servers, select the servers and select an action in the lower-left corner.

21.5.3.1.5. Configure vulnerability management policies

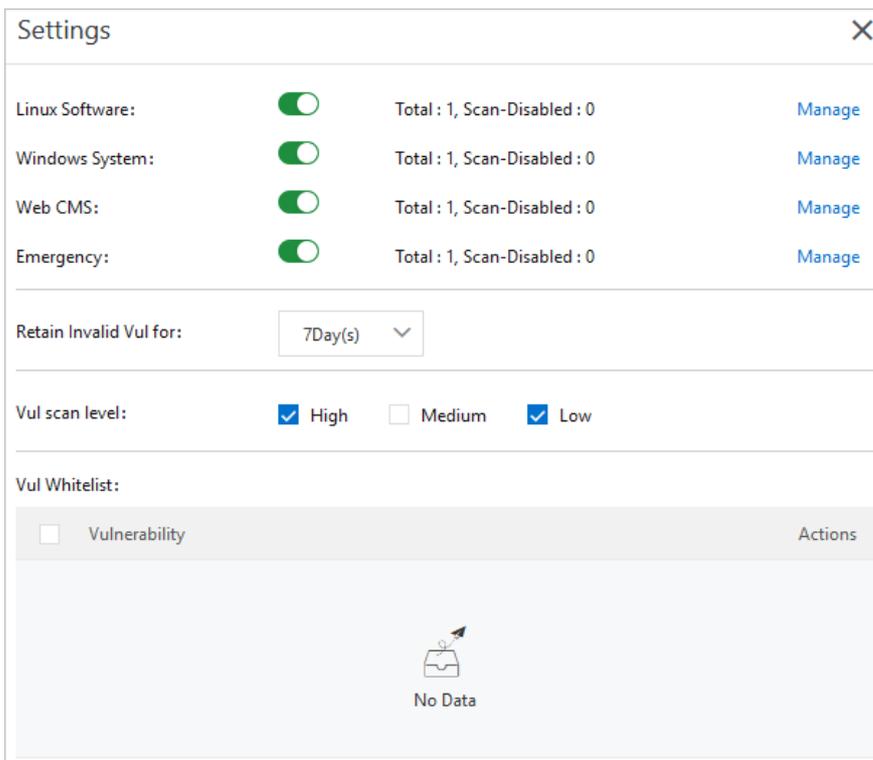
You can enable or disable automatic detection for different types of vulnerabilities and enable vulnerability detection for specific servers. You can also set a time duration for which invalid vulnerabilities are retained and configure a vulnerability whitelist.

Context

A vulnerability whitelist allows you to exclude vulnerabilities from the detection list. You can add multiple vulnerabilities in the vulnerability list to the whitelist. The system does not detect vulnerabilities that are added to the whitelist. You can manage the vulnerability whitelist on the vulnerability settings page.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Threat Prevention > Vulnerabilities**.
3. In the upper-right corner, click **Settings** to configure vulnerability management policies.



- Select a vulnerability type and enable or disable detection for vulnerabilities of this type.
- Click **Manage** next to a vulnerability type and specify the servers on which vulnerabilities of this type are detected.
- Select a time duration during which invalid vulnerabilities are retained. For example, you can select 7 days, 30 days, or 90 days.

 **Note** If you do not take an action on a detected vulnerability, the system determines that the alert is invalid. The system deletes the vulnerability after the specified duration.

- Select the vulnerability severities for scanning.
 - **High:** Vulnerabilities of this severity must be fixed as soon as possible.
 - **Medium:** Vulnerabilities of this severity can be fixed later.
 - **Low:** Vulnerabilities of this severity are less urgent and can be fixed later.
- Select vulnerabilities in the whitelist and click **Remove** to enable the system to detect these vulnerabilities and generate alerts for these vulnerabilities.

21.5.3.2. Baseline check

21.5.3.2.1. Baseline check overview

The baseline check feature automatically checks the security configurations on servers and provides detailed check results and suggestions for baseline reinforcement.

Description

After you enable the baseline check feature, Apsara Stack Security automatically checks for risks related to the operating systems, accounts, databases, passwords, and security compliance configurations of your servers, and provides reinforcement suggestions. For more information, see [Baselines](#).

By default, a full baseline check is automatically performed from 00:00 to 06:00 every day. You can create and manage scan policies for baseline checks. When you create or modify a policy, you can specify the baselines, interval, and time period, and select the servers to which you want to apply this policy. For more information, see [Add a custom baseline check policy](#).

Precautions

By default, the following baselines are disabled. To check these baselines, make sure that these baselines do not affect your business and select them when you customize a scan policy.

- Baselines related to weak passwords for specific applications such as MySQL, PostgreSQL, and SQL Server

Note If these baselines are enabled, the system attempts to log on to servers with weak passwords. The logon attempts consume server resources and generate a large number of logon failure records.

- Baselines related to China classified protection of cybersecurity
- Baselines related to the Center for Internet Security (CIS) standard

Baselines

| Category | Baseline |
|-------------------|---|
| High risk exploit | <ul style="list-style-type: none"> • High risk exploit - CouchDB unauthorized access high exploit risk • High risk exploit - Docker unauthorized access high vulnerability risk • High risk exploit - Elasticsearch unauthorized access high exploit vulnerability risk • High risk exploit - Memcached unauthorized access high exploit vulnerability risk • High risk exploit - Apache Tomcat AJP File Read/Inclusion Vulnerability • High risk exploit - ZooKeeper unauthorized access high exploit vulnerability risk |

| Category | Baseline |
|----------|--|
| | <p>Security baseline check against the Alibaba Cloud standard:</p> <ul style="list-style-type: none"> • Alibaba Cloud Standard-Aliyun Linux 2 Security Baseline Check • Alibaba Cloud Standard - CentOS Linux 6 Security Baseline Check • Alibaba Cloud Standard - CentOS Linux 7 Security Baseline Check • Alibaba Cloud Standard - Debian Linux 8 Security Baseline • Alibaba Cloud Standard - Redhat Linux 6 Security Baseline Check • Alibaba Cloud Standard - Redhat Linux 7 Security Baseline Check • Alibaba Cloud Standard - Ubuntu Security Baseline Check • Alibaba Cloud Standard - Windows Server 2008 R2 Security Baseline Check • Alibaba Cloud Standard - Windows 2012 R2 Security Baseline • Alibaba Cloud Standard - Windows 2016/2019 R2 Security Baseline <p>Security baseline check against the CIS standard:</p> <ul style="list-style-type: none"> • Alibaba Cloud Aliyun Linux 2 CIS Benchmark • CIS CentOS Linux 6 LTS Benchmark • CIS CentOS Linux 7 LTS Benchmark • CIS Debian Linux 8 Benchmark • CIS Ubuntu Linux 14 LTS Benchmark • CIS Ubuntu Linux 16/18 LTS Benchmark • CIS Microsoft Windows Server 2008 R2 Benchmark • CIS Microsoft Windows Server 2012 R2 Benchmark • CIS Microsoft Windows Server 2016/2019 R2 Benchmark |

| Category | Baseline |
|--|--|
| <p>CIS and China's Protection of Cybersecurity</p> | <p>Baseline check on compliance of China classified protection of cybersecurity level III:</p> <ul style="list-style-type: none"> • Aliyun Linux 2 Baseline for China classified protection of cybersecurity-Level III • CentOS Linux 6 Baseline for China classified protection of cybersecurity-Level III • CentOS Linux 7 Baseline for China classified protection of cybersecurity-Level III • Debian Linux 8 Baseline for China classified protection of cybersecurity-Level III • Redhat Linux 6 Baseline for China classified protection of cybersecurity-Level III • Redhat Linux 7 Baseline for China classified protection of cybersecurity-Level III • SUSE Linux 10 Baseline for China classified protection of cybersecurity-Level III • SUSE Linux 11 Baseline for China classified protection of cybersecurity-Level III • SUSE Linux 12 Baseline for China classified protection of cybersecurity-Level III • Ubuntu 14 Baseline for China classified protection of cybersecurity-Level III • Waiting for Level 3-Ubuntu 16/18 compliance regulations inspection • China's Level 3 Protection of Cybersecurity - Windows Server 2008 R2 Compliance Baseline Check • Windows 2012 R2 Baseline for China classified protection of cybersecurity-Level III • Windows 2016/2019 R2 Baseline for China classified protection of cybersecurity-Level III |

| Category | Baseline |
|-------------------------|--|
| Best security practices | <ul style="list-style-type: none"> • Alibaba Cloud Standard-Aliyun Linux 2 Security Baseline Check • Alibaba Cloud Standard - Apache Security Baseline Check • Alibaba Cloud Standard - CentOS Linux 6 Security Baseline Check • Alibaba Cloud Standard - CentOS Linux 7/8 Security Baseline Check • Alibaba Cloud Standard - Debian Linux 8 Security Baseline • Alibaba Cloud Standard - IIS 8 Security Baseline Check • Alibaba Cloud Standard - Memcached Security Baseline Check • Alibaba Cloud Standard - MongoDB 3.x Security Baseline Check • Alibaba Cloud Standard - Mysql Security Baseline Check • Alibaba Cloud Standard - Nginx Security Baseline Check • Alibaba Cloud Standard - Redhat Linux 6 Security Baseline Check • Alibaba Cloud Standard - Redhat Linux 7 Security Baseline Check • Alibaba Cloud Standard - Redis Security Baseline Check • Alibaba Cloud Standard - Ubuntu Security Baseline Check • Alibaba Cloud Standard - Windows Server 2008 R2 Security Baseline Check • Alibaba Cloud Standard - Windows 2012 R2 Security Baseline • Alibaba Cloud Standard - Windows 2016/2019 R2 Security Baseline • Alibaba Cloud Standard-Apache Tomcat Security Baseline |

| Category | Baseline |
|---------------|---|
| Weak password | <ul style="list-style-type: none"> Weak Password-MongoDB Weak Password baseline(support version 2. X) Weak password - Ftp login weak password baseline Weak password - Linux system login weak password baseline Weak password - MongoDB login weak password baseline Weak password - SQL Server DB login weak password baseline Weak password - Mysql DB login weak password baseline Weak password - Mysql DB login weak password baseline(Windows version) Weak password - PostgreSQL DB login weak password baseline Weak password - Redis DB login weak password baseline Weak password - rsync login weak password baseline Weak password - svn login weak password baseline |

21.5.3.2.2. Configure baseline check policies

This topic describes how to create, modify, and delete baseline check policies and how to specify baseline check levels.

Context

By default, the baseline check feature uses the **default policy** to check the baseline risks of assets. You can also customize baseline check policies based on your business requirements. For example, you can customize a baseline check policy to check the compliance with China classified protection of cybersecurity-Level II.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Threat Prevention > Baseline Check**.
3. In the upper-right corner of the page that appears, click **Manage Policies**. In the **Manage Policies** panel, create, modify, or delete a baseline check policy. You can also modify the default policy.
 - o In the upper-right corner of the panel, click **Create Policy** to customize a baseline check policy. Then, click **Ok**.

| Parameter | Description |
|--------------------|---|
| Policy Name | Enter a policy name. |
| Schedule | Set the time interval for scheduled scan tasks to 1 Day(s), 3 Day(s), 7 Day(s), or 30 Day(s). Then, select one of the following time ranges for scheduled scan tasks: 00:00 to 06:00, 06:00 to 12:00, 12:00 to 18:00, and 18:00 to 24:00. |
| Check Items | Select the baseline items that need to be checked from the following categories: High risk exploit, Container security, CIS and China's Protection of Cybersecurity, Best security practices, and Weak password. |

| Parameter | Description |
|-----------|--|
| Servers | Select the server groups to which you want to apply the baseline check policy. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> ? Note By default, newly purchased servers are classified into the Default group under Asset Groups. To apply this policy to new servers, select Default. </div> |

- Click **Edit** or **Delete** next to the new policy to modify or delete it.

? **Note** You cannot restore a policy after you delete it.

- Find the **Default** policy and click **Edit** in the **Actions** column to modify the server groups to which the default policy is applied.

? **Note** You cannot delete the default policy or modify the baseline items of the default policy. You can only modify the server groups to which the default policy is applied.

- In the lower part of the **Manage Policies** panel, specify the baseline check levels. Valid values: High, Medium, and Low.

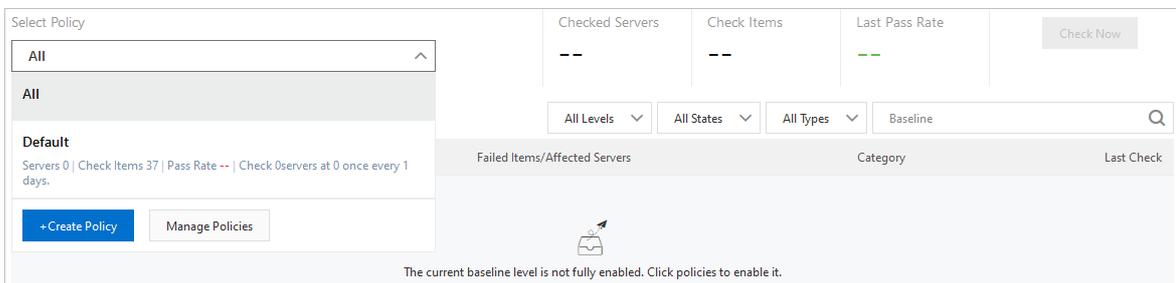
4. Click **OK**.

21.5.3.2.3. View baseline check results and manage failed check items

Apsara Stack Security Center provides detailed baseline check results and suggestions on how to manage failed check items. This topic describes how to view baseline check results and manage failed check items in Apsara Stack Security Center. The check results include affected assets, checked items, and suggestions.

View the summary of baseline check results

- Log on to [Apsara Stack Security Center](#).
- In the left-side navigation pane, choose **Server Security > Threat Prevention > Baseline Check**.
- In the upper part of the **Baseline Check** page, view the summary of baseline check results. You can filter data by policy.



You can select a policy from the **Select Policy** drop-down list to view the following information:

- Checked Servers:** The number of servers on which the baseline check runs. These servers are specified in the selected baseline check policy.
- Check Items:** The number of **check items** specified in the selected baseline check policy.
- Last Pass Rate:** The pass rate of the last baseline check.

If the number below **Last Pass Rate** is green, the pass rate of the checked servers is high. If the number is red, a large number of failed check items have been detected on the checked servers. We recommend that you view the check result details and manage the failed check items.

View all baselines

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Threat Prevention > Baseline Check**.
3. Select **All** from the **Select Policy** drop-down list.
The **Baseline Check** page displays details about all baselines, including **Baseline**, **Checked Item**, **Failed Items/Affected Servers**, **Category**, and **Last Check**.

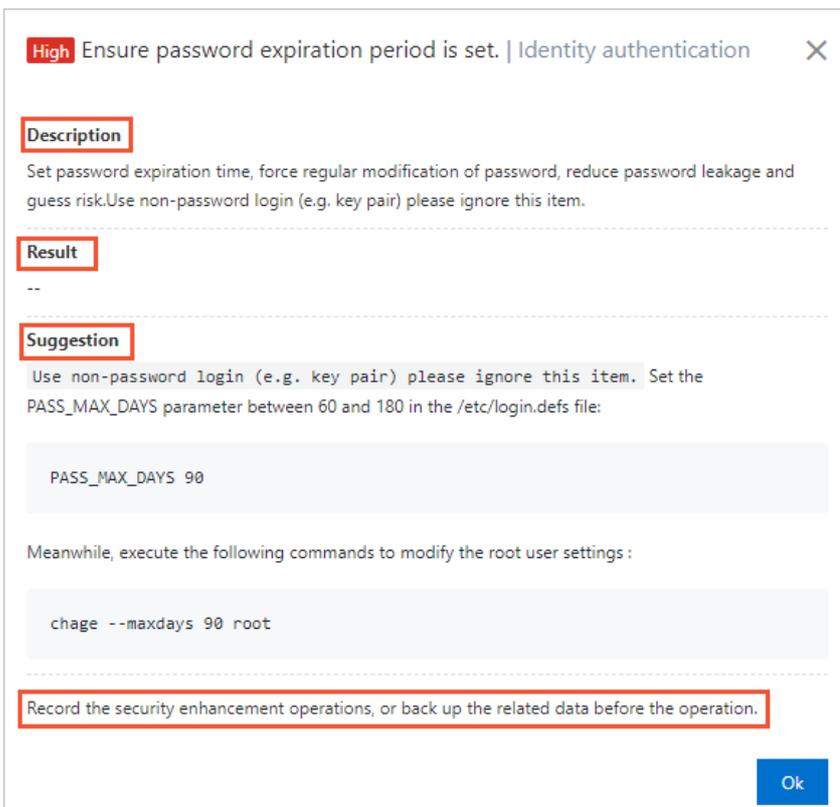
 **Note** You can also select a baseline check policy from the **Select Policy** drop-down list to view the baselines specified in this policy.

View details about a baseline

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Threat Prevention > Baseline Check**.
3. In the **Baseline** column, click a baseline to view its details.
Baseline details include the affected assets and numbers of **Passed Items** and **At-Risk Items**.
4. In the details panel, manage the failed check items.
 - Find the asset that you want to manage and click **View** in the **Actions** column to open the **At-Risk Items** panel.
 - You can click **Verify** in the **Actions** column to check whether the baseline risks of an asset have been managed. If the verification is passed, the number of **At-Risk Items** is reduced, and the status of the check items change to **Passed**.

View details about a failed check item

1. Find the baseline that you want to manage and click it. In the panel that appears, find the asset that you want to manage and click **View** in the **Actions** column to view failed check items.
You can view the check items of the asset and the statuses of the check items. The status can be **Passed** or **Failed**.
2. You can click **Details** in the **Actions** column to view the description, result, and suggestion for this check item.



Note We recommend that you follow the suggestions to manage check items whose status is **Failed** at the earliest opportunity, especially high-risk check items.

Manage failed check items

In the At-Risk Items panel, manage failed check items.

- **Add check items to the whitelist**

If you want to disable alerts for a check item, click **Whitelist** to add the check item to the whitelist. Check items in the whitelist do not trigger alerts.

Note You can also select multiple check items and click **Whitelist** in the lower-left corner to add the check items to the whitelist at a time.

- **Remove check items from the whitelist**

If you want to enable alerts for a check item in the whitelist, you can click **Remove** to remove the check item from the whitelist. You can remove one or more check items from the whitelist at a time. After a check item is removed from the whitelist, the check item triggers alerts again.

- **Verify the fix of a failed check item**

After you fix a baseline risk, you can click **Verify** to check whether the risk has been fixed. After you click **Verify**, the status of the check item changes to **Verifying**.

If you do not manually perform the verification, Apsara Stack Security automatically verifies the fix based on the detection interval specified in the policy.

If the verification is passed, the **Status** of the check item changes to **Passed**.

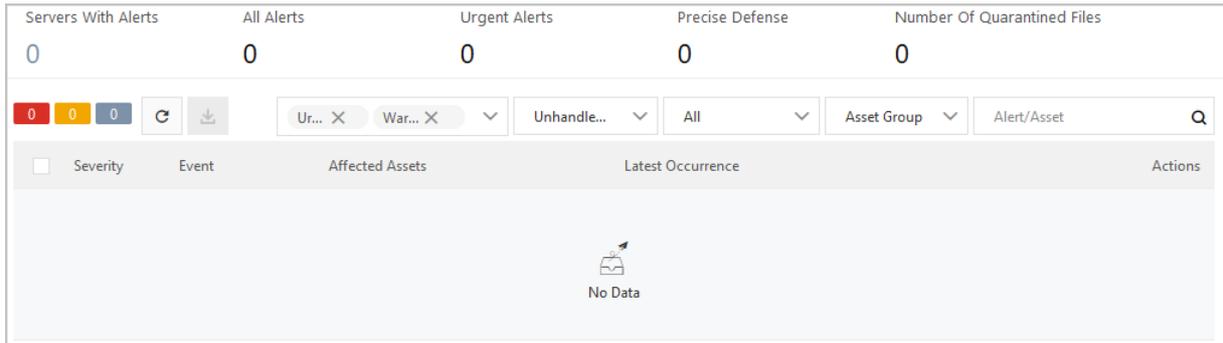
21.5.4. Intrusion prevention

21.5.4.1. Intrusion events

21.5.4.1.1. Intrusion event types

If the Server Security feature detects sensitive file tampering, webshells, unusual logons, suspicious processes, and malicious processes, it generates alerts. Based on these alerts, you can monitor the security status of your assets and handle potential threats at the earliest opportunity.

Apsara Stack Security provides statistics based on alerts. This helps you understand the overall protection situation. To view statistics on alerts, you must choose **Server Security > Intrusion Prevention > Intrusions**.



Alert types

The following table describes the alert items.

| Alert type | Description |
|---------------------|---|
| Threat intelligence | <p>Identify potential threats in assets based on the threat intelligence of Apsara Stack Security. Threat intelligence can correlate threat information to help you analyze and process the information. If threats are detected, alerts are reported. This helps improve the detection efficiency and response speed. Intelligence types include the following items:</p> <ul style="list-style-type: none"> • Malicious domain names • Malicious IP addresses • IP addresses of dark web services • IP addresses of command and control (C&C) servers • IP addresses of mining pools • Malicious URLs • Malicious download sources |
| Unusual Logon | <p>Detect unusual logons to your servers. You can specify approved logon IP addresses, time periods, and accounts. Logons from disapproved IP addresses, accounts, or time periods trigger alerts. You can manually add approved logon locations or configure the system to automatically update approved logon locations. You can also specify assets on which alerts are triggered when disapproved logon locations are detected.</p> <p>Apsara Stack Security can detect the following events:</p> <ul style="list-style-type: none"> • Logons to Elastic Compute Service (ECS) instances from disapproved IP addresses • Logons to ECS instances from disapproved locations • Execution of unusual commands after SSH-based logons to ECS instances • Brute-force attacks on SSH passwords of ECS instances |

| Alert type | Description |
|---|---|
| Webshell | <p>Use engines developed by Alibaba Cloud to scan common webshell files. Apsara Stack Security supports scheduled scan tasks, provides real-time protection, and quarantines webshell files.</p> <ul style="list-style-type: none"> • Apsara Stack Security scans the entire web directory early in the morning on a daily basis. A change made to files in the web directory triggers dynamic detection. • You can specify the assets on which Apsara Stack Security scans for webshells. • You can quarantine, restore, or ignore detected trojan files. |
| Precision defense | The antivirus feature provides precise protection against common ransomware, DDoS trojans, mining programs, trojans, malicious processes, webshells, and computer worms. |
| Suspicious Account | Detect logons to your assets from disapproved accounts. |
| Cloud threat detection | Detect threats in other cloud services. |
| Persistence | Detect suspicious scheduled tasks on servers and generates alerts when advanced persistent threats (APTs) on the servers are detected. |
| Unusual Network Connection | Detect disconnections or suspicious network connections. |
| Suspicious Process | Detect whether suspicious processes exist. |
| Malicious Process | <p>Detect your servers in real time. An agent is used to collect process information, and the information is uploaded to the cloud for detection. You can manage detected viruses in Apsara Stack Security Center.</p> <p>Apsara Stack Security can detect the following malicious activities and processes:</p> <ul style="list-style-type: none"> • Accesses from malicious IP addresses • Mining programs • Self-mutating trojans • Malicious programs • Trojans |
| Sensitive File Tampering | Check whether sensitive files on your server are maliciously modified. The sensitive files include preloaded configuration files in Linux shared libraries. |
| Other | Detect other types of attacks, such as DDoS attacks. |
| Web Application Threat Detection | Detect server intrusions that use web applications. |
| Application intrusion event | Detect server intrusions that use system application components. |

21.5.4.1.2. View and handle detected intrusion events

This topic describes how to view and handle detected intrusion events on the Intrusions page.

Background information

After intrusion events are detected, the intrusion events are displayed on the Intrusions page. You can choose **Server Security > Intrusion Prevention > Intrusions** to go to the Intrusions page.

If the intrusion events are not handled, they are displayed in the **Unhandled Alerts** list on the **Intrusions** page. After the intrusion events are handled, the status changes from **Unhandled Alerts** to **Handled**.

 **Note** Apsara Stack Security retains the records of **Unhandled Alerts** and **Handled** on the **Intrusions** page. By default, the records of **Unhandled Alerts** are displayed.

View intrusion events

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Intrusion Prevention > Intrusions**.
3. On the page that appears, search for or view intrusion events. You can also view the details about the events.

Handle intrusion events

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Intrusion Prevention > Intrusions**.
3. On the **Intrusions** page, find the intrusion event that you want to handle and click **Process** in the **Actions** column. In the dialog box that appears, set Process Method and click **Process Now**.

 **Note** If the intrusion event contains multiple correlated exceptions, on the page that appears after you click **Process**, you can handle the exceptions.

- **Ignore**: If you ignore the intrusion event, the status of the intrusion event changes to **Handled**. Server Guard no longer generates alerts for the event.
- **Add To Whitelist**: If the intrusion event is a false positive, you can add the intrusion event to the whitelist. Then, the status of the intrusion event changes to **Handled**. Server Guard no longer generates alerts for the event. In the **Handled** list, you can click **Cancel whitelist** to remove a specific intrusion event from the whitelist.

 **Note** A false positive indicates that Server Guard has generated a false alert on a normal process. A common false positive is a **suspicious process that sends TCP packets**. The false positive notifies you that suspicious scans on other devices are detected on your servers.

- **Batch unhandled**: This method allows you to handle multiple intrusion events at a time. Before you handle multiple intrusion events at a time, we recommend that you view the details of the intrusion events.
4. (Optional) If you confirm that one or more intrusion events are false positives or need to be ignored, go to the **Intrusions** page. Then, select the intrusion events and click **Ignore Once** or **Whitelist**.

Export intrusion events

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Intrusion Prevention > Intrusions**.
3. In the upper-left corner above the intrusion event list on the **Intrusions** page, click the  icon to export the list.
After the list is exported, the **Done** message appears in the upper-right corner.
4. In the **Done** notification of the **Alerts** page, click **Download**.
The alert list is downloaded to your computer.

21.5.4.1.3. View exceptions related to an alert

Server Guard supports automatic analysis of exceptions related to an alert. You can click an alert in the alert list to view and manage all exceptions that are related to the alert. You can also view the results of automatic attack tracing to analyze and handle the exceptions.

Context

- Security Center automatically associates alerts with exceptions in real time to detect potential threats.
- Exceptions related to an alert are listed in chronological order. This allows you to analyze and handle the exceptions to improve the emergency response mechanism of your system.
- An automatically correlated alert is identified by the  icon.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Intrusion Prevention > Intrusions**.
3. On the Intrusions page, click the **name of the alert** that you want to manage. The alert details panel appears.
4. In the alert details panel, view the details and related exceptions of the alert. Then, handle the exceptions.
 - View alert details
You can view the assets that are affected by the alert, the first and latest time when the alert was triggered, and the details about the related exceptions.
 - View affected assets
You can move the pointer over the name of an **affected asset** to view the details about the asset. The details include information about all the alerts, vulnerabilities, baseline risks, and asset fingerprints on the asset.
 - View and manage **related exceptions**
In the **Related Exceptions** section, you can view the details about all the exceptions that are related to the alert. You can also view suggestions on how to handle the exceptions.
 - Click **Note** to the right of an exception to add a note for the exception.
 - Click the  icon to the right of a note to delete the note.

21.5.4.1.4. Use the file quarantine feature

Sever Guard can quarantine malicious files. Quarantined files are listed in the quarantine box on the Intrusions page. The system automatically deletes a quarantined file 30 days after the file is quarantined. You can restore a quarantined file with a few clicks before the file is deleted. This topic describes how to view and restore quarantined files.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Intrusion Prevention > Intrusions**.
3. In the upper-right corner of the **Intrusions** page, click **Quarantine**.
In the **Quarantine** panel, you can perform the following operations:
 - View information about quarantined files. The information includes server IP addresses, paths that store the files, file status, and modification time.
 - Click **Restore** in the **Actions** column to remove a file from the quarantine box. The restored file appears in the alert list again.

21.5.4.1.5. Configure security alerts

This topic describes how to configure security alerts, which allows you to specify approved logon locations and custom web directories to scan.

Context

Server Guard supports advanced logon settings. You can configure more fine-grained logon detection rules. For example, you can specify approved logon IP addresses, logon time, and logon accounts to block unauthorized requests sent to your assets.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Intrusion Prevention > Intrusions**.
3. In the upper-right corner, click **Settings**.

Configure parameters on different tabs.

o Add an approved logon location

- a. In the **Login Location** section, click **Management** on the right.
- b. Select the logon location that you want to add and select the servers that allow logons from the added location.
- c. Click **OK**.

Server Guard allows you to **edit** or **delete** approved logon locations.

- Find a specific logon location and click **Edit** on the right to change the servers that allow logons from this location.
- Find a specific logon location and click **Delete** on the right to delete the logon location.

o Configure advanced logon settings

 **Note** When you configure advanced logon settings, you can specify the IP addresses, accounts, and time ranges that are allowed for logons to your assets. After the advanced logon settings are configured, Server Guard sends you alerts if your assets receive unauthorized logon requests. The procedure of configuring advanced logon settings is similar to that of configuring **Login Location**. You can **add**, **edit**, or **delete** advanced logon settings in a similar way.

- Turn on or turn off **Uncommon IP Alert** to the right of **Common Login IPs**. If you turn on **Uncommon IP Alert** and your assets receive logon requests from unauthorized IP addresses, alerts are triggered.
- Turn on or turn off **Uncommon Time Alert** to the right of **Common Login Time**. If you turn on **Uncommon Time Alert** and your assets receive logon requests in unauthorized time ranges, alerts are triggered.
- Turn on or turn off **Uncommon Account Alert** to the right of **Common Login Accounts**. If you turn on **Uncommon Account Alert** and your assets receive logon requests from unauthorized accounts, alerts are triggered.

o Add web directories to scan

Server Guard automatically scans web directories of data assets in your servers and runs dynamic and static scan tasks. You can also manually add other web directories.

- a. In the **Add Scan Targets** section, click **Management** on the right.
- b. Specify a valid web directory and select the servers on which the specified web directory is scanned.

 **Note** To ensure the scan performance and efficiency, we recommend that you do not specify a root directory.

- c. Click **OK**.

21.5.4.1.6. Cloud threat detection

The cloud threat detection feature provided by Server Guard is integrated with widely-used antivirus engines. The feature detects viruses based on large amounts of threat intelligence data provided by Alibaba Cloud and the exception detection model designed by Alibaba Cloud. This model is designed based on machine learning and deep learning. This way, the cloud threat detection feature can provide full-scale and dynamic antivirus protection to safeguard your servers.

The cloud threat detection feature scans hundreds of millions of files on a daily basis and protects millions of servers on the cloud.

Detection capabilities

The cloud threat detection feature uses the Server Guard agent to collect process information and scans the retrieved data for viruses in the cloud. If a malicious process is detected, you can stop the process and quarantine the source files.

The cloud threat detection feature provides the following capabilities:

- **Deep learning engine developed by Alibaba Cloud:** The deep learning engine is built on deep learning technology and a large number of attack samples. The engine detects malicious files on the cloud and automatically identifies potential threats to supplement traditional antivirus engines.
- **Cloud sandbox developed by Alibaba Cloud:** The cloud sandbox feature allows you to simulate cloud environments and monitor attacks launched by malicious samples. The cloud sandbox feature automatically detects threats and offers dynamic analysis and detection capabilities based on big data analytics and machine learning modeling techniques.
- **Integration with major antivirus engines:** The cloud threat detection feature is integrated with major antivirus engines and updates its virus library in real time.
- **Threat intelligence detection:** The cloud threat detection feature works with the exception detection module to detect malicious processes and operations based on threat intelligence data provided by Alibaba Cloud Security.

Detectable virus types

The cloud threat detection feature is developed based on the security technologies and expertise of Alibaba Cloud. The feature provides end-to-end security services, including threat intelligence collection, data masking, threat identification, threat analysis, and malicious file quarantine and restoration. You can quarantine and restore files that contain viruses in the Security Center console.

The cloud threat detection feature can detect the following types of viruses.

| Virus | Description |
|----------------|---|
| Mining program | A mining program consumes server resources and mines cryptocurrency without authorization. |
| Computer worm | A computer worm uses computer networks to replicate itself and spread to a large number of computers within a short period of time. |
| Ransomware | Ransomware, such as WannaCry, uses encryption algorithms to encrypt files and prevent users from accessing the files. |
| Trojan | A trojan is a program that allows an attacker to access information about servers and users, gain control of the servers, and consume system resources. |
| DDoS trojan | A DDoS trojan hijacks servers and uses zombie servers to launch DDoS attacks, which interrupts your service. |
| Backdoor | A backdoor is a malicious program injected by an attacker. Then, the attacker can use the backdoor to control the server or launch attacks. |

| Virus | Description |
|-------------------|--|
| Computer virus | A computer virus inserts malicious code into normal programs and replicates the code to infect the whole system. |
| Malicious program | A malicious program may pose threats to system and data security. |

Benefits

- **Self-developed and controllable:** The cloud threat detection feature is based on deep learning, machine learning, and big data analytics with a large number of attack and defense practices. The feature uses multiple detection engines to dynamically protect your assets against viruses.
- **Lightweight:** The cloud threat detection feature consumes only 1% of CPU resources and 50 MB of memory.
- **Dynamic:** The cloud threat detection feature dynamically retrieves startup logs of processes to monitor the startup of viruses.
- **Easy to manage:** You can manage all servers and view their status at any time in the Security Center console.

Threat detection limits

Apsara Stack Security Center allows you to detect and process security alerts, scan for and fix vulnerabilities, analyze attacks, and check security settings. Apsara Stack Security Center can analyze alerts and automatically trace attacks. This allows you to protect your assets. Apsara Stack Security supports a wide range of protection features. We recommend that you install the latest system patches on your assets. We also recommend that you use security services, such as Cloud Firewall and Web Application Firewall (WAF), to better protect your assets against attacks.

 **Note** Attacks and viruses are evolving, and security breaches may occur in various business environments. We recommend that you use the alerting, vulnerability detection, baseline check, and configuration assessment features provided by Apsara Stack Security to better protect your assets against attacks.

21.5.4.2. Website tamper-proofing

21.5.4.2.1. Overview

Tamper protection monitors website directories in real time, restores modified files or directories, and protects websites from trojans, hidden links, and uploads of violent and illicit content.

Background information

To make illegal profits or conduct business attacks, attackers exploit vulnerabilities in websites to insert illegal hidden links and tamper with the websites. Defaced web pages affect normal user access and may lead to serious economic losses, damaged brand reputation, or political risks.

Tamper protection allows you to add Linux and Windows processes to the whitelist and update protected files in real time.

How tamper protection works

The Security Center agent automatically collects the list of processes that attempt to modify files in the protected directories of the protected servers. It identifies unusual processes and file changes in real time and blocks unusual processes.

The alert list is displayed on the Tamper Protection page. You can view unusual file changes, the corresponding processes, and the number of attempts made by each process in the alert list. If a file is modified by a trusted process, you can add the process to the whitelist. After the process is added to the whitelist, tamper protection no longer blocks the process. In scenarios where the content of websites, such as news and education websites, is frequently modified, the whitelist saves you the effort of frequently enabling and disabling tamper protection.

Versions of operating systems and kernels supported by tamper protection

| OS | Supported operating system version | Supported kernel version |
|---------|--|--|
| Windows | Windows Server 2008 and later | All versions |
| CentOS | 6.5, 6.6, 6.7, 6.8, 6.9, 6.10, 7.0, 7.1, 7.2, 7.3, 7.4, 7.5, and 7.6 | <ul style="list-style-type: none"> • 2.6.32-x • 3.10.0-x |
| Ubuntu | 14, 16, and 18 | <ul style="list-style-type: none"> • 3.13.0-32-generic • 3.13.0-86-generic • 4.4.0-62-generic • 4.4.0-63-generic • 4.4.0-93-generic • 4.4.0-151-generic • 4.4.0-117-generic • 4.15.0-23-generic • 4.15.0-42-generic • 4.15.0-45-generic • 4.15.0-52-generic |

 **Note**

- The preceding table lists kernel versions supported by tamper protection. Servers that use an unsupported kernel version cannot use tamper protection. Make sure that your server uses a supported kernel version. If a kernel version is not supported, you must upgrade it to a supported version. Otherwise, you cannot add processes to the whitelist.
- Before you upgrade the server kernel, back up your asset data.

21.5.4.2.2. Configure tamper protection

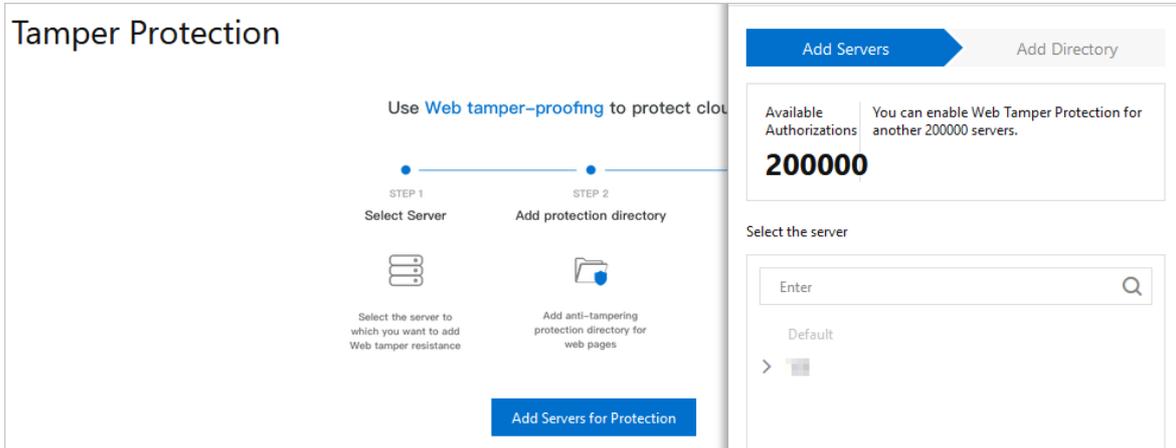
The Server Security feature allows you to configure tamper protection for web pages.

Limits

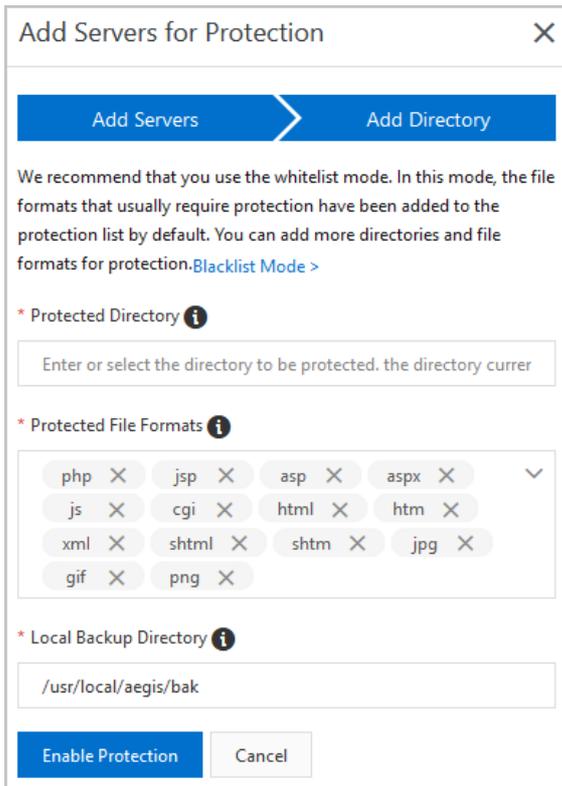
- For each server, you can add a maximum of 10 directories for protection.
- If you want to add directories that are on a Windows server, the directories must meet the following requirements: The size of each directory does not exceed 20 GB. Each directory contains no more than 2,000 folders. The number of directory levels does not exceed 20. The size of each file does not exceed 3 MB.
- If you want to add directories that are on a Linux server, the directories must meet the following requirements: The size of each directory does not exceed 20 GB. Each directory contains no more than 3,000 folders. The number of directory levels does not exceed 20. The size of each file does not exceed 3 MB.
- Before you add a directory for protection, make sure that the directory meets the preceding requirements.
- We recommend that you exclude file formats that do not require protection, such as *LOG*, *PNG*, *JPG*, *MP4*, *AVI*, and *MP3*. Multiple file formats can be separated by semicolons (;).

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Intrusion Prevention > File Tamper Protection**.
3. On the **Tamper Protection** page, click **Add Servers for Protection**.
4. In the **Add Servers for Protection** panel, select a server that you want to protect.



5. Click **Next** to go to the **Add Directory** step.
6. In the **Add Directory** step, configure the parameters.



Select a protection mode. The settings of other parameters vary based on the protection mode. You can select **Whitelist Mode** or **Blacklist Mode**. In whitelist mode, tamper protection is enabled for the specified directories and file formats. In blacklist mode, tamper protection is enabled for the subdirectories, file formats, and files that are not specified. By default, **Whitelist Mode** is selected.

- o The following table describes the parameters that you must configure if you select **Whitelist Mode**.

| Parameter | Description |
|------------------------|--|
| Protected Directory | Enter the path of the directory that you want to protect. <div style="background-color: #e0f2f7; padding: 5px; border: 1px solid #ccc;"> ? Note Servers that run Linux or Windows operating systems use different path formats. Enter a valid directory path based on the type of your operating system. </div> |
| Protected File Formats | Select file formats that you want to protect from the drop-down list, such as <i>js</i> , <i>html</i> , <i>xml</i> , and <i>jpg</i> . |
| Local Backup Directory | The default path in which backup files of the protected directories are stored. By default, Apsara Stack Security assigns <code>/usr/local/aegis/bak</code> to Linux servers and <code>C:\Program Files (x86)\Alibaba\Aegis\bak</code> to Windows servers. You can change the default path based on your business requirements. |

- o The following table describes the parameters that you must configure if you select Blacklist Mode.

| Parameter | Description |
|--------------------------|--|
| Protected Directory | Enter the path of the directory that you want to protect. |
| Excluded Sub-Directories | Enter the subdirectories that you do not want to protect. Click Add Sub-Directory to add more subdirectories. Apsara Stack Security does not provide tamper protection for files in the excluded subdirectories. |
| Excluded File Formats | Select file formats that you do not want to protect from the drop-down list. Valid values: log , txt , and ldb . Apsara Stack Security does not provide tamper protection for the files in the excluded formats. |
| Excluded Files | Enter the path of the files for which you do not want to protect. Click Add File to add more files. Apsara Stack Security does not provide tamper protection for the excluded files. |
| Local Backup Directory | The default path in which backup files of the protected directories are stored. By default, Apsara Stack Security assigns <code>/usr/local/aegis/bak</code> to Linux servers and <code>C:\Program Files (x86)\Alibaba\Aegis\bak</code> to Windows servers. You can change the default path based on your business requirements. |

7. Click **Enable Protection**.

After you enable this feature for a server, the server is displayed on the Management tab of the **Tamper Protection** page.

? **Note** By default, tamper protection is in the **Off** state for the server. To enable tamper protection for the server, you must **turn on** the switch in the Protection column on the **Tamper Protection** page.

- 8. On the **Tamper Protection** page, find the server that you add. Then, click the **Management** tab and turn on the switch in the **Protection** column to enable protection for the server.

 **Note** By default, tamper protection is in the **Off** state for the server. To enable tamper protection for the server, you must turn on the switch in the Protection column on the **Tamper Protection** page.

After tamper protection is enabled, the status of the server changes to **Running**.

 **Note** If the status of the server is **Exception**, move the pointer over **Exception** in the Status column to view the cause and click **Retry** to enable tamper protection again.

What to do next

After you enable tamper protection for a server, you can go to the **Alerts** page and select **Webpage Tampering** from the alert type drop-down list to view the alerts generated upon tampering events.

Note

Tamper protection does not take effect immediately after you configure the protected directory, and you can still write files to the directory. In this case, you must go to the **Management** page, disable **Protection** for the server where the directory is located, and then enable **Protection** again.

Handling suggestions for abnormal protection states

| State | Description | Suggestion |
|-----------------|---|---|
| Initializing | Tamper protection is being initialized. | If this is your first time enabling tamper protection for a server, the protection status becomes Initializing . Tamper protection will be enabled in a few seconds. |
| Running | Tamper protection is enabled and running as expected. | None. |
| Exception | An error occurred when tamper protection was enabled. | Move the pointer over Exception in the Status column to view the exception cause and click Retry . |
| Not Initialized | Web tamper protection is disabled. | Turn on the switch in the Protection column to enable tamper protection. |

21.5.4.2.3. View protection status

This topic describes how to view the status of tamper protection for your assets.

Context

The tamper protection feature monitors changes to the files in website directories in real time and blocks suspicious file changes. To view the status of and details about the tamper protection feature, you must log on to Apsara Stack Security Center and choose **Server Security > Intrusion Prevention > File Tamper Protection**. The following information is displayed:

- Tamper protection overview

You can view the numbers of files that are changed on the current day and in the last 15 days, the number of protected servers, and the number of protected directories.

- Distribution of protected file types

Protected file types include TXT, PNG, MSI, and ZIP. You can also add more types of files for tamper protection based on your business requirements.

 **Note** All types of files for tamper protection can be added.

- **Top five files**

This section shows the names and paths of the top five files that are ranked based on the number of changes to files in descending order in the last 15 days.

- **Tamper protection alerts**

This section lists the alerts generated for blocked suspicious changes to files for your assets. You can view details about the alerts, including the severity, alert name, affected assets, paths of files with suspicious changes, and protection status.

 **Note**

- If an alert is reported more than 100 times, we recommend that you handle the alert at your earliest opportunity.
- Only alerts at the **Medium** level are displayed in the console.
- Only alerts in the **Defended** state are displayed. These alerts are triggered when the tamper protection feature blocks suspicious processes that attempt to modify files without authorization.

21.5.4.3. Configure the anti-virus feature

Server Guard provides the anti-virus feature. This feature allows you to customize settings for virus and webshell detection.

Detect and remove viruses

The anti-virus feature can automatically quarantine common Internet viruses, such as common trojans, ransomware, mining programs, and DDoS trojans. Apsara Stack Security experts check and verify all the automatically quarantined viruses to ensure a minimum false positive rate.

If the virus blocking feature is disabled, Server Guard generates alerts when viruses are detected. You must manually manage detected viruses in Apsara Stack Security Center. We recommend that you enable the virus blocking feature to improve the security of your servers.

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Intrusion Prevention > Virus Defence**.
3. On the **Anti-virus** tab, click **Scan**.
4. In the dialog box that appears, select the servers that you want to scan.
5. Click **Scan**.
6. On the **Anti-virus** page, click the **Real-time protection** tab and enable the virus blocking feature.

After the virus blocking feature is enabled, Server Guard automatically quarantines common viruses that are detected. Quarantined viruses are listed on the Alerts page. To filter quarantined viruses, you can select the **Precision defense** type.

Detect and remove webshells

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Intrusion Prevention > Virus Defence**.
3. Specify servers for webshell detection.
 - i. In the **Webshell Detection** section, click **Manage**.

- ii. Select the servers for which you want to enable webshell detection.
- iii. Click OK to complete the configuration.

21.5.5. Log retrieval

21.5.5.1. Log retrieval overview

The log retrieval function provided by Server Security allows you to manage logs scattered in various systems of Apsara Stack in a centralized manner, so that you can easily identify the causes of issues that occur on your servers.

The log retrieval function supports storage of logs for 180 days and query of logs generated within 30 days.

Benefits

The log retrieval function provides the following benefits:

- **End-to-end log retrieval platform:** Allows you to retrieve logs of various Apsara Stack services in a centralized manner and trace issues easily.
- **Cloud-based SaaS service:** Allows you to query logs on all servers in Apsara Stack without additional installment and deployment.
- Supports TB-level data retrieval. It also allows you to add a maximum of 50 inference rules (Boolean expressions) in a search condition and obtain full-text search results within several seconds.
- Supports a wide range of log sources.
- Supports log shipping, which allows you to import security logs to Log Service for further analysis.

Scenarios

You can use log retrieval to meet the following requirements:

- **Security event analysis:** When a security event is detected on a server, you can retrieve the logs to identify the cause and assess the damage and affected assets.
- **Operation audit:** You can audit the operation logs on a server to identify high-risk operations and serious issues in a meticulous way.

Supported log types

Log types

| Log type | Description |
|-------------------------|---|
| Logon history | Log entries about successful system logons |
| Brute-force attack | Log entries about system logon failures that are generated during brute-force attacks |
| Process snapshot | Log entries about processes on a server at a specific time |
| Listening port snapshot | Log entries about listening ports on a server at a specific time |
| Account snapshot | Log entries about account logon information on a server at a specific time |
| Process initiation | Log entries about process initiation on a server |
| Network connection | Log entries about active connections from a server to external networks |

21.5.5.2. Query logs

This topic describes how to search for and view server logs.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Log Retrieval**.
3. Specify search conditions.

| Search condition | Description |
|------------------|--|
| Log source | The log source that you want to query. For more information, see Log sources . |
| Field | The field that is recorded for the log source. For more information, see Log sources . |
| Keyword | The keyword of the field. |
| Logical operator | The equality operator. |
| + | The inference rules in a search condition for a log source. |
| Add conditions | The search conditions for different log sources. |

4. Click **Search** and view the search result.
 - **Reset**: Click **Reset** to clear the search condition configurations.
 - **Save Search**: Click **Save Search** to save the search condition configurations which you can use to search for logs in the future.
 - **Saved Searches**: Click **Saved Searches** to select and use a search condition that you saved.

21.5.5.3. Supported log sources and fields

This topic describes the log sources and fields that are supported by the log retrieval feature.

The log retrieval feature allows you to query the following types of log sources. You can click a log source link to view the fields that can be retrieved.

| Log source | Description |
|--|--|
| Logon history | Log entries about successful system logons |
| Logs of brute-force attacks | Log entries about failed system logons during brute-force attacks |
| Process snapshot logs | Log entries about processes on a server at a specific point in time |
| Logs of listening port snapshots | Log entries about listening ports on a server at a specific point in time |
| Account snapshot logs | Log entries about account-based logons on a server at a specific point in time |
| Process startup logs | Log entries about process startups on a server |
| Network connection logs | Log entries about active connections from a server to the Internet. |

Logon history

The following table describes the fields that you can use to query the logon history.

| Field | Data type | Description |
|------------|-----------|---|
| uuid | string | The ID of the client. |
| IP | string | The IP address of the server. |
| warn_ip | string | The source IP address used for the logon. |
| warn_port | string | The logon port. |
| warn_user | string | The username used for the logon. |
| warn_type | string | The logon type. |
| warn_count | string | The number of logon attempts. |

Logs of brute-force attacks

The following table describes the fields that you can use to query logs of brute-force attacks.

| Field | Data type | Description |
|------------|-----------|--|
| uuid | string | The ID of the client. |
| IP | string | The IP address of the server. |
| warn_ip | string | The source IP address of the attack. |
| warn_port | string | The target port of the attack. |
| warn_user | string | The target username of the attack. |
| warn_type | string | The attack type. |
| warn_count | string | The number of brute-force attack attempts. |

Process startup logs

The following table describes the fields that you can use to query process startup logs.

| Field | Data type | Description |
|-----------|-----------|-------------------------------|
| uuid | string | The ID of the client. |
| IP | string | The IP address of the server. |
| pid | string | The ID of the process. |
| groupname | string | The user group. |
| ppid | string | The ID of the parent process. |
| uid | string | The ID of the user. |
| username | string | The username. |
| filename | string | The file name. |

| Field | Data type | Description |
|-----------|-----------|--------------------------------------|
| pfilename | string | The name of the parent process file. |
| cmdline | string | The command line. |
| filepath | string | The path of the process file. |
| pfilepath | string | The path of the parent process file. |

Logs of listening port snapshots

The following table describes the fields that you can use to query logs about listening port snapshots.

| Field | Data type | Description |
|-----------|-----------|-------------------------------|
| uuid | string | The ID of the client. |
| IP | string | The IP address of the server. |
| src_port | string | The listening port. |
| src_ip | string | The listening IP address. |
| proc_path | string | The path of the process file. |
| pid | string | The ID of the process. |
| proc_name | string | The name of the process. |
| proto | string | The protocol. |

Account snapshot logs

The following table describes the fields you can use to query account snapshot logs.

| Field | Data type | Description |
|-----------|-----------|---|
| uuid | string | The ID of the client. |
| IP | string | The IP address of the server. |
| perm | string | Indicates whether the user has root permissions. |
| home_dir | string | The home directory. |
| warn_time | string | The time when a password expiration notification is sent. |
| groups | string | The group to which the user belongs. |
| login_ip | string | The IP address of the last logon. |
| last_chg | string | The time when the password was last changed. |
| shell | string | The Linux shell command. |

| Field | Data type | Description |
|----------------|-----------|--|
| domain | string | The Windows domain. |
| tty | string | The logon terminal. |
| account_expire | string | The time when the account expires. |
| passwd_expire | string | The time when the password expires. |
| last_logon | string | The last logon time. |
| user | string | The username. |
| status | string | The account status. Valid values: <ul style="list-style-type: none"> • 0: disabled • 1: normal |

Process snapshot logs

The following table describes the fields that you can use to query process snapshot logs.

| Field | Data type | Description |
|------------|-----------|---|
| uuid | string | The ID of the client. |
| IP | string | The IP address of the server. |
| path | string | The path of the process file. |
| start_time | string | The time when the process was started. |
| uid | string | The ID of the user. |
| cmdline | string | The command line. |
| pname | string | The name of the parent process. |
| name | string | The name of the process. |
| pid | string | The ID of the process. |
| user | string | The username. |
| md5 | string | The MD5 hash value of the process file. If the size of the process file exceeds 1 MB, the system does not calculate the MD5 hash value of the process file. |

Network connection logs

The following table describes the fields that you can use to query network connection logs.

| Field | Data type | Description |
|-------|-----------|-----------------------|
| uuid | string | The ID of the client. |

| Field | Data type | Description |
|-----------|-----------|-------------------------------|
| IP | string | The IP address of the server. |
| src_ip | string | The source IP address. |
| src_port | string | The source port. |
| proc_path | string | The path of the process file. |
| dst_port | string | The destination port. |
| proc_name | string | The name of the process. |
| dst_ip | string | The destination IP address. |
| status | string | The status. |

21.5.5.4. Logical operators

The log retrieval feature supports multiple search conditions. You can add multiple logical operators to one search condition for one log source, or combine multiple search conditions for several log sources by using different logical operators. This topic describes the logical operators that are supported in log retrieval. Examples are provided to help you understand these operators.

The following table describes the logical operators that are supported in log retrieval.

Logical operators

| Logical operator | Description |
|------------------|---|
| and | <p>Binary operator.</p> <p>This operator is in the format of <code>query 1 and query 2</code>, which indicates the intersection of the query results of <code>query 1</code> and <code>query 2</code>.</p> <p>Note If no logical operators are used for multiple keywords, the default operator is AND.</p> |
| or | <p>Binary operator.</p> <p>This operator is in the format of <code>query 1 or query 2</code>, which indicates the union of the query results of <code>query 1</code> and <code>query 2</code>.</p> |
| not | <p>Binary operator.</p> <p>This operator is in the format of <code>query 1 not query 2</code>, which indicates the results that match <code>query 1</code> but do not match <code>query 2</code>. This format is equivalent to <code>query 1 - query 2</code>.</p> <p>Note If you use only <code>not query 1</code>, the log data that does not contain the query results of <code>query 1</code> is returned.</p> |

21.5.6. Settings

21.5.6.1. Install the Server Guard agent

This topic describes how to install the Server Guard agent.

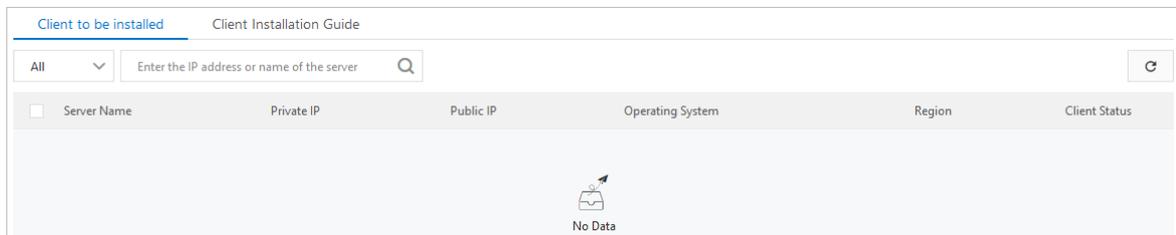
Context

To use the protection services provided by Server Guard, you must install the Server Guard agent on the operating system of your server.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Server Security > Server Settings > Client Installation**.
3. (Optional) On the Client Installation page, click the **Client to be installed** tab to view the number of servers on which the Server Guard agent is not installed. On this tab, you can also view the information about the servers.

You can specify the operating system type, IP address, or server name to search for a server.



4. Click the **Client Installation Guide** tab.
5. Obtain and install the Server Guard agent based on the operating system type of your server.
 - o **Windows**
 - a. In the left-side pane of the page, click **Click to download** to download the installation package to your computer.
 - b. Upload the installation package to your server. For example, you can use an FTP client to upload the installation package to your server.
 - c. Run the installation package on your server as an administrator.

Note If you install the agent on a server that is not in Alibaba Cloud, you are prompted to enter the installation verification key. You can find the installation verification key on the Client Installation Guide tab.

- o **Linux**
 - a. In the right-side pane of the page, select **Alibaba Cloud Server** or **Non-Alibaba Server**.
 - b. Select the installation command for your 32-bit or 64-bit operating system and click **Copy** to copy the command.
 - c. Log on to your Linux server as an administrator.
 - d. Run the installation command on your Linux server to download and install the Server Guard agent.

21.5.6.2. Manage protection modes

This topic describes how to manage protection modes for each server to improve the efficiency and security of the server.

Procedure

1. Log on to [Apsara Stack Security Center](#).

2. In the left-side navigation pane, choose **Server Security > Server Settings > Protection Mode**.
3. On the Protection Mode page, click **Manage**.
Configure protection modes for servers.
 - **Business First Mode**: The peak CPU utilization is less than 10%, and the peak memory usage is less than 50 MB.
 - **Protection First Mode**: The peak CPU utilization is less than 20%, and the peak memory usage is less than 80 MB.
4. Click **OK**.

21.6. Physical server security

21.6.1. Create and grant permissions to a security administrator account

The physical server security feature is used to ensure the security of physical servers on the platform side. This feature requires you to use a dedicated security administrator account for the platform. This topic describes how to create and grant permissions to a security administrator account.

Procedure

1. Log on to the Apsara Uni-manager Management Console as a system administrator.
For more information, see the "**Log on to the Apsara Uni-manager Management Console**" topic of *Apsara Uni-manager Management Console User Guide*.
2. Create a dedicated organization that is used to manage the security of physical servers, and obtain the primary key of the organization.

 **Notice** Make sure that the organization is used only to manage the security of physical servers. Do not add Elastic Compute Service (ECS) instances to the organization.

- i. Create the dedicated organization.
For more information, see **Enterprise Center > Organization Management > Create Organization** in *Apsara Uni-manager Management Console User Guide*.
 - ii. Obtain the **primary key** of the newly created organization.
For more information, see **Enterprise Center Organization Management Obtain the AccessKey pair of an organization** in *Apsara Uni-manager Management Console User Guide*.
3. Create a dedicated account to manage the security of physical servers.
For more information, see **Enterprise Center > User Management > System User Management > Create User** in *Apsara Uni-manager Management Console User Guide*.

 **Note** When you create the account, take note of the following points for the organization and role:

- In the **Organization** section, select the organization that is created in the previous step.
 - In the **Role** section, select **Platform Security Configuration Administrator** and **Security System Configuration Administrator**.
4. Log on to Apsara Stack Security Center by using the newly created account.
For more information, see [Log on to Apsara Stack Security Center](#).
 5. Add the **primary key** of the newly created organization to the protection configuration of physical servers.

- i. In the left-side navigation pane, choose **Security Management Center (SOC) > System Configuration > Global Settings**.
- ii. On the **Global Settings** page, click the **Physical Server Protection** tab.
- iii. Click **Add Account**.
- iv. In the **Add Account** dialog box, specify the **Username** and **Primary Key** parameters.
 - **Username**: Enter the account that you created in Step 3.
 - **Primary Key**: Enter the primary key that you obtained in Step 2.
- v. Click **OK**.

Result

After you complete the configuration, the **Physical Server Security** feature appears in the left-side navigation pane of Apsara Stack Security. Then, you can use the account created in this topic to maintain the security of physical servers.

21.6.2. View the information on the Overview page

Security administrators can view the security status of all physical servers on the Overview page of Apsara Stack Security.

Procedure

1. Log on to [Apsara Stack Security Center](#) by using an Apsara Stack tenant account.

 **Note** For more information about the Apsara Stack tenant account, see [Create and grant permissions to a security administrator account](#).

2. In the left-side navigation pane, choose **Physical Server Security > Overview**.
3. On the **Overview** page, view the security status of your physical servers.
 - **Security event statistics**: displays the numbers of security events on the physical servers. Security events include unusual logons, webshells, and server exceptions.
 - **Events**: displays the trends of security events on the physical servers. A security event is an intrusion event that is detected on a physical server.
 - **Protection Status**: displays the number of physical servers that are protected and the number of offline physical servers.
 - **Recent Important Events**: displays the recent important security events that are detected on the physical servers. You can click a security event to view its details.

21.6.3. Physical servers

21.6.3.1. Manage physical server groups

This topic describes how to manage physical server groups. To facilitate the security management of physical servers, you can add the physical servers to groups and view their security events by group.

Context

By default, physical servers do not belong to a server group. You must add your physical servers to a server group. If you delete a group, all the physical servers in the group are retained but no longer belong to a server group.

Procedure

1. Log on to [Apsara Stack Security Center](#) by using an Apsara Stack tenant account.

 **Note** For more information about the Apsara Stack tenant account, see [Create and grant permissions to a security administrator account](#).

2. In the left-side navigation pane, choose **Physical Server Security > Servers**.

3. In the left-side group pane, manage sever groups.

- o Create a group.

Click the Add Subgroup icon next to **All Servers** or a specific group, enter a group name, and click **OK**.

 **Note** The system supports a maximum of three levels of groups.

- o Modify a group.

Click the Modify Group Name icon next to the target group, enter a new name, and click **OK**.

- o Delete a group.

Click the Delete icon next to the target group. In the message that appears, click **OK**.

 **Note** After you delete a group, all servers in the group are automatically moved to the default group.

- o Sort groups.

Click **Manage Groups** to sort groups in descending order by priority.

4. Change the server group of specific physical servers.

- Select servers from the list on the right.
- Click **Change Group**.
- In the Change Group dialog box that appears, select a group from the drop-down list.
- Click **OK**.

21.6.3.2. Manage physical servers

This topic describes how to manage servers. On the Servers page, you can view the status of servers protected by Server Guard.

Procedure

1. Log on to [Apsara Stack Security Center](#) by using an Apsara Stack tenant account.

 **Note** For more information about the Apsara Stack tenant account, see [Create and grant permissions to a security administrator account](#).

2. In the left-side navigation pane, choose **Physical Server Security > Servers**.

3. (Optional) Search for a server.

To view the agent status of a server, enter the server IP address in the search bar, and click **Search**. Detailed server information, such as security information, is displayed.

4. View the agent status and detailed security information of the server.

Click



in the upper-right corner of the page to select the information columns you want to display. The following table lists the information categories.

| Category | Information |
|---------------------|--|
| Basic information | <ul style="list-style-type: none"> ◦ Server IP/Name ◦ Tag ◦ OS ◦ Region |
| Agent status | Agent Status |
| Threat prevention | <ul style="list-style-type: none"> ◦ Vulnerability ◦ Baseline Risk |
| Intrusion detection | <ul style="list-style-type: none"> ◦ Unusual Logons ◦ Webshells ◦ Suspicious Servers |
| Server fingerprints | <ul style="list-style-type: none"> ◦ Processes ◦ Ports ◦ Root Accounts/Total Accounts |

5. Manage servers.

| Action | Description |
|-------------------------|---|
| Change Group | Select servers and click Change Group to add the selected servers to a new group. |
| Modify Tag | Select servers and click Modify Tag to modify tags for the servers. |
| Security Inspection | Select servers and click Security Inspection to select the items to be checked. |
| Delete External Servers | Select external servers, and choose More > Delete External Servers . |
| Disable Protection | Select the servers whose agent status is Online , and choose More > Disable Protection . This temporarily disables protection for these servers to reduce server resource consumption. |
| Enable Protection | Select the servers whose agent status is Disable Protection , and choose More > Enable Protection . This enables protection for these servers. |

21.6.4. Intrusion detection

21.6.4.1. Configure policies to identify unusual logons

This topic describes how to configure logon security. You can set approved locations, IP addresses, time periods, and accounts for logons.

Procedure

1. Log on to [Apsara Stack Security Center](#) by using an Apsara Stack tenant account.

 **Note** For more information about the Apsara Stack tenant account, see [Create and grant permissions to a security administrator account](#).

2. In the left-side navigation pane, choose **Physical Server Security > Intrusion Detection**.
3. Click the **Unusual Logons** tab.
4. Click **Logon Security** in the upper-right corner of the page.
5. Set approved logon locations.

To add an approved logon location, follow these steps:

- i. Click **Add**.
- ii. Select a logon location from the drop-down list.
- iii. Specify the servers on which the selected logon location takes effect.
 - Click **All Servers** to select specific servers.
 - Click **Server Groups** to select servers by group.
- iv. Click **OK**.

 **Note** Click **Modify** or **Delete** to modify or delete an approved logon location.

6. Set approved logon IP addresses.

Turn on the **Disapproved IP Alert** switch. The switch is turned on if it turns green.

To add an approved logon IP address, follow these steps:

- i. Click **Add**.
- ii. In the **Specify an Approved Logon IP** section, enter an IP address.
- iii. Specify the servers on which the specified IP address takes effect.
 - Click **All Servers** to select specific servers.
 - Click **Server Groups** to select servers by group.
- iv. Click **OK**.

Click **Modify** or **Delete** to modify or delete an approved logon IP address.

7. Set approved logon time periods.

Turn on the **Disapproved Time Alert** switch. The switch is turned on if it turns green.

To add an approved logon time period, follow these steps:

- i. Click **Add**.
- ii. In the **Specify an Approved Logon Duration** section, specify a time period.
- iii. Specify the servers on which the specified logon time period takes effect.
 - Click **All Servers** to select specific servers.
 - Click **Server Groups** to select servers by group.
- iv. Click **OK**.

Click **Modify** or **Delete** to modify or delete an approved logon time period.

8. Set approved accounts.

Turn on the **Disapproved Account Alert** switch. The switch is turned on if it turns green.

To add an approved account, follow these steps:

- i. Click **Add**.
- ii. In the **Specify an Approved Account** section, enter an account.

- iii. Select the servers on which the specified account takes effect.
 - Click **All Servers** to select specific servers.
 - Click **Server Groups** to select servers by group.
 - iv. Click **OK**.
- Click **Modify** or **Delete** to modify or delete an approved account.

What's next

After you configure policies to identify unusual logons, you can view and handle unusual logons on the **Unusual Logons** tab. For more information, see [Handle unusual logons](#).

21.6.4.2. Handle unusual logons

This topic describes how to handle unusual logons. The unusual logons include logons from disapproved locations or IP addresses, logons by using brute-force attacks or disapproved accounts, and logons at a disapproved time.

Context

For more information about the policies that are configured to identify unusual logons, see [Configure policies to identify unusual logons](#).

Procedure

1. Log on to [Apsara Stack Security Center](#) by using an Apsara Stack tenant account.

 **Note** For more information about the Apsara Stack tenant account, see [Create and grant permissions to a security administrator account](#).

2. In the left-side navigation pane, choose **Physical Server Security > Intrusion Detection**.
3. Click the **Unusual Logons** tab.
4. (Optional) Specify the filter conditions to locate unusual logons. The filter conditions include **Asset**, **Alert Type**, and **Status**.

 **Note** If you skip this step, all unusual logons appear in the event list.

5. In the event list, view and handle unusual logons.
 - Check whether an unusual logon is a false positive.
 - If the unusual logon is a false positive, click **Ignore**.
 - If the unusual logon is not a false positive, harden the security of the physical server and click **Ignore**. To harden the security of the physical server, you can specify a more complex password, fix vulnerabilities, or check configuration items.

21.6.5. Server fingerprints

21.6.5.1. Configure data refresh frequencies

You can configure the frequencies at which the data of running processes, system accounts, listening ports, and software versions is collected and refreshed.

Procedure

1. Log on to [Apsara Stack Security Center](#) by using an Apsara Stack tenant account.

 **Note** For more information about the Apsara Stack tenant account, see [Create and grant permissions to a security administrator account](#).

2. In the left-side navigation pane, choose **Physical Server Security > Server Fingerprints**.
3. In the upper-right corner of the Server Fingerprints page, click **Configure**.
4. In the **Configure** dialog box, configure the data refresh frequencies for the listening ports, running processes, system accounts, and software versions.
5. Click **OK**.

21.6.5.2. View listening ports

Apsara Stack Security regularly collects information about listening ports on a server.

Procedure

1. Log on to [Apsara Stack Security Center](#) by using an Apsara Stack tenant account.

 **Note** For more information about the Apsara Stack tenant account, see [Create and grant permissions to a security administrator account](#).

2. In the left-side navigation pane, choose **Physical Server Security > Server Fingerprints**.
3. Click the **Listening Port** tab.
4. (Optional)Specify **Port Number** and **Process Name**. Then, click **Search** to search for the specified port or process.

 **Note** If you skip this step, all listening ports appear in the port list.

5. In the port list, view **Port Number**, **Network Protocol**, and **Servers**.
6. Click a port number to view the information about the assets and processes that are associated with the port.

21.6.5.3. View running processes

Apsara Stack Security regularly collects the process information of a server.

Procedure

1. Log on to [Apsara Stack Security Center](#) by using an Apsara Stack tenant account.

 **Note** For more information about the Apsara Stack tenant account, see [Create and grant permissions to a security administrator account](#).

2. In the left-side navigation pane, choose **Physical Server Security > Server Fingerprints**.
3. Click the **Running Process** tab.
4. (Optional)Specify **Process Name** and **User** and click **Search** to search for the specified process or user.

 **Note** If you skip this step, all processes appear in the process list.

5. In the process list, view **Process Name** and **Servers**.
6. Click the name of a running process to view the information about the assets, paths, and running users that are associated with the process.

21.6.5.4. View account information

Apsara Stack Security regularly collects the account information of a server.

Procedure

1. Log on to [Apsara Stack Security Center](#) by using an Apsara Stack tenant account.

 **Note** For more information about the Apsara Stack tenant account, see [Create and grant permissions to a security administrator account](#).

2. In the left-side navigation pane, choose **Physical Server Security > Server Fingerprints**.
3. Click the **Account Information** tab.
4. (Optional)Specify **Username** and click **Search** to search for the specific account.

 **Note** If you skip this step, all accounts appear in the account list.

5. In the account list, view **Account** and **Servers**.
6. Click the name of an account to view the information about the assets and user groups that are associated with the account. You can also check whether the account has root permissions

21.6.5.5. View software versions

This topic describes how to view the information about the software versions of physical servers. The information facilitates the management of software assets.

Context

This topic covers the following scenarios:

- Check for software assets that are installed without authorization.
- Check for outdated versions of software assets.
- Find affected assets if vulnerabilities are detected.

Procedure

1. Log on to [Apsara Stack Security Center](#) by using an Apsara Stack tenant account.

 **Note** For more information about the Apsara Stack tenant account, see [Create and grant permissions to a security administrator account](#).

2. In the left-side navigation pane, choose **Physical Server Security > Server Fingerprints**.
3. Click the **All Versions** tab.
4. (Optional)Specify **Software Name**, **Version Name**, and **Software Installation Path**. Then, click **Search** to search for specific software.

 **Note** If you skip this step, all software appears in the software list.

5. In the software list, view the information in the **Software Name** and **Hosts** columns.
6. Click the name of the software to view the host, version, and installation directory of the software.

21.6.6. Log retrieval

21.6.6.1. Supported log sources and fields

This topic describes the log sources and fields that are supported by the log retrieval feature.

The log retrieval feature allows you to query the following types of log sources. You can click a log source link to view the fields that can be retrieved.

| Log source | Description |
|--|--|
| Logon history | Log entries about successful system logons |
| Logs of brute-force attacks | Log entries about failed system logons during brute-force attacks |
| Process snapshot logs | Log entries about processes on a server at a specific point in time |
| Logs of listening port snapshots | Log entries about listening ports on a server at a specific point in time |
| Account snapshot logs | Log entries about account-based logons on a server at a specific point in time |
| Process startup logs | Log entries about process startups on a server |
| Network connection logs | Log entries about active connections from a server to the Internet. |

Logon history

The following table describes the fields that you can use to query the logon history.

| Field | Data type | Description |
|------------|-----------|---|
| uuid | string | The ID of the client. |
| IP | string | The IP address of the server. |
| warn_ip | string | The source IP address used for the logon. |
| warn_port | string | The logon port. |
| warn_user | string | The username used for the logon. |
| warn_type | string | The logon type. |
| warn_count | string | The number of logon attempts. |

Logs of brute-force attacks

The following table describes the fields that you can use to query logs of brute-force attacks.

| Field | Data type | Description |
|-----------|-----------|--------------------------------------|
| uuid | string | The ID of the client. |
| IP | string | The IP address of the server. |
| warn_ip | string | The source IP address of the attack. |
| warn_port | string | The target port of the attack. |

| Field | Data type | Description |
|------------|-----------|--|
| warn_user | string | The target username of the attack. |
| warn_type | string | The attack type. |
| warn_count | string | The number of brute-force attack attempts. |

Process startup logs

The following table describes the fields that you can use to query process startup logs.

| Field | Data type | Description |
|-----------|-----------|--------------------------------------|
| uuid | string | The ID of the client. |
| IP | string | The IP address of the server. |
| pid | string | The ID of the process. |
| groupname | string | The user group. |
| ppid | string | The ID of the parent process. |
| uid | string | The ID of the user. |
| username | string | The username. |
| filename | string | The file name. |
| pfilename | string | The name of the parent process file. |
| cmdline | string | The command line. |
| filepath | string | The path of the process file. |
| pfilepath | string | The path of the parent process file. |

Logs of listening port snapshots

The following table describes the fields that you can use to query logs about listening port snapshots.

| Field | Data type | Description |
|-----------|-----------|-------------------------------|
| uuid | string | The ID of the client. |
| IP | string | The IP address of the server. |
| src_port | string | The listening port. |
| src_ip | string | The listening IP address. |
| proc_path | string | The path of the process file. |
| pid | string | The ID of the process. |
| proc_name | string | The name of the process. |

| Field | Data type | Description |
|-------|-----------|---------------|
| proto | string | The protocol. |

Account snapshot logs

The following table describes the fields you can use to query account snapshot logs.

| Field | Data type | Description |
|----------------|-----------|---|
| uuid | string | The ID of the client. |
| IP | string | The IP address of the server. |
| perm | string | Indicates whether the user has root permissions. |
| home_dir | string | The home directory. |
| warn_time | string | The time when a password expiration notification is sent. |
| groups | string | The group to which the user belongs. |
| login_ip | string | The IP address of the last logon. |
| last_chg | string | The time when the password was last changed. |
| shell | string | The Linux shell command. |
| domain | string | The Windows domain. |
| tty | string | The logon terminal. |
| account_expire | string | The time when the account expires. |
| passwd_expire | string | The time when the password expires. |
| last_logon | string | The last logon time. |
| user | string | The username. |
| status | string | The account status. Valid values: <ul style="list-style-type: none">• 0: disabled• 1: normal |

Process snapshot logs

The following table describes the fields that you can use to query process snapshot logs.

| Field | Data type | Description |
|-------|-----------|-------------------------------|
| uuid | string | The ID of the client. |
| IP | string | The IP address of the server. |
| path | string | The path of the process file. |

| Field | Data type | Description |
|------------|-----------|---|
| start_time | string | The time when the process was started. |
| uid | string | The ID of the user. |
| cmdline | string | The command line. |
| pname | string | The name of the parent process. |
| name | string | The name of the process. |
| pid | string | The ID of the process. |
| user | string | The username. |
| md5 | string | The MD5 hash value of the process file. If the size of the process file exceeds 1 MB, the system does not calculate the MD5 hash value of the process file. |

Network connection logs

The following table describes the fields that you can use to query network connection logs.

| Field | Data type | Description |
|-----------|-----------|-------------------------------|
| uuid | string | The ID of the client. |
| IP | string | The IP address of the server. |
| src_ip | string | The source IP address. |
| src_port | string | The source port. |
| proc_path | string | The path of the process file. |
| dst_port | string | The destination port. |
| proc_name | string | The name of the process. |
| dst_ip | string | The destination IP address. |
| status | string | The status. |

21.6.6.2. Logical operators

The log retrieval feature supports multiple search conditions. You can add multiple logical operators to one search condition for one log source, or combine multiple search conditions for several log sources by using different logical operators. This topic describes the logical operators that are supported in log retrieval. Examples are provided to help you understand these operators.

The following table describes the logical operators that are supported in log retrieval.

Logical operators

| Logical operator | Description |
|------------------|---|
| and | <p>Binary operator.</p> <p>This operator is in the format of <code>query 1 and query 2</code>, which indicates the intersection of the query results of <code>query 1</code> and <code>query 2</code>.</p> <p>Note If no logical operators are used for multiple keywords, the default operator is AND.</p> |
| or | <p>Binary operator.</p> <p>This operator is in the format of <code>query 1 or query 2</code>, which indicates the union of the query results of <code>query 1</code> and <code>query 2</code>.</p> |
| not | <p>Binary operator.</p> <p>This operator is in the format of <code>query 1 not query 2</code>, which indicates the results that match <code>query 1</code> but do not match <code>query 2</code>. This format is equivalent to <code>query 1 - query 2</code>.</p> <p>Note If you use only <code>not query 1</code>, the log data that does not contain the query results of <code>query 1</code> is returned.</p> |

21.6.6.3. Query logs

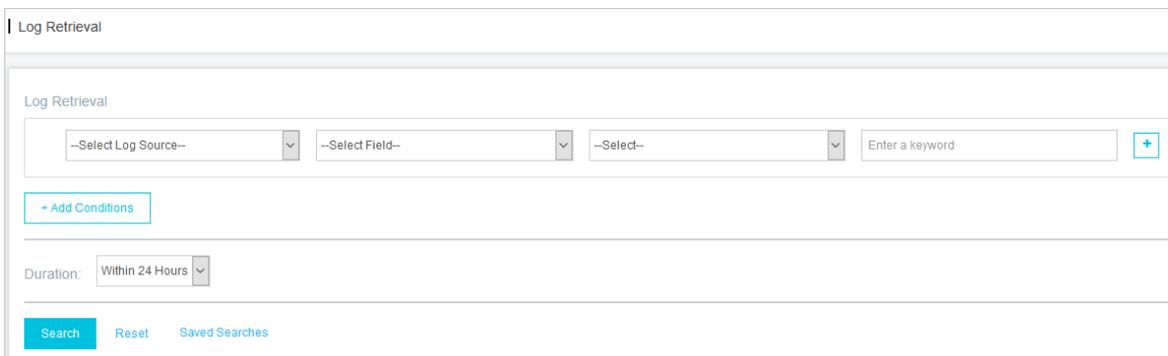
This topic describes how to search for and view physical server logs.

Procedure

1. Log on to [Apsara Stack Security Center](#) by using an Apsara Stack tenant account.

Note For more information about the Apsara Stack tenant account, see [Create and grant permissions to a security administrator account](#).

2. In the left-side navigation pane, choose **Physical Server Security > Log Retrieval**.
3. Specify search conditions.



Note For more information about log sources, log fields, and logical operators, see [Supported log sources and fields](#) and [Inference rules and logical operators](#).

4. Click **Search** and view the search result.

- **Reset**: Click **Reset** to clear the search condition configurations.
- **Save Search**: Click **Save Search** to save the search condition configurations which you can use to search for logs in the future.
- **Saved Searches**: Click **Saved Searches** to select and use a search condition that you saved.

21.6.7. Configure security settings for physical servers

This topic describes how to configure security settings for physical servers. You can enable or disable periodic trojan scans. You can also specify the working mode of the Server Guard agent.

Procedure

1. Log on to [Apsara Stack Security Center](#) by using an Apsara Stack tenant account.

 **Note** For more information about the Apsara Stack tenant account, see [Create and grant permissions to a security administrator account](#).

2. In the left-side navigation pane, choose **Physical Server Security > Settings**.
3. Enable periodic trojan scans for physical servers.
 - i. In the Trojan Scan section, click **Manage**.
 - ii. In the All Servers section, select the physical servers on which you want to perform periodic trojan scans. Then, click the rightwards arrow.
 - iii. Click **OK**.
4. On the Protection Mode page, click **Manage**.
Configure protection modes for servers.
 - **Business First Mode**: The peak CPU utilization is less than 10%, and the peak memory usage is less than 50 MB.
 - **Protection First Mode**: The peak CPU utilization is less than 20%, and the peak memory usage is less than 80 MB.

21.7. Application security

21.7.1. Quick start

This topic helps you get started with the features of Web Application Firewall (WAF).

WAF uses intelligent semantic analysis algorithms to identify web attacks. WAF also uses a learning model to enhance its analysis capabilities and meet your daily security protection requirements without relying on traditional rule libraries.

The following content describes the procedure for using WAF:

1. Customize WAF protection rules.

WAF provides default protection policies. You can also customize policies that suit your business requirements.

 - For more information about how to configure protection policies, see [Configure protection policies](#).
 - For more information about how to configure custom rules, see [Create a custom rule](#).
 - For more information about how to configure HTTP flood protection rules, see [Configure an HTTP flood protection rule](#).
2. Add websites that you want to protect.

WAF can protect Internet websites and virtual private cloud (VPC) websites.

- For more information about how to add an Internet website to WAF for protection, see [Add an Internet website for protection](#).
- For more information about how to add a VPC website to WAF for protection, see [Add a VPC website for protection](#).

3. Configure Domain Name System (DNS) resolution.

For more information about how to change the DNS-resolved source IP address for a website to a virtual IP address assigned by WAF, see [Modify DNS resolution settings](#).

4. View WAF protection results.

- For more information about how to view the protection overview, see [View protection overview](#).
- For more information about how to view the service access information, see [View Web service access information](#).
- For more information about how to view the detection logs for web attacks, see [View attack detection logs](#).
- For more information about how to view the detection logs for HTTP flood attacks, see [View HTTP flood protection logs](#).

21.7.2. Detection overview

21.7.2.1. View protection overview

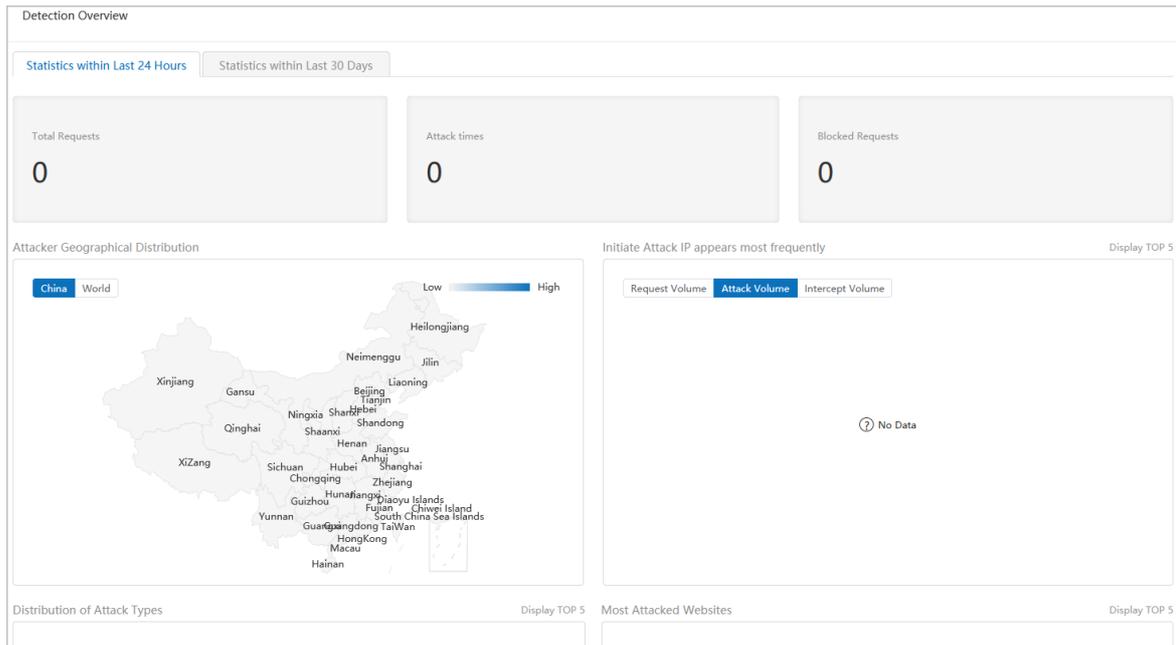
This topic describes how to view the Web Application Firewall (WAF) protection overview.

Context

The Detection Overview page displays information such as the statistics of previous attacks, the geographical distribution of attackers, the total number of requests, and the number of blocked requests. You can also view details about the attacks. This way, you can customize rules to protect your web services.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Application Security > WAF**.
3. Choose **Statistic > Detection Overview**.
4. On the **Detection Overview** page, view data on the **Statistics within Last 24 Hours** and **Statistics within Last 30 Days** tabs.



- **Total Requests**
Displays the total number of requests.
- **Attack times**
Displays the total number of attacks.
- **Blocked Requests**
Displays the number of blocked requests.
- **Attacker Geographical Distribution**
Displays the distribution of attackers on a map. You can select a map of China or a map of the world.
- **Initiate Attack IP appears most frequently (Display TOP 5)**
Displays the top five IP addresses from which the most attacks are launched in a bar chart. The x-axis indicates the numbers of requests. The y-axis indicates the IP addresses.
- **Distribution of Attack Types (Display TOP 5)**
Displays the distribution of the top five attack types and the number of attacks of each type in a bar chart.
- **Most Attacked Websites (Display TOP 5)**
Displays the top five attacked websites and the number of attacks on each website in a bar chart.

21.7.2.2. View access information

This topic describes how to view access information about web services.

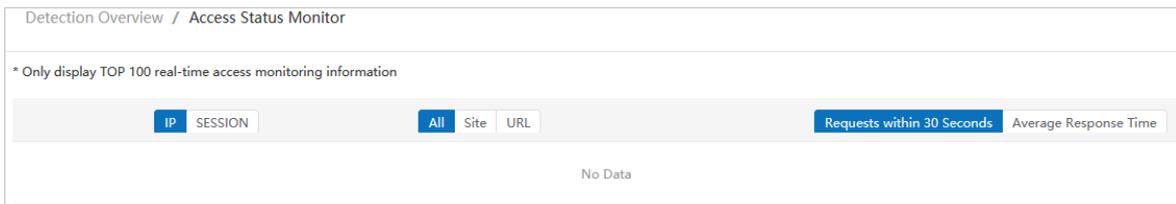
Context

Web Application Firewall (WAF) monitors the access of web services. This way, security administrators can analyze the service access information to detect vulnerabilities and improve security of the services.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Application Security > WAF**.
3. In the left-side navigation tree of the **WAF** page, choose **Statistic > Access Status Monitor**.

4. Filter access records to view details.



21.7.3. Protection logs

21.7.3.1. View attack detection logs

This topic describes how to view attack detection logs.

Context

These logs allow you to analyze attacks on your web services. You can update the protection policies and custom rules, and fix the web service vulnerabilities based on the analysis results.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Application Security > WAF**.
3. Choose **Detection Logs > Attack Detection Logs**.
4. Click **Filter**, specify filter conditions, and then click **OK**.

Note If you specify multiple conditions, they are evaluated by using a logical AND. The system returns the required logs only when all the conditions are met.

5. View the attack detection logs.

21.7.3.2. View HTTP flood protection logs

This topic describes how to view HTTP flood protection logs.

Context

These logs allow you to analyze HTTP flood attacks on your web services. In addition, you can update the HTTP flood protection rules and HTTP flood whitelist, and fix the web service vulnerabilities based on the analysis results.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Application Security > WAF**.
3. Choose **Detection Logs > HTTP Flood Detection Logs**.
4. Click **Filter**, specify filter conditions, and then click **OK**.

Note If you specify multiple conditions, they are evaluated by using a logical AND. The system returns the required logs only when all the conditions are met.

5. View the HTTP flood detection logs.

The blocked HTTP flood attacks, related rules, and attack time are displayed.

21.7.3.3. View system operation logs

This topic describes how to view system operation logs.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Application Security > WAF.**
3. In the left-side navigation tree, choose **Detection Logs > System operation log.**
4. View the system operation logs.

The usernames, content, IP addresses, and creation time are displayed.

21.7.3.4. View access logs

This topic describes how to view access logs.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Application Security > WAF.**
3. Choose **Detection Logs > Access Log.**
4. Click **Filter**, specify filter conditions, and then click **OK.**

 **Note** If you specify multiple conditions, they are evaluated by using a logical AND. The system returns the required logs only when all the conditions are met.

5. View the access logs.

The requested addresses, destination IP addresses, source IP addresses, methods, response status codes, and time are displayed.

21.7.4. Protection configuration

21.7.4.1. Configure protection policies

This topic describes how to configure Web Application Firewall (WAF) protection policies.

Context

WAF provides a default protection policy. You can also customize protection policies to suit your business requirements.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Application Security > WAF.**
3. In the left-side navigation tree of the **WAF** page, choose **Protection Configuration > Website Protection Policies.**
4. Click **Add a protection policy.** In the panel that appears, specify **Policy name** and click **Confirm.**
5. In the **Operation** column of the new protection policy, click the  icon to view or edit the policy.

Default Policy Set

Technical Details

Decode

Attack Detecti... URL Decode JSON Parse Base64 Decode Hexadecimal Conversion

Other Modules Backslash Unescape XML Parse PHP Deserialization UTF-7 Decode

Block Options

HTTP Respons... SQL Injection Detection Module Only ForbidHigh Risk XSS Detection Module Only ForbidHigh Risk

HTTP Request ... Intelligence Module Only ForbidHigh Risk CSRF Detection Module Only ForbidHigh Risk

Detection Tim... SSRF Detection Module Only ForbidHigh Risk PHP Deserialization Detection Module Only ForbidHigh Risk

ASP Code Injection Detection Module Only ForbidHigh Risk SSTI Detection Module Only ForbidHigh Risk

Java Deserialization Detection Module Only ForbidHigh Risk File Upload Attack Detection Module Only ForbidHigh Risk

File Inclusion Attack Detection Module Only ForbidHigh Risk PHP Code Injection Detection Module Only ForbidHigh Risk

Java Code Injection Detection Module Only ForbidHigh Risk Command Injection Detection Module Only ForbidHigh Risk

Server Response Detection Module Disabled Robot Detection Module Disabled

Other Modules

None

Block Options

Block Return 405

| Parameter | Description |
|------------------------------------|---|
| Decode | Select algorithms that you want to use to decode the requests. |
| Attack Detection Modules | Specify the types of attacks that you want to detect and the risk levels of attacks that you want to block. |
| Block Options | Specify the HTTP status code and image that you want WAF to return when it blocks an attack. |
| HTTP Response Processing | Specify Enable HTTP response processing and Response Detection Max Body Size . |
| HTTP Request Body Detection | Specify Response Detection Max Body Size . |
| Detection Timeout | Specify Enable Detection Timeout and Timeout Threshold . |

For example, perform the following steps to configure modules in the **Attack Detection Modules** section:

- i. Move the pointer over a specific module in the **Attack Detection Modules** section. In this example, move the pointer over **SQL Injection Detection Module** and click the **modify** icon.
- ii. In the **SQL Injection Detection Module** panel, configure the following parameters.

| Parameter | Description |
|--------------------------------|--|
| Enabled | Specify whether to enable the detection module. |
| Blocking Threshold | Valid values: NotForbid , Only ForbidHigh Risk , ForbidMedium or High Risk , and Forbid All . |
| Record Threshold | Valid values: Notrecord , Onlyrecord High Risk , recordMedium or High Risk , and record All . |
| Detect Non-Injected SQL | Specify whether to enable detection for NoSQL injection vulnerabilities. |

- iii. Click **OK**.

6. Manage protection policies.

To delete a protection policy, select the protection policy. Then, in the upper-right corner, choose **More > Delete Selected Protection Policies**. In the message that appears, click **OK**.

Note You cannot delete the default protection policy.

21.7.4.2. Create a custom rule

This topic describes how to create a custom rule for Web Application Firewall (WAF).

Context

You can create custom rules to meet different requirements for intrusion detection. You can create, edit, or delete custom rules as an administrator. You can use custom rules to filter out requests that meet specific conditions.

Multiple custom rules are evaluated by using a logical **OR**. If two custom rules use the same conditions but trigger different actions such as blocking traffic or allowing traffic, WAF runs the first rule.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Application Security > WAF**.
3. In the left-side navigation tree of the **WAF** page, choose **Protection Configuration > Customized Rules**.
4. In the upper-right corner, click **Add Rule**. In the **Add Customized Rules** panel, configure the parameters.

Add Customized Rules

For the newly created custom rule, it is recommended to set it to the observation mode first, observe for a period of time and find no false positives, and then turn on the intercept mode.

Type: Block Enabled

Comment *

Risk level: No threat

Matching Pattern *
Parameter value

Add Pattern

Apply to Websites

Advanced

Cancel Confirm

Parameters used to create a custom rule

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|----------------------|---|
| Type | <p>The operating mode of the rule. Valid values: Block, Allow, Monitor, and Detection module control.</p> <ul style="list-style-type: none"> ◦ Block: If an HTTP request meets the conditions of the rule, the HTTP request is blocked. ◦ Allow: If an HTTP request meets the conditions of the rule, the HTTP request is allowed. ◦ Monitor: If an HTTP request meets the conditions of the rule, the HTTP request is recorded and allowed. ◦ Detection module control |
| Comment | The remarks about the rule. We recommend that you enter the purpose of the rule. |
| Risk level | The risk level. Valid values: No threat , Low Risk , Medium Risk , and High Risk . |
| Matching Pattern | <p>The conditions that trigger the rule.</p> <p>Click Add Pattern to specify more than one condition. Multiple conditions are evaluated by using a logical AND. The custom rule takes effect only when all conditions are met.</p> |
| Apply to Websites | The websites that you want the rule to protect. |
| Log Recording Option | Specifies whether to record a log when the rule is triggered. The default value is Enable Log Recording. After Log Recording Option is set to Enable Log Recording, all interception events are recorded in the intrusion detection logs. |
| Attack Type | The type of attack that you want the rule to block. |
| Expiration Time | The time at which the rule expires. |

5. Click **Confirm**.

6. Manage custom rules.

- Edit a rule.

To edit a rule, click the  icon in the **Actions** column.

- Enable a rule.

To enable a rule that is disabled, select the rule and choose **More > Enable Selected Rules**.

- Disable a rule.

To disable a rule that is enabled, select the rule and choose **More > Disable Selected Rules**.

- Export a rule.

To export a rule, select the rule and choose **More > Export Selected Rules**.

- Delete a rule.

To delete a rule that you no longer need, select the rule and choose **More > Delete Selected Rules**.

21.7.4.3. Configure an HTTP flood protection rule

This topic describes how to configure an HTTP flood protection rule.

Context

An HTTP flood attack is a type of DDoS attack that targets the application layer. Attackers use proxy servers or zombies to overwhelm targeted web servers by sending a large number of HTTP requests.

Create an HTTP flood protection rule

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Application Security > WAF.**
3. In the left-side navigation tree of the WAF page, choose **Protection Configuration > HTTP Flood Detection.**
4. Click **Add Rule.** The **Add HTTP Flood Detection Rules** panel appears.
5. Configure parameters and click **Confirm.**

Add HTTP Flood Detection Rules
✕

Rule Mode

Observe
 Blocking Mode

Rule Types

Restrict Users by Policy
 Restrict Known Users

Rule Name *

Target Type

IP
 SESSION

Restricted IP List *
 Fill IP

One IP address or IP address segment per line
 If it is an IP address segment, please use "IP address/subnet mask" format such as
 192.168.100.200

Restriction Mode *

Frobidden
▼

Restricted URL Address *

URL Prefix
▼

http://
▼

Restriction Time

↑
↓

sec
▼

Restrict known users access **http://**

Cancel

Confirm

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|---|---|
| Rule Mode | <p>The action on requests after the HTTP flood protection rule is triggered. Valid values: Blocking Mode and Observe.</p> <ul style="list-style-type: none"> ◦ Blocking Mode: limits the requests that trigger the HTTP flood protection rule. ◦ Observe: records the requests that trigger the HTTP flood protection rule, but does not limit the requests. |
| Rule Types | <p>The type of the HTTP flood protection rule. Valid values: Restrict Users by Policy and Restrict Known Users. The difference between the two types is determined by whether requests of users are initiated from a specific IP address or in a specific session.</p> <ul style="list-style-type: none"> ◦ Restrict Users by Policy: limits requests that meet all the configuration items of the HTTP flood protection rule. Configuration items include Restriction Trigger Threshold, Restricted URL Address, Restriction Mode, Restriction Time, and Statistical Range of Visits in the Advanced section. ◦ Restrict Known Users: limits requests that are initiated from specific IP addresses or in specific sessions based on the HTTP flood protection rule. To achieve this purpose, you must configure the IP address or session list and the limit mode. After you configure the list, the HTTP flood protection rule limits requests based on the list. |
| Rule Name | The name of the HTTP flood protection rule. |
| Target Type | <p>The type of source for requests that are limited. Valid values: IP and SESSION.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note If you set Target Type to SESSION, you can apply the HTTP protection rule only to a website whose User Identification is set to WAF User System. For more information, see Add an Internet website for protection.</p> </div> |
| Restriction Trigger Threshold | If you set Rule Types to Restrict Users by Policy , you must configure the triggering conditions for the HTTP flood protection rule. |
| Restricted URL Address | <p>If you set Rule Types to Restrict Users by Policy, you must specify the URL addresses that are protected based on the HTTP flood protection rule.</p> <ul style="list-style-type: none"> ◦ URL Prefix ◦ URL ◦ Record all IP addresses |
| Restricted IP List or Restricted SESSION List | If you set Rule Types to Restrict Known Users , you must enter the IP addresses or sessions from which you want to limit requests based on the setting of Target Type . You can enter only one IP address or session in each line. |

| Parameter | Description |
|-----------------------------|---|
| Restricted URL Address | <p>If you set Rule Types to Restrict Known Users, you must specify the URL addresses that are protected based on the HTTP flood protection rule.</p> <ul style="list-style-type: none"> ◦ URL Prefix ◦ URL ◦ Restrict user access to all addresses |
| Restriction Mode | <p>The mode in which the HTTP flood protection rule limits requests. Valid values:</p> <ul style="list-style-type: none"> ◦ Forbidden: The rule blocks specific sources from accessing the specified URL address. ◦ Frequency control: The rule limits the frequency at which specific sources access the specified URL address. |
| Restriction Time | <p>The time at which the action specified in the HTTP flood protection rule takes effect.</p> |
| Statistical Range of Visits | <p>If you set Rule Type to Restrict Users by Policy, you can specify the range of data records to limit requests in the Advanced section.</p> <ul style="list-style-type: none"> ◦ Statistics Full Access Data: If you select this option, the frequency of requests is limited when the requests are forwarded to WAF and meet the HTTP flood protection rule, regardless of which data records the requests access. This decreases system performance. ◦ Statistics TOP Access Data: If you select this option, the frequency of requests is limited when the requests meet the preceding conditions and access the top 100 data records. This option helps minimize the decrease in system performance. You can select this option when the number of accessed data records is larger than 100. Note that the top 100 data records are measured based on real-time monitoring. |

Manage HTTP flood protection rules

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Application Security > WAF**.
3. In the rule list, manage existing HTTP flood protection rules. Modification is allowed.
 - Search for a rule.

Click **Filter**. In the **Filter Item** panel, specify filter conditions.
 - Enable a rule.

Select a rule that is in the Disabled state and choose **More > Enable Selected Rules**.
 - Disable a rule.

Select a rule that is in the Enabled state and choose **More > Disable Selected Rules**.
 - Delete a rule.

Select a rule and choose **More > Delete Selected Rules**.

21.7.4.4. Configure the HTTP flood whitelist

This topic describes how to configure the HTTP flood whitelist.

Context

If a request source is trusted, you can add this request source to the HTTP flood whitelist to allow all requests from this source.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Application Security > WAF.**
3. In the left-side navigation tree of the **WAF** page, choose **Protection Configuration > HTTP Flood Detection.**
4. On the **HTTP Flood Detection Whitelist** tab, click **Add Whitelist Item**, add a trusted request source, and then click **Confirm.**

| Parameter | Description |
|---------------|---|
| Type | Select the type of the request source. Valid values: IP and SESSION . |
| IP or SESSION | Specify the IP addresses or sessions based on the setting of Type . You can enter only one IP address or session in each line. |
| Comment | Enter remarks for the request source. |

5. Manage request sources in the whitelist.
 - Search for a request source in the whitelist.
Click **Filter**. In the panel that appears, specify a filter condition or click **Add Filter Item** to specify more filter conditions.
 - Remove a request source from the whitelist.
Select the request source and choose **More > Delete Selected Items**.

21.7.4.5. Manage SSL certificates

This topic describes how to upload or delete SSL certificates.

Context

After you upload an SSL certificate on the **SSL Certificate Management** page, you can select this certificate when you add an HTTPS website for protection on the **Protection Site Management** page.

 **Note** When you add an HTTPS website for protection on the **Protection Site Management** page, you must select the SSL certificate that corresponds to the domain name of the HTTPS website.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Application Security > WAF**.
3. Choose **Protection Configuration > SSL Certificate Management**.
4. Upload a new certificate.
 - i. Click **Upload SSL Certificate**.
 - ii. In the **Name** field, enter a **name for the new certificate**.

We recommend that you enter the domain name for easier management.

 **Note** If your Certificate Authority (CA) certificate and private key are in the same file, select **Include private key in certificate file**.

- iii. In the **File** section, upload the CA certificate file and private key file.
 - iv. Specify **Certificate Password**.
 - v. Click **Confirm**.
5. (Optional) Delete the uploaded SSL certificate.

If an uploaded SSL certificate expires, you can delete it.

 - i. In the SSL certificate list, select the certificate that you want to delete.
 - ii. Choose **More > Delete selected SSL certificate**.
 - iii. In the message that appears, click **OK**.

21.7.4.6. Add Internet websites for protection

This topic describes how to add Internet websites to Web Application Firewall (WAF).

Context

WAF can protect the following types of websites:

- Internet websites.
- Virtual Private Cloud (VPC) websites. For more information about how to add VPC websites to WAF, see [Add a VPC website for protection](#).

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Application Security > WAF**.
3. Choose **Protection Configuration > Protection Site Management**. On the page that appears, click the **Internet Websites** tab.
4. In the upper-right corner, click **Add a site**. The Add Protected Site panel appears.
5. In the Monitoring Information step, configure parameters and click **Next**.

Specify the Internet website that you want WAF to protect. WAF can protect both HTTP and HTTPS websites.

Add Protected Site
✕

1 Monitoring Information

Configure Protected Site Information on WAF

Protected Website Name *

Domain Name *

Remarks

Port Settings * Enable SSL

[Add a group of ports](#)

Create Virtual IP Method

[Next](#)

| Parameter | Description |
|-------------------------------|---|
| Protected Website Name | The name of the website that you want WAF to protect. |
| Domain Name | The domain name of the website. <ul style="list-style-type: none"> ◦ You can use an asterisk (*) as a wildcard. ◦ If you specify multiple domain names, separate them with commas (,). |
| Port Settings | The port that WAF listens on. <ul style="list-style-type: none"> ◦ If the website supports HTTPS requests, select Enable SSL and upload an HTTPS certificate. ◦ If the website can be accessed over multiple ports, click Add a group of ports to add the required ports. |

| Parameter | Description |
|--------------|---|
| Cert Setting | <p>The HTTPS certificate of the website. Valid values: <i>Upload a New Certificate</i> and <i>Choose an Existing Certificate</i>.</p> <ul style="list-style-type: none"> ◦ <i>Upload a New Certificate</i>: If the HTTPS certificate of the website has not been uploaded to WAF, select this option. <p>By default, the HTTPS certificate and private key are separately uploaded. If you select Include private key in certificate file, upload only one file that contains both the HTTPS certificate and private key.</p> <ul style="list-style-type: none"> ◦ <i>Choose an Existing Certificate</i>: If the HTTPS certificate of the website has been uploaded to WAF, select this option. Then, select the required HTTPS certificate from the drop-down list. <p>Note This parameter is required only if you select Enable SSL next to the listening port field.</p> |
| Name | <p>The name of the HTTPS certificate.</p> <p>Note This parameter is required only if you select Enable SSL next to the listening port field.</p> |
| File | <p>The HTTPS certificate and private key.</p> <p>By default, the HTTPS certificate and private key are separately uploaded. If you select Include private key in certificate file next to Name, upload only one file that contains both the HTTPS certificate and private key.</p> <p>Note This parameter is required only if you select Enable SSL next to the listening port field.</p> |
| Virtual IP | <p>The IP address type and virtual IP address.</p> <p>Note You can select an IPv6 address as the virtual IP address for WAF.</p> <p>By default, WAF provides 10 virtual IP addresses. You can add more virtual IP addresses based on your business requirements.</p> <p>Note A virtual IP address is available only for the department to which the creator of the virtual IP address belongs.</p> |

6. In the Request Processing Method step, configure parameters and click **Next**.

Add Protected Site
✕

✓ **Monitoring Information**
Configure Protected Site Information on WAF

2 **Request Processing Method**
Configure WAF server response method

Request Processing Method Forward to Backend Server Redirect
 Respond with Specified Content

Load Balancing Algorithm Weighted Round Robin Least Connections Method
 Source Address Hash

Backend Server Address *

Fill in the back-to-source address Return to the back-to-source instance

http:// : 80 Weight

| Response mode | Parameter | Description |
|---------------------------|--------------------------|---|
| Forward to Backend Server | Load Balancing Algorithm | The algorithm for load balancing. Valid values: Weighted Round Robin , Source Address Hash , and Least Connections Method . |
| | Backend Server Address | The address of the origin server to which WAF forwards inbound traffic. Valid values: Fill in the back-to-source address and Return to the back-to-source instance . <ul style="list-style-type: none"> ◦ Fill in the back-to-source address: Enter the address of the origin server. If you enter multiple addresses, load balancing is performed based on the specified load balancing algorithm. ◦ Return to the back-to-source instance: Enter the address of a specific ECS or SLB instance. If you enter multiple addresses, load balancing is performed based on the specified load balancing algorithm. |
| | X-Forwarded-For | The pass-through mode of the actual source IP address. The X-Forwarded-For (XFF) header is used to identify the actual source IP address of an HTTP client. The header is used for traffic forwarding services, such as HTTP proxy and load balancing. |

| Response mode | Parameter | Description |
|--------------------------------|----------------------|---|
| Redirect | Response Status Code | The HTTP status code that WAF returns when it redirects inbound traffic to a specified address. Valid values: 301, 302, and 307. <ul style="list-style-type: none"> 301: The requested page is permanently moved to another URL. 302: The requested page is temporarily moved to another URL. The requester must continue to use the original URL for future requests. 307: The requested page is temporarily moved to another URL. The requester must continue to use the original URL for future requests. |
| | Redirect address | The destination URL for redirection. |
| Respond with Specified Content | Response Status Code | The HTTP status code that WAF returns when it returns specified content. Valid values: 200, 404, and 503. |
| | Response | The content to return. For example, you can upload an image for the Response parameter. If a user visits the website, WAF returns the uploaded image. |

7. In the Protection Policy step, configure parameters and click Next. Then, go to the **Finish** step.

Note You can configure a protection policy only if you set **Request Processing Method** to **Forward to Backend Server**.

| Parameter | Description |
|---------------------|---|
| Protection Policy | Select a WAF protection policy. For more information, see Configure protection policies . |
| User Identification | Specify whether to enable the user identification feature. <p>Note If you enabled HTTP flood protection for the protected website and set Target Type to SESSION when you configured the HTTP flood protection rule, you must set User Identification to WAF User System.</p> |

21.7.4.7. Add VPC websites for protection

This topic describes how to add Virtual Private Cloud (VPC) websites to Web Application Firewall (WAF) for protection.

Context

WAF can protect the following types of websites:

- Internet websites. For more information about how to add an Internet website for protection, see [Add an Internet website for protection](#).
- VPC websites.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Application Security > WAF > Protection Configuration > Protection Site Management**. On the page that appears, click the **VPC Websites** tab.
3. In the upper-right corner, click **Add a site**. The **Add Protected Site** panel appears.
4. In the Monitoring Information step, configure parameters and click **Next**.

Specify the VPC website that you want WAF to protect. WAF can protect both HTTP and HTTPS websites.

| Parameter | Description |
|-------------------------------|---|
| Protected Website Name | The name of the website that you want WAF to protect. |
| Domain Name | The domain name of the website. <ul style="list-style-type: none"> ◦ You can use an asterisk (*) as a wildcard. ◦ If you specify multiple domain names, separate them with commas (,). |
| Port Settings | The port that WAF listens on. <ul style="list-style-type: none"> ◦ If the website supports HTTPS requests, select Enable SSL and upload an HTTPS certificate. ◦ If the website can be accessed over multiple ports, click Add a group of ports to add the specific ports. |

| Parameter | Description |
|----------------------|--|
| Cert Settings | <p>The HTTPS certificate of the website. Valid values: Upload a New Certificate and Choose an Existing Certificate.</p> <p>Note Specify this parameter only if you select Enable SSL next to Port Settings.</p> <ul style="list-style-type: none"> ◦ Upload a New Certificate: If the HTTPS certificate used by the website has not been uploaded to WAF, select this option. By default, the HTTPS certificate and private key are separately uploaded. If you select include private key in certificate file, upload only a file that contains both the HTTPS certificate and private key. ◦ Choose an Existing Certificate: If the HTTPS certificate used by the website is uploaded to WAF, select this option. Then, select the specific HTTPS certificate from the drop-down list. |
| Name | <p>The name of the HTTPS certificate.</p> <p>Note Specify this parameter only if you select Enable SSL next to Port Settings and set Cert Setting to Upload a New Certificate.</p> |
| File | <p>The HTTPS certificate and private key to upload. By default, the HTTPS certificate and private key are separately uploaded. If you select include private key in certificate file next to Name, upload only a file that contains both the HTTPS certificate and private key.</p> <p>Note Specify this parameter only if you select Enable SSL next to Port Settings and set Cert Setting to Upload a New Certificate.</p> |

5. In the set up VPC step, configure parameters and click **Next**.

Add Protected Site
✕

✓ **Monitoring Information**
Configure Protected Site Information on WAF

2 **set up VPC**
Configure the VPC and related parameters of the protection site

Protected VPC *

Virtual Switch *

Create Virtual IP Method

VPC Virtual IP *

3 **Request Processing Method**
Configure WAF server response method

Previous
Next

| Parameter | Description |
|---------------------------------|--|
| Protected VPC | The VPC to which the website belongs. |
| Virtual Switch | The vSwitch associated with the specified VPC. |
| Create Virtual IP Method | The method to create a virtual IP address. Valid values: Select an existing virtual IP and Create virtual IP . |
| VPC Virtual IP | <ul style="list-style-type: none"> ◦ If you set Create Virtual IP Method to Select an existing virtual IP, select an existing virtual IP address from the VPC Virtual IP drop-down list. ◦ If you set Create Virtual IP Method to Create virtual IP, click Click to Create Vip next to VPC Virtual IP to generate a virtual IP address. |

6. In the Request Processing Method step, configure parameters and click **Next**.

| Response mode | Parameter | Description |
|---------------|---------------------------------|--|
| | Load Balancing Algorithm | The algorithm for load balancing. Valid values: Weighted Round Robin , Source Address Hash , and Least Connections Method . |

| Response mode | Parameter | Description |
|--------------------------------|------------------------|---|
| Forward to Backend Server | Backend Server Address | <p>The IP address of the origin server to which WAF forwards inbound traffic. Valid values: Fill in the back-to-source address and Return to the back-to-source instance.</p> <ul style="list-style-type: none"> ◦ Fill in the back-to-source address: Enter the address of the origin server. If you enter multiple addresses, load balancing is performed based on the specified load balancing algorithm. ◦ Return to the back-to-source instance: Enter the address of a specific Elastic Compute Service (ECS) or Server Load Balancer (SLB) instance. If you enter multiple addresses, load balancing is performed based on the specified load balancing algorithm. |
| | X-Forwarded-For | <p>The passthrough mode of the actual source IP address.</p> <p>The X-Forwarded-For (XFF) header is used to identify the actual source IP address of an HTTP client. The header is used for request forwarding services, such as HTTP proxy and load balancing.</p> |
| Redirect | Response Status Code | <p>The HTTP status code that WAF returns when it forwards inbound traffic to a specified address.</p> <p>Valid values: 301, 302, and 307.</p> <ul style="list-style-type: none"> ◦ 301: The requested page is permanently moved to another URL. ◦ 302: The requested page is temporarily moved to another URL. The requester must continue to use the original URL for future requests. ◦ 307: The requested page is temporarily moved to another URL. The requester must continue to use the original URL for future requests. |
| | Redirect address | The required URL for redirection. |
| Respond with Specified Content | Response Status Code | <p>The HTTP status code that WAF returns when it returns specified content.</p> <p>Valid values: 200, 404, and 503.</p> |
| | Response | <p>The content to return.</p> <p>For example, you can upload an image for the Response parameter. If a user visits the website, WAF returns the uploaded image.</p> |

7. In the Protection Policy step, configure parameters and click Next. Then, go to the **Finish** step.

| Parameter | Description |
|---------------------|---|
| Protection Policy | Select a WAF protection policy. For more information, see Configure protection policies . |
| User Identification | Specify whether to enable the user identification feature. |

21.7.4.8. Verify the configurations of a website on your on-premises server

This topic describes how to verify the configurations of a website on your on-premises server.

Context

Before you use Web Application Firewall (WAF) to scrub traffic destined for a website, we recommend that you verify the configurations of the website on your on-premises server. After you add the virtual IP address and the domain of a website to the hosts file on your on-premises server, the request to access the domain from a local browser passes through WAF first.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. Add the virtual IP address and domain name to the `hosts` file on your on-premises server.

If your computer runs Windows 7, the hosts file is stored in the following path: `C:\Windows\System32\drivers\etc\hosts`.

- i. Open the hosts file by using a text editor, such as Notepad.
- ii. Add the following content to the end of the file: `<The virtual IP address that is assigned by WAF><Protected domain name>`.

```
# localhost name resolution is handled within DNS itself.
# ->::1 localhost
# ->4.115.11.11 localhost
# ->4.115.11.11 example.com
```

Note The IP address preceding the domain name is the virtual IP address that is assigned by WAF.

3. Ping the protected domain name from your on-premises server.
The returned IP address must be the virtual IP address that is assigned by WAF in the hosts file. If the returned IP address is still the IP address of the origin server, refresh the local Domain Name System (DNS) cache.
4. Enter the domain name in the address bar of your browser and press Enter.
If the access configurations on WAF are correct, you can visit the website.
5. Verify the protection capability of WAF.
Simulate a web attack request and check whether WAF blocks the request.
For example, add `?alert(xss)` after the URL. If you try to visit `www.example.com/?alert(xss)`, WAF is expected to block the request.

21.7.4.9. Modify DNS resolution settings

This topic describes how to modify the Domain Name System (DNS) resolution settings to connect your website to Web Application Firewall (WAF).

Context

Before you can modify the DNS resolution settings, you must verify the settings on your computer and make sure that the settings are correct. Then, the traffic destined for your website can be redirected to WAF after you modify the settings.

The domain name of a protected website may not be resolved by a DNS provider. For example, a website may use a Server Load Balancer (SLB) instance to connect to the Internet. In this case, you can perform the following operation to protect the website by using WAF: Specify the virtual IP address that WAF assigns to your website as the back-to-origin IP address of the SLB instance.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Application Security > WAF**.
3. In the left-side navigation tree of the WAF page, choose **Protection Configuration > Protection Site Management**.
4. Find the website whose DNS resolution settings you want to modify and click the  icon in the **Operation** column.
5. On the **Basic Information** tab, record the virtual IP address assigned to the website.
6. Log on to the console of the DNS provider and find the DNS resolution settings for the domain name of the website. Then, change the IP address in the A record to the virtual IP address assigned to the website.

 **Note** We recommend that you set the TTL to 600 seconds in DNS resolution settings. The larger the TTL is, the longer it takes to synchronize and update DNS records.

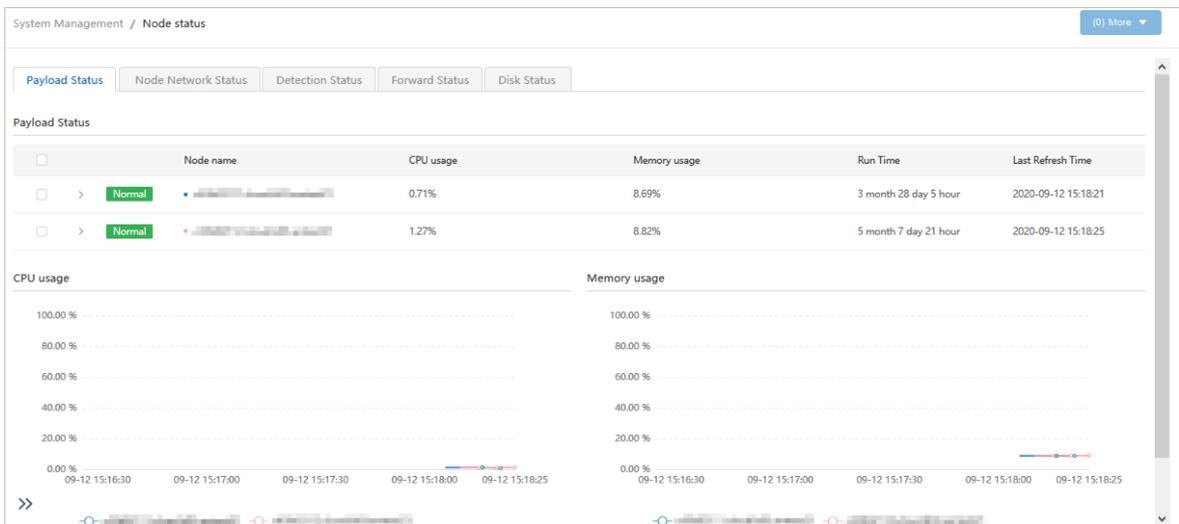
21.7.5. System management

21.7.5.1. View the load status of nodes

This topic describes how to view the load status of Web Application Firewall (WAF) nodes. The status information includes CPU utilization and memory usage. You can identify faults based on the status and check whether scale-out or scale-up is required.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Application Security > WAF**.
3. In the left-side navigation tree of the WAF page, choose **System Info > Node status**.
4. On the **Payload Status** tab, view the load status of WAF nodes.



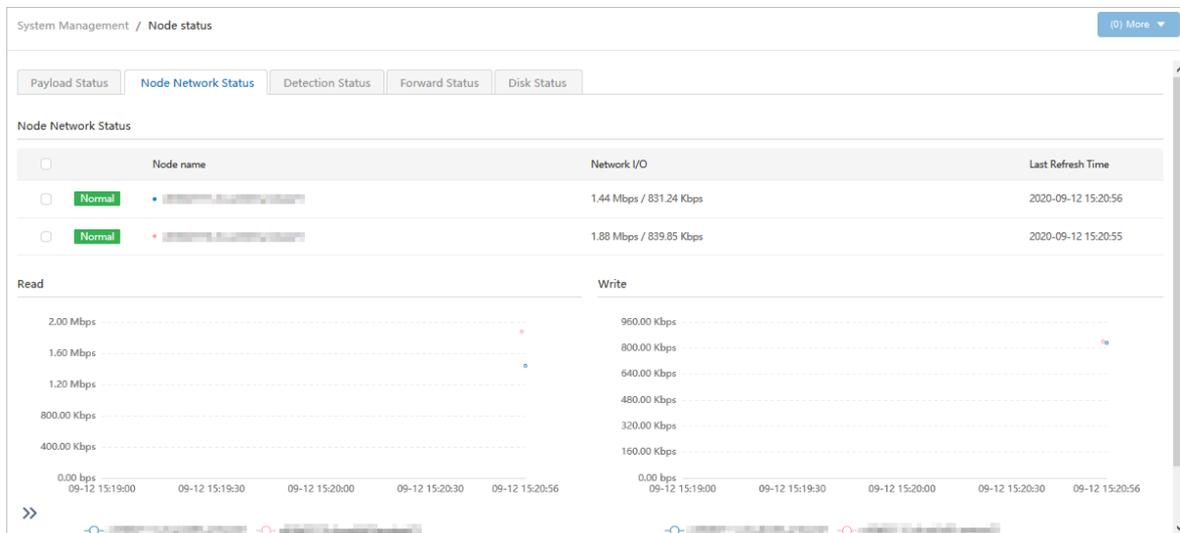
In the Payload Status section, you can view the CPU utilization and memory usage of WAF nodes. In the CPU usage and Memory usage sections, you can view the changes in CPU utilization and memory usage over a specific period of time.

21.7.5.2. View the network status of nodes

This topic describes how to view the network status of Web Application Firewall (WAF) nodes. The status information includes network I/O, traffic detection status, and traffic forwarding status.

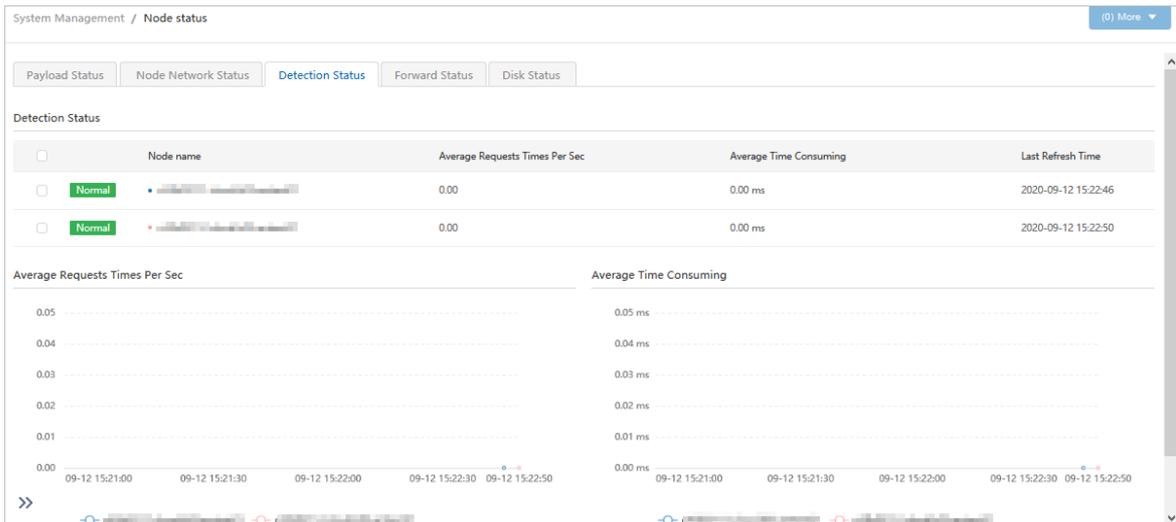
Node network status

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Application Security > WAF**.
3. Choose **System Info > Node status**.
4. Click the **Node Network Status** tab.
5. View the network I/O of WAF nodes.



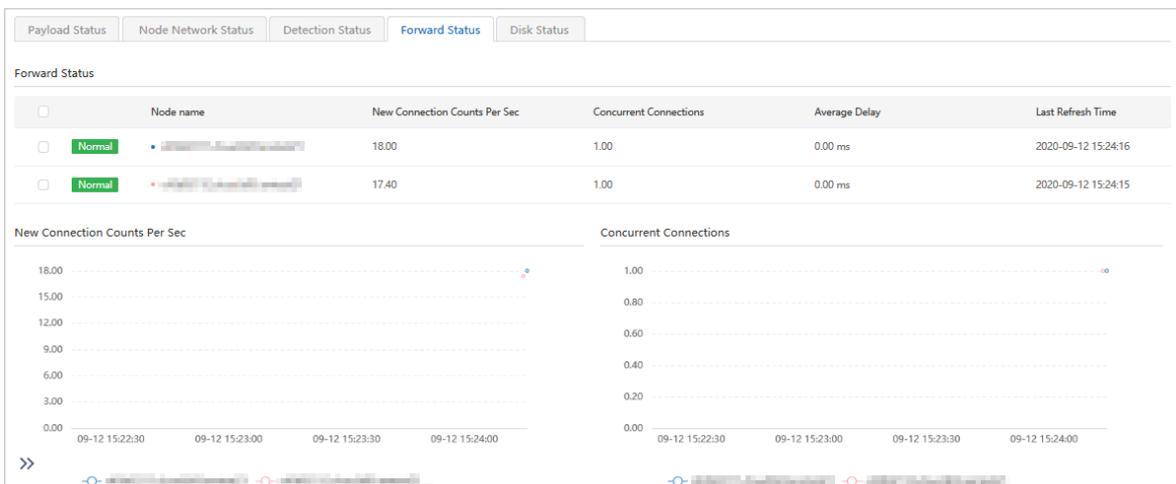
Traffic detection status

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Application Security > WAF > System Info > Node status**.
3. Click the **Detection Status** tab.
4. View the traffic detection status of WAF nodes.



Traffic forwarding status

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Application Security > WAF > System Info > Node status.**
3. Click the **Forward Status** tab.
4. View the traffic forwarding status of WAF nodes.

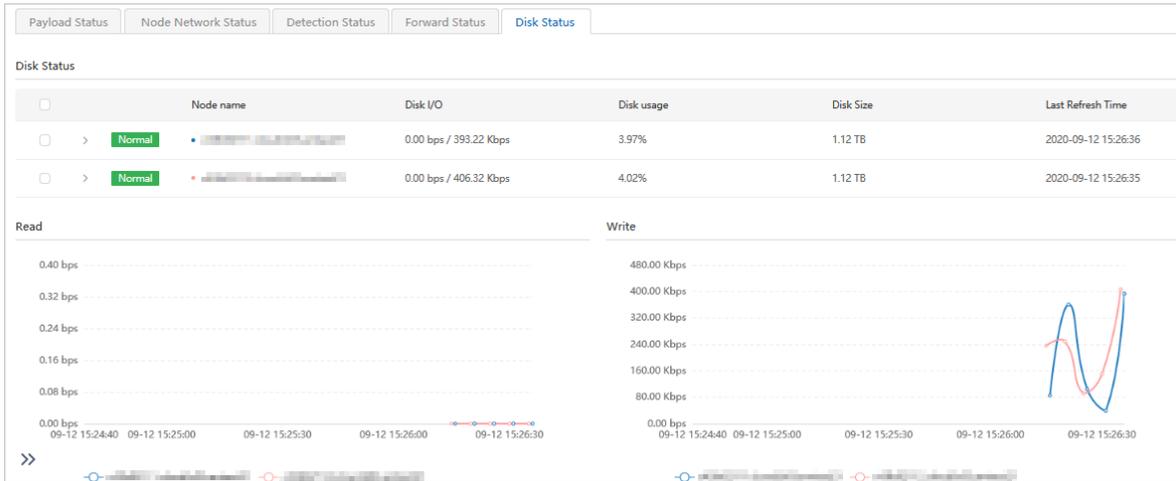


21.7.5.3. View the disk status of nodes

This topic describes how to view the disk status of Web Application Firewall (WAF) nodes. You can identify faults based on the status and check whether scale-out or scale-up is required.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Application Security > WAF.**
3. In the left-side navigation tree of the **WAF** page, choose **System Info > Node status.**
4. Click the **Disk Status** tab to view the disk status of WAF nodes.



In the Disk Status section, you can view the disk I/O and disk usage of WAF nodes. In the **Read** and **Write** sections, you can view the changes in disk reads and writes over a specific period of time.

21.7.5.4. Configure alerts

This topic describes how to add a syslog server to Web Application Firewall (WAF). After the syslog server is added, WAF alert logs can be pushed to the syslog server over the syslog protocol.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Application Security > WAF**.
3. In the left-side navigation tree of the WAF page, choose **System Settings > Syslog Configuration**.
4. On the **Alarm Service Configuration** tab, click **Add alarm service**.
5. In the **Add Alarm Service** panel, configure parameters.

Add Alarm Service ✕

Alarm Type Syslog

Syslog Server* :

RFC RFC3164 ▼

Protocol TCP ▼

Comment

| Parameter | Description |
|---------------|---|
| Syslog Server | The IP address and port number of the syslog server. |
| RFC | The Request for Comments (RFC) document that defines the syslog protocol. Valid values: RFC3164 and RFC5424 . |
| Protocol | The transmission protocol. Valid values: TCP and UDP . |
| Comment | The description of the syslog server. This information facilitates subsequent identification and management. |
| General | The type of alert. Valid values: System Management and System Monitor and Alarm . |

| Parameter | Description |
|-----------|--|
| Security | The module whose alert logs are sent to the syslog server. |

6. Click **Confirm**. The newly added syslog server appears in the list of the Alarm Service Configuration tab.
7. Find the newly added syslog server and click the  icon in the **Operation** column to test whether alerts are sent.
 - o If a message appears, indicating that the alert test is successful, the syslog server is added.
 - o If an error message appears, WAF cannot connect to the syslog server.

21.7.5.5. Configure alert thresholds

This topic describes how to configure alert thresholds.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Application Security > WAF**. On the page that appears, choose **System Settings > Syslog Configuration**.
3. Click the **Alarm Threshold Configuration** tab and click the  icon next to the threshold that you want to modify.
4. In the panel that appears, specify the threshold.

Alarm Service Configuration
Alarm Threshold Configuration



Alarm Service Configuration

System alarm configuration

The system alarm configuration affects the global alarm threshold, please modify it carefully.

| | |
|---------------------------|--|
| Queries per second | No alarm when the number of queries per second is too high  |
| Number of new connections | No alarms if the number of new connections is too high  |
| CPU usage is too high | Continuous CPU usage 1 min over 80 %  |
| Memory usage is too high | Continuous memory usage 1 min over 80 %  |
| Disk usage is too high | Disk usage exceeded 80 %  |

| Threshold | Description |
|---------------------------|--|
| Queries per second | If queries per second exceed this threshold, alerts are sent. If this threshold is set to 0, no alerts are sent. |
| Number of new connections | If a large number of new connections exist, no alerts are sent. |

| Threshold | Description |
|--------------------------|--|
| CPU usage is too high | If CPU utilization exceeds this threshold in a specific period of time, alerts are sent. |
| Memory usage is too high | If memory usage exceeds this threshold in specific a period of time, alerts are sent. |
| Disk usage is too high | If disk usage exceeds this threshold, alerts are sent. |

5. Click OK.

21.8. Security Operations Center (SOC)

21.8.1. View the dashboard

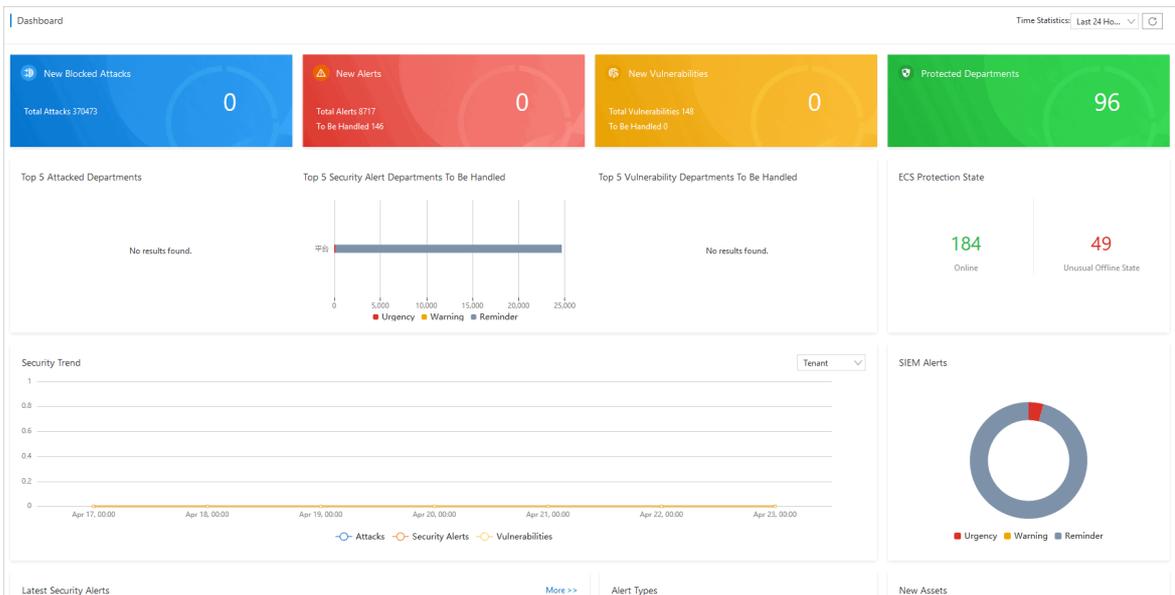
This topic describes how to view the overall security information about the Apsara Stack network environment.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Overview**.
3. In the upper-right corner of the **Dashboard** page, select a time range from the **Time Statistics** drop-down list.

Valid values: *Last 24 Hours, Last 7 Days, and Last 30 Days.*

4. View the overall security information.



The Dashboard page displays the following information:

- **New Blocked Attacks, New Alerts, New Vulnerabilities, and Protected Departments**
- **Top 5 Attacked Departments, Top 5 Security Alert Departments To Be Handled, and Top 5 Vulnerability Departments To Be Handled**
- **Security Trend**, which supports a switchover between Tenant and Platform
- **Latest Security Alerts and Alert Types**
- **Latest Attacks and Attack Types**
- **ECS Protection State, New Assets, and Protected Assets**

21.8.2. Security Monitoring

21.8.2.1. View security monitoring data of tenants

This topic describes how to view the security monitoring data of tenants on the Attack Protections, Security Alerts, and Vulnerabilities tabs.

Attack Protections

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Security Monitoring > Tenant Security Monitoring**.
3. On the **Attack Protections** tab, view all attack events.

You can specify search conditions to search for attack events based on the following table.

| Search condition | Description |
|------------------------------|---|
| Department | The department to which the assets affected by the attack belong. |
| Data source | The data source. |
| Status | The attack status. |
| Attack type | The attack type. |
| Start time and end time | The time range to query the attack event. |
| Attack name or asset keyword | The attack name or the keywords of affected assets. |

4. View details in the attack event list.
5. Click the icons in the upper-left corner to refresh or export the list.
The following list describes how to perform the operations:
 - Click the  icon to refresh the attack event list.
 - Click the  icon to export the attack event list.
6. Find an attack event. In the Actions column, block requests from a specific IP address, create a tag for the event, or view logs and details of the event.
The following list describes how to perform the operations:
 - Block requests from a specific IP address: Click **Block IP Addresses**. In the **Block IP Addresses** dialog box, configure parameters to block requests from a specific IP address. For more information, see [Block IP Addresses](#).
In the upper-right corner, click **View Blocked IPs** to view details about the blocked IP addresses.
 - Create a tag: Click **Tag**. In the **Customize Tag** dialog box, create a tag for the attack event and click **OK**.
 - View logs: Click **View Log**. On the **Cloud Tenant Logs** tab of the **Log Audit** page, view the logs of tenants.
 - View details: Click **Details** to view the details of the attack event.

Security Alerts

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Security Monitoring > Tenant Security Monitoring**. On the page that appears, click the **Security Alerts** tab.
3. (Optional) Specify search conditions.

 **Note** If you want to view all security alerts, skip this step.

All Departments ▾ All Data Sources ▾ Reminder × Warning × Urgency × ▾ All States ▾ All Alert Types ▾ Alert Name/Assets 🔍

| Search condition | Description |
|-----------------------------|--|
| Department | The department to which the assets affected by the alert belong. |
| Data source | The data source. |
| Level | The alert level. You can select one or more alert levels. Valid values: <ul style="list-style-type: none">○ Urgency○ Warning○ Reminder |
| Status | The alert status. |
| Type | The alert type. You can select All Alert Types or a specific alert type. |
| Start time and end time | The time range to query the alert. |
| Alert name or asset keyword | The alert name or the keywords of affected assets. |

4. View details in the security alert list.
5. Click the icons in the upper-left corner to refresh or export the list.



The following list describes how to perform the operations:

- Click the  icon to refresh the security alert list.
- Click the  icon to export the security alert list.

Vulnerabilities

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Security Monitoring > Tenant Security Monitoring**. On the page that appears, click the **Vulnerabilities** tab.
3. Click the **Vulnerabilities** or **Server Configurations** tab.
 - **Vulnerabilities**: provides vulnerability information.
 - **Server Configurations**: provides information about server baseline risks.
4. Specify search conditions.

Note If you want to view all vulnerabilities or server baseline risks, skip this step.

| Search condition | Description |
|--|---|
| Department | The department to which the assets affected by the vulnerability or server baseline risk belong. |
| Level | The level of the vulnerability or server baseline risk. |
| Type | The type of the vulnerability or server baseline risk. |
| Status | The status of the vulnerability or server baseline risk. |
| Start time and end time | The time range to query the vulnerability or server baseline risk. |
| Vulnerability or risk name, asset keyword, or CVE ID keyword | The name of the vulnerability or server baseline risk, or the keywords of affected assets or CVE IDs. |

5. Click the icons in the upper-left corner to refresh or export the list of vulnerabilities or server baseline risks.



The following list describes how to perform the operations:

- Click the icon to refresh the list of vulnerabilities or server baseline risks.
- Click the icon to export the list of vulnerabilities or server baseline risks.

21.8.2.2. View security monitoring data of the Apsara Stack platform

This topic describes how to view the security monitoring data of the Apsara Stack platform on the Attack Protections, Security Alerts, and Vulnerabilities tabs.

Attack Protections

- Log on to [Apsara Stack Security Center](#).
- In the left-side navigation pane, choose **Security Operations Center (SOC) > Security Monitoring > Platform Security Monitoring**.
- On the **Attack Protections** tab, view all attack events.

You can specify search conditions to search for attack events based on the following table.

| Search condition | Description |
|-------------------------|---|
| Data source | The data source. |
| Status | The attack status. |
| Attack type | The attack type. |
| Start time and end time | The time range to query the attack event. |

| Search condition | Description |
|------------------------------|---|
| Attack name or asset keyword | The attack name or the keywords of affected assets. |

- View details in the attack event list.
- Click the icons in the upper-left corner to refresh or export the list.

The following list describes how to perform the operations:

- Click the  icon to refresh the attack event list.
- Click the  icon to export the attack event list.

- Find an attack event. In the Actions column, block requests from a specific IP address, create a tag for the event, or view logs and details of the event.

The following list describes how to perform the operations:

- Block requests from a specific IP address: Click **Block IP Addresses**. In the **Block IP Addresses** dialog box, configure parameters to block requests from a specific IP address. For more information, see [Block IP Addresses](#).
 In the upper-right corner, click **View Blocked IPs** to view details about the blocked IP addresses.
- Create a tag: Click **Tag**. In the **Customize Tag** dialog box, create a tag for the attack event and click **OK**.
- View logs: Click **View Logs**. On the **Cloud Platform Logs** tab of the **Log Audit** page, view the logs of the platform.
- View details: Click **Details** to view the details of the attack event.

Security Alerts

- Log on to [Apsara Stack Security Center](#).
- In the left-side navigation pane, choose **Security Operations Center (SOC) > Security Monitoring > Platform Security Monitoring**. On the page that appears, click the **Security Alerts** tab.
- (Optional)Specify search conditions.

 **Note** If you want to view all security alerts, skip this step.

| Search condition | Description |
|-----------------------------|--|
| Data source | The data source. |
| Level | The alert level. You can select one or more alert levels. Valid values: <ul style="list-style-type: none"> Urgency Warning Reminder |
| Status | The alert status. |
| Type | The alert type. You can select All Alert Types or a specific alert type. |
| Start time and end time | The time range to query the alert. |
| Alert name or asset keyword | The alert name or the keywords of affected assets. |

4. View details in the security alert list.
5. Click the icons in the upper-left corner to refresh or export the list.



The following list describes how to perform the operations:

- Click the  icon to refresh the security alert list.
- Click the  icon to export the security alert list.

Vulnerabilities

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Security Monitoring > Platform Security Monitoring**. On the page that appears, click the **Vulnerabilities** tab and the **Platform Baseline** tab.
3. Specify search conditions.

 **Note** If you want to view all baseline risks of the platform, skip this step.

| Search condition | Description |
|----------------------------|---|
| Level | The level of the baseline risk. |
| Type | The type of the baseline risk. |
| Status | The status of the baseline risk. |
| Start time and end time | The time range to query the baseline risk. |
| Risk name or asset keyword | The risk name or the keywords of affected assets. |

4. Click the icons in the upper-left corner to refresh or export the list of platform baseline risks.



The following list describes how to perform the operations:

- Click the  icon to refresh the list of platform baseline risks.
- Click the  icon to export the list of platform baseline risks.

21.8.2.3. View the global traffic

This topic describes how to view the global traffic, including the average traffic, peak traffic, overall traffic trends, traffic of tenants, and traffic of platforms.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Security Monitoring > Global Traffic Analysis**.
3. On the **Global Traffic Analysis** page, view the global traffic.

You can view the following information on this page:

- o View the average traffic and peak traffic
 - a. In the upper-right corner of the **Global Traffic Analysis** page, select a time range and traffic direction.

Valid values of time ranges: **Last 6 Hours**, **Last 24 Hours**, and **Last 7 Days**.

Valid values of traffic directions: **Inbound** and **Outbound**.
 - b. In the upper-left corner of the Global Traffic Analysis page, view the average and peak traffic of the specified traffic direction within the specified time range.
- o View traffic trends
 - a. In the upper-right corner of the **Global Traffic Analysis** page, select a traffic type.
 - b. View the overall traffic trends of each traffic type within the specified time range.
- o View the traffic of tenants on the **Tenant Traffic** tab
- o View the traffic of platforms on the **Platform Traffic** tab

21.8.3. Asset Management

21.8.3.1. View tenant assets

This topic describes how to view the assets of users. The assets include Elastic Compute Service (ECS) instances, ApsaraDB RDS instances, Object Storage Service (OSS) buckets, Server Load Balancer (SLB) instances, and elastic IP addresses (EIPs).

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Asset Management > Tenant Assets**.
3. Select the required service. Example: **Elastic Compute Service (ECS)**.
4. Specify search conditions to view a specific asset.

 **Note** If you want to view all assets, skip this step.

| Search condition | Description |
|---------------------------|---|
| Department | The department to which the asset belongs. |
| VPC | The virtual private cloud (VPC) to which the asset belongs. |
| Status | The running status of the asset. |
| New | Specifies whether the asset to query is newly added. |
| Server name or IP address | The keywords of the asset name. |

5. View asset information in the asset list.

21.8.3.2. View platform assets

This topic describes how to view the assets of the platform.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Asset Management > Platform Assets.**
3. Specify search conditions to view a specific asset.

 **Note** If you want to view all assets, skip this step.

4. View asset information in the asset list.

21.8.4. Log Analysis

21.8.4.1. View the Log Overview page

This topic describes how to view the logs that are displayed in the widgets on the Log Overview page.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Log Overview.**
3. On the **Log Overview** page, view the widgets of the logs.

You can perform the following operations to **modify** or **delete** the widgets on the **Log Overview** page:

- To modify a widget, click **Modify** in the upper-right corner of the widget.
 - a. In the **Modify** dialog box, reconfigure the **Chart Type**, **Category**, and **Value** parameters. When you configure the **Category** and **Value** parameters, take note of the following points:
 - **Category:** If you set **Chart Type** to **Bar Chart**, **Line Chart**, **Pie Chart**, or **Sheet**, you must specify this parameter.
 - **Value:** If you set **Chart Type** to **Pie Chart** or **Individual Value Plot**, you must specify this parameter.
 - b. Click **Refresh** to preview the widget in the right side of the Modify dialog box.
 - c. Above the widget, enter a new name to rename the widget.
 - d. Click **OK**. The widget is updated on the **Log Overview** page.
- To delete a widget, click **Delete** in the upper-right corner of the widget.

The widgets on the **Log Overview** page are created on the **Log Audit** page. To create a widget, click **Please go to the log audit page to add a chart** in the **Add custom visualization chart** section in the lower part of the **Log Overview** page. On the **Log Audit** page, create a custom widget. For more information, see [View global logs.](#)

21.8.4.2. View global logs

This topic describes how to view global logs. Global logs are classified into logs of tenants, logs of the platform, and logs of data centers based the department to which the logs belong.

View a log widget

1. Log on to [Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Log Audit.**
3. Click the tab on which you want to view the logs of a department. For example, if you want to view the logs of tenants, click the **Cloud Tenant Logs** tab.
4. Specify search conditions and click the  icon. Then, you can view the logs that meet the search conditions in

the sections of log distribution chart and log list.

| Search condition | Description |
|-------------------------|---|
| Department | If you want to view the logs of a tenant, you can specify this search condition. Select the department to which the tenant belongs. |
| Log source | Select the type of the system from which you want to collect the logs, the name of the system from which you want to collect the logs, and the log type from the drop-down list. |
| Duration | Select the time range within which you want to collect the logs. Valid values: Last 15 Minutes , Last 30 Minutes , Last 24 Hours , Last 7 Days , Last 30 Days , and Custom . |
| Start time and end time | Specify the start and end time within which you want to collect the logs. |
| Log content | Enter the log content in the search box. If you want to save the search conditions as frequently used search conditions, click Save Search Condition . If you want to use these frequently used search conditions, click Historical Records . Then, you can view the logs that meet the search conditions. |

If you want to search for logs by using JSON domain-specific language (DSL) statements, click **Advanced Search**. In the **Advanced Search** dialog box, enter JSON DSL statements and click **Submit**.

Create a log widget

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Log Audit**.
3. Click the tab on which you want to view the logs of a department. For example, if you want to view the logs of tenants, click the **Cloud Tenant Logs** tab.
4. In the lower part of the Log Audit page, click the **Visualization** tab.
5. On the **Visualization** tab, configure parameters in the **Chart** section.

| Parameter | Description |
|-----------------------|---|
| Chart Type | The type of widget that you want to display on the Overview page. Valid values: Bar Chart , Line Chart , Pie Chart , Individual Value Plot , and Sheet . |
| Category | This parameter is required only if you select Bar Chart , Line Chart , Pie Chart , or Sheet for Chart Type . The type of item that you want to display in the horizontal axis or the column header of the widget. |
| Value Category | The type of item that you want to display in the vertical axis or the row header of the widget. |
| Value Type | The type of value that you want to display in the widget. Valid values: count , max , min , avg , sum , unique_count , and median . If you want to display multiple value types in the widget, click Add Value Field . Then, you can configure the Value Category and Value Type parameters. |

6. In the upper-left corner of the preview section on the Visualization tab, enter a name for the widget.

7. Click **Update**.

After the widget is created, you can view the widget on the **Log Overview** page.

To configure the content to be displayed in the log list, you can perform the following steps: In the lower part of the Log Audit page, click the **Log** tab. Then, click the  icon. To export the log list, click the  icon.

21.8.4.3. Log configurations

21.8.4.3.1. Manage log sources

This topic describes how to view and manage the log sources that are connected to Security Operations Center (SOC).

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Log Configurations**.
3. Click the **Log Sources** tab. On the **Log Sources** tab, view the overview log information and log list.

This tab provides the following information:

- In the upper section of the **Log Sources** tab, you can view the total volume of log data and the total number of connected log sources.
- In the upper-right corner above the log source list, you can specify conditions to search for specific log sources.

| Search condition | Description |
|---------------------------|--|
| Access system type | The type of the source system in the log source that is connected to SOC. Valid values: Host , Storage , Application , Networking , Data , Security , and Other . |
| Log Type | The type of log that is collected by SOC. Valid values: Operations Log , Operational Log , Alert Log , and Others . |
| Access mode | The mode that is used to collect logs. Valid values: Custom and Built-in . |
| Status | The status of the log source that is connected to SOC. Valid values: On and Off . In the log source list, find a log source and click the  icon in the Status column to set the log source to on or off. In the Tips message, click OK . |

4. Find a log source and click **Modify** in the Actions column.
5. In the **Edit** dialog box, configure the Storage Days and View Permissions parameters. Then, click **OK**.

21.8.4.3.2. Create a log collection task

This topic describes how to create a log collection task.

Prerequisites

If you use Logtail to collect logs, make sure that the following conditions are met:

- The Logtail agent is installed.
- A server group for which you want to configure Logtail is created.

For more information, see [Manage log collectors](#).

1. Log on to [Apsara Stack Security Center](#).

2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Log Configurations**.
3. On the **Log Configurations** page, click the **Log Access Task** tab.
4. On the **Log Access Task** tab, click **Add**.
5. In the **Configure Log Source** step of the **Create Task** wizard, configure parameters and click **Next**.

| Parameter | Description |
|---------------------------|---|
| Task Name | The name of the log collection task. |
| Access System Name | The name of the source system from which you want to collect the logs. Example: Windows operating system. |
| Access System Type | The type of the source system from which you want to collect the logs. Valid values: Host, Storage, Application, Networking, Data, Security, and Other . |
| Log Source Name | The subtype of the log type. |
| Log Type | The type of log that you want to collect. Valid values: Operations Log, Operational Log, Alert Log, and Others . |
| Source | The department to which the logs belong. Valid values: Cloud Tenant, Cloud Platform, and On-Premises Data Center . |

6. In the left-side navigation tree of the **Configure Access Method** step, select a data source type, configure parameters, and then click **Next**.

Valid values of Data Source Type: Syslog, SLS, and Logtail.

- o If you select Syslog, you must configure the following parameters.

| Parameter | Description |
|---------------------------|--|
| IP Address | The IP address or CIDR block used to report syslog logs. |
| Protocol | The network protocol used when the logs are collected. Valid values: UDP and TCP . |
| Access System Type | The type of source system from which you want to collect the logs. Valid values: Host, Application, Networking, and Other . |
| Keyword | The keyword of log data. |

- o If you select SLS, you must configure the following parameters.

| Parameter | Description |
|--------------------|---|
| Log Project | The name of the project in Log Service. |
| Log Store | The name of the Logstore in Log Service. |
| Endpoint | The endpoint used to connect to the project in a specific region. |
| accessKey | The AccessKey ID of your account used to access Log Service. |
| secretKey | The AccessKey secret of your account used to access Log Service. |

- o If you select Logtail, you must configure the following parameters.

| Parameter | Description |
|--|--|
| Name | The name of the log collector. |
| Log type | The type of log that you want to collect. Valid values: JSON , Apsara , Separator , and Regular Expression . |
| Log Pattern | The format of log file names. Example: <i>access*.log</i> . |
| Log Path | The path used to store the logs that you want to collect. Absolute paths and relative paths are supported. You can use wildcards in relative paths. |
| Log Sample | A sample log entry from the logs that you want to collect. |
| Regular Expression to Match First Log Entry | After you enter the sample log entry , click Generate Automatically . A regular expression is generated to match the first line of the log entry. |
| Regular Expression | Click Regular Expression . In the Generate Regular Expression dialog box, select the fields that you want to extract and click Generate Regular Expression . After the regular expression is generated, click OK . |
| Date Format | The date format that is automatically generated based on the extracted time fields. |
| Apply to Server Group | The server group for which you want to configure Logtail. |

7. In the **Parse and Normalize Data** step, configure parameters and click **Next**.

| Parameter | Description |
|--------------------------------|---|
| Data acquisition method | The method used to obtain the sample log entry. Valid values: Automatically Obtain and Manual Input . |
| Data Sample | If you select Manual Input as the data acquisition method, you must enter the sample log entry. |
| Resolver | The parser, which can be selected based on the sample log entry. Valid values: JSON and jsonArray . |
| Extract Values | The content that can be extracted from the logs. Click Add Field . In the Add Field dialog box, configure the Original Extract Key , Original Extract Value , Enrich Data , and TargetType parameters. Then, click OK . Valid values of Enrich Data : Retain Original Field , Tag Field , Delete Field , Rename Field , Automatically Fill System Time , Assign Value to Constant , and Assign Value to Variable . |

8. In the **Generate Task** step, configure the **Log Source Name**, **Storage Duration**, and **View Permissions** parameters.

9. Click **Save**.

After the log collection task is created, the task appears in the list of log collection tasks. You can publish, modify, or delete the task in the **Actions** column of the task.

- **Publish**: After the log collection task is created, logs are not automatically collected. You must click **Publish**. After the task is published, the  icon appears in the **Access Status** column of the task.

You can click the switch in the **Access Status** column to change the status of the log collection task.

- **Edit**: You can click **Edit** to change the configurations in the **Parse and Normalize Data** and **Generate**

Task steps. You cannot change the data source type or log collector.

- **Delete:** If you no longer need the log collection task, you can click **Delete** in the Actions column of the task.

21.8.4.3.3. Manage log collectors

Before you can use Logtail to collect logs, you must install the Logtail agent, add log collectors, and create server groups. This topic describes how to install the Logtail agent, add log collectors, and create server groups.

Install the Logtail agent

If the Logtail agent has been installed, the system automatically uninstalls the existing version of the Logtail agent, deletes the `/usr/local/ilogtail` directory, and then reinstalls the Logtail agent. After the new version of the Logtail agent is installed, it automatically runs, and a startup application is added to the registry.

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Log Configurations**.
3. On the **Log Configurations** page, click the **Collectors** tab.
4. Install the Logtail agent.

On the **Collectors** tab, install the Logtail agent based on the operating system of the server and the installation method that is supported by the Logtail agent.

- If the server runs a Linux operating system and the Logtail agent supports manual installation, perform the following steps:
 - a. In the **Linux Operating Systems** section, click **Download Installation Package** and determine whether to download a 32-bit or 64-bit installation file to your computer.
 - b. Log on to the server as an administrator. Then, run the installation command to install the Logtail agent.
- If the server runs a Windows operating system and the Logtail agent supports manual installation, perform the following steps:
 - a. In the **Windows Operating Systems** section, click **Download Installation Package** to download the installation file to your computer.
 - b. Upload the installation file to the server. For example, you can use an FTP client to upload the installation file to the server.
 - c. Run the installation file on the server as an administrator.

After the Logtail agent is installed, it automatically runs, and a startup application is added to the registry.

Add a log collector

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Log Configurations**.
3. In the **Collectors** section, click **Add Collector**.
4. In the **Add Collector** dialog box, configure the following parameters and click **OK**.

| Parameter | Description |
|-----------------------------|--|
| Collector IP Address | The IP address of the source whose logs you want to collect. |
| Server Name | The name of the server where the Logtail agent is installed. |
| Operating System | The operating system of the server where the Logtail agent is installed. |

After the log collector is added, you can view the information of the log collector in the **Collectors** section.

The information includes **Collector IP Address** and **Server Name**.

Create a server group

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Log Configurations**.
3. In the **Server Groups** section, click **Create Server Group**.
4. In the **Create Server Group** dialog box, configure the following parameters and click **OK**.

| Parameter | Description |
|--------------------------|--|
| Server Group Name | The name of the server group. |
| IP Addresses | <p>The IP addresses of the servers that you want to add to the server group. To add a server to the server group, perform the following steps:</p> <ol style="list-style-type: none"> i. In the Ungrouped Servers section, select the server that you want to add to the group. ii. Click the  icon to add the server to the Servers to Be Grouped section. |

After the server group is created, you can view basic information of the server group in the **Server Groups** section. If you want to view detailed information of a server group or delete a server group, perform the following steps:

- o Find the server group. Then, click **Details** in the Actions column to view the detailed information of the server group.
- o Find the server group. Then, click **Delete** in the Actions column to delete the server group.

21.8.4.3.4. Manage storage policies

This topic describes how to view and configure storage policies for logs.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Log Configurations**.
3. On the **Log Configurations** page, click the **Policies** tab.
4. On the **Policies** tab, view the storage policy for logs.
5. Click **Modify**.
6. On the **Policies** tab, configure parameters such as the storage periods for logs and monitoring data. Then, click **Modify**.

| Section | Parameter | Description |
|--|-----------------------------|--|
| Log Storage (Advanced) | Maximum Log Size | The upper limit of the storage space that can be used to store logs. |
| | Maximum Storage Period | The maximum number of days during which logs can be stored. |
| Security Monitoring Data Storage (Advanced) | ApsaraDB RDS Storage Period | The maximum number of days during which monitoring data can be stored. |

21.8.4.4. Security Audit

21.8.4.4.1. Overview

A security audit refers to the systemic and independent inspection and verification of activities and behavior in the computer network environment. Delegated by property owners and authorized by management authorities, professional auditors give their assessments according to relevant laws and regulations. When the administrator needs to back track system operations, the administrator can perform a security audit.

Security audits are long-term security management activities throughout the lifecycle of cloud services. The security audit feature of Apsara Stack Security can collect system security data, analyze weaknesses in system operations, report audit events, and classify audit events into important, moderate, and low risk levels. The security administrator views and analyzes audit events to continuously improve the system and ensure the security and reliability of cloud services.

21.8.4.4.2. View security audit overview

This topic describes how to view the summarized information about security audit.

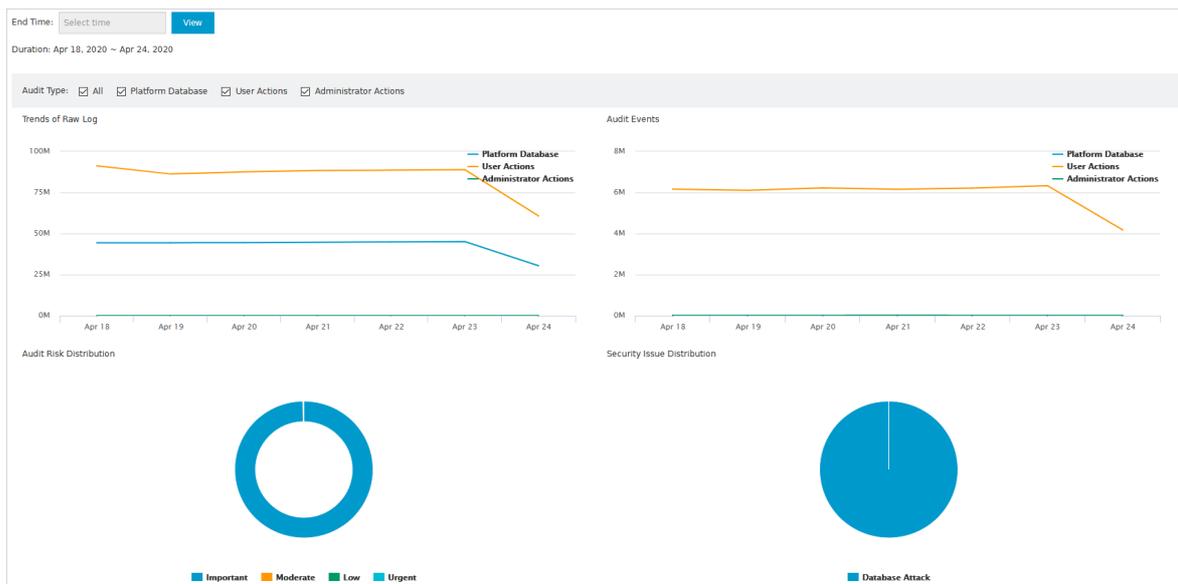
Context

The **Overview** tab provides reports on the raw log trend, audit event trend, audit risk distribution, and security event distribution. The reports are displayed in run charts or pie charts to help security administrators analyze the trend of risks in your cloud services.

On the **Overview** tab, security administrators can check the number of log entries and the storage usage in a specific time range.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Security Audit**. On the page that appears, click the **Overview** tab.
3. On the **Overview** tab, view the audit summary for the last seven days.



- o Trends of Raw Log

This chart displays the trend of logs generated by physical servers, network devices, ApsaraDB RDS instances, Elastic Compute Service (ECS) instances, and API calls in the last seven days. Security administrators can analyze the trend to check whether the number of log entries is at a normal level.

- Audit Events

This chart displays the trend of audit events that are generated by physical servers, network devices, ApsaraDB RDS instances, ECS instances, and API calls in the last seven days. Security administrators can analyze the trend to check whether the number of audit events is at a normal level.

- Audit Risk Distribution

This chart displays the percentage distribution of audit events at different risk levels in the last seven days. Risk levels are important, moderate, and low. Security administrators can analyze the trend to check whether the audit events are at acceptable risk levels.

- Security Issue Distribution

This chart displays the percentage distribution of different event types in the last seven days. Security administrators can analyze this chart to check for the most frequent audit events and identify high-risk events to improve security protection.

- Log Size

This chart displays the volume of online logs and offline logs. If these logs consume many storage resources, we recommend that you back up required audit logs and delete unnecessary logs.

- Audit Log Size

This chart displays the size of logs for each audit type.

4. View the audit summary in a specific time range.

- i. Specify **End Time** as the end of the time range to query.

- ii. In **Audit Type**, select the audit types to query.

- iii. Click **View** to view the audit summary in the last seven days before the specified end time.

21.8.4.4.3. Query audit events

This topic describes how to query audit events.

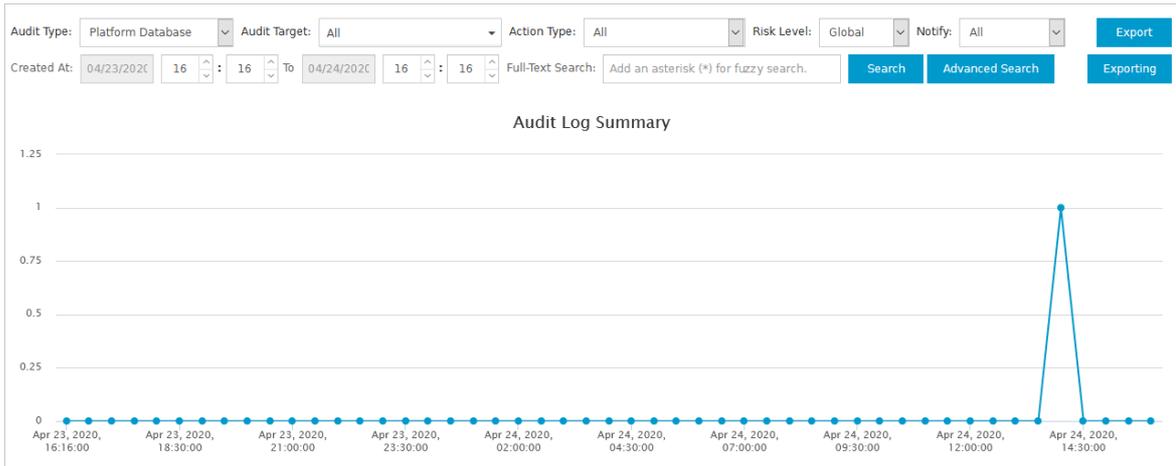
Context

On the **Audit Query** tab, you can view the details of audit events, including the log creation time, audit type, audit object, action type, risk level, and log content.

The system matches the logs that are collected by a security audit module with audit rules. If the log content matches the regular expression in an audit rule, an audit event is reported. For more information about audit rules, see [Add an audit policy](#).

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Security Audit**. On the page that appears, click the **Audit Query** tab.
3. On the **Audit Query** tab, configure query conditions to view audit events within the specified time range.



- o Basic query
 - a. Configure **Audit Type**, **Audit Target**, **Action Type**, **Risk Level**, and **Notify**.
 - b. Specify a time range to query.
 - c. In the **Full-Text Search** search box, enter a keyword.
 - d. Click **Search**.
 - o Advanced query

In addition to the basic query conditions, you can configure advanced query conditions.

 - a. Configure basic query conditions.
 - b. Click **Advanced Search**.
 - c. Below **Filter Condition**, configure **User**, **Target**, **Action**, **Result**, and **Cause**.
 - d. Click **Save**.
4. Click **Export** to export the audit events.
- Download the exported file for analysis. For more information, see [Manage export tasks](#).

21.8.4.4.4. View raw logs

This topic describes how to view raw audit logs.

Context

On the **Raw Log** tab, you can view the raw logs generated by a running audit object. Raw logs contain information that is required for debugging. Security administrators can use these raw logs to troubleshoot system failures.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Security Audit**. On the page that appears, click the **Raw Log** tab.
3. On the **Raw Log** tab, configure query conditions to view the log summary chart and raw logs within a specific time range.
 - i. Specify **Audit Type** and **Audit Target**.
 - ii. Enter a keyword.
 - iii. Specify a time range to query.
 - iv. Click **Search**.
4. Click **Export** to export the data.

Download the exported file for analysis. For more information, see [Manage export tasks](#).

21.8.4.4.5. Manage log sources

This topic describes how to view and manage log sources.

Context

You can view the number of log entries by log type or log source. You can also specify whether to display logs.

- The Log Types sub-tab provides the number of all log entries for a specific audit object of a specific device instance.
- The Log Sources sub-tab provides the number of log entries for all audit objects of a specific device instance.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Security Audit**. On the page that appears, click the **Log Sources** tab.
3. Click the **Log Types** tab and view the number of log entries for each audit object.

You can view the number of log entries that are recorded on the current day and the number of log entries that are recorded during the last 30 days for each audit object.

If you do not want to display the log entries for an audit object, perform the following steps:

- i. Find the audit object and click **Hide** in the **Actions** column.
- ii. In the Note message, click **Confirm**.

 **Note** The process that is used to display the log entries for an audit object is similar to the process that is used to hide the log entries.

4. Click the **Log Sources** sub-tab and view the number of log entries for each device instance.

You can view the number of log entries that are recorded on the current day and the number of log entries that are recorded during the last 30 days for each device instance.

If you do not want to display the log entries for an audit object from a specific device instance, perform the following steps:

- i. Find the device instance and click **Hide** in the **Actions** column.
- ii. In the Note message, click **Confirm**.

 **Note** The process that is used to display the log entries for an audit object is similar to the process that is used to hide the log entries.

21.8.4.4.6. Policy settings

21.8.4.4.6.1. Manage audit rules

This topic describes how to create, modify, or delete an audit rule.

Context

If a log entry matches an audit rule, an audit event is reported. You can specify regular expressions in an audit rule to match log entries. A regular expression defines a matching pattern for character strings and can be used to check whether a string contains a specific substring. The following table provides examples about the pattern.

| Regular expression | Description |
|--------------------------|--|
| <code>^d{5,12}\$</code> | Matches the consecutive numbers from the fifth number to the twelfth number. |
| <code>load_file\{</code> | Matches the "load_file{" string. |

The security audit module defines the default audit rule based on the string that is generated in the log. This applies when an audit event is reported. The security administrator can also customize audit rules based on the string that is generated in the log. This applies when the system encounters an attack.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Security Audit**. On the page that appears, click the **Policies** tab.
3. On the **Policies** tab, click the **Audit Rules** sub-tab.
4. Create an audit rule.
 - i. Click **New** in the upper-right corner.
 - ii. In the **Add Policy** dialog box, configure parameters.

Add Policy

Policy Name:

Audit Type:

Audit Target:

Action Type: Risk Level:

Notify:

Filter Condition:

| | | | |
|--------|-------|--------------------------|---|
| User | Equal | Enter a user | x |
| | | + | |
| Target | Equal | Enter a target | x |
| | | + | |
| Action | Equal | Enter a command | x |
| | | + | |
| Result | Equal | Search by result keyword | |

Add Cancel

- iii. Click **Add**.

The system sends an alert email to the specified alert recipient after you create an audit rule. This applies if one string in an audit log of the specified audit type, audit object, or risk level matches the regular expression of the audit rule.

5. Manage audit rules.

You can create, query, disable, enable, and delete audit rules.

- Query audit rules

Specify **Audit Type** and **Audit Target**. Enter a keyword in the search box and click **Search**.

- Disable an audit rule

Find the audit rule that you want to disable and click **Disable** in the **Actions** column.

- Enable an audit rule

Find the audit rule that has been disabled and click **Enable** in the **Actions** column.

- Delete an audit rule

Find the audit rule that you want to delete and click **Delete** in the **Actions** column.

 **Note** You can delete only custom rules.

21.8.4.4.6.2. Configure alert recipients

This topic describes how to configure the recipients of alerts on audit events.

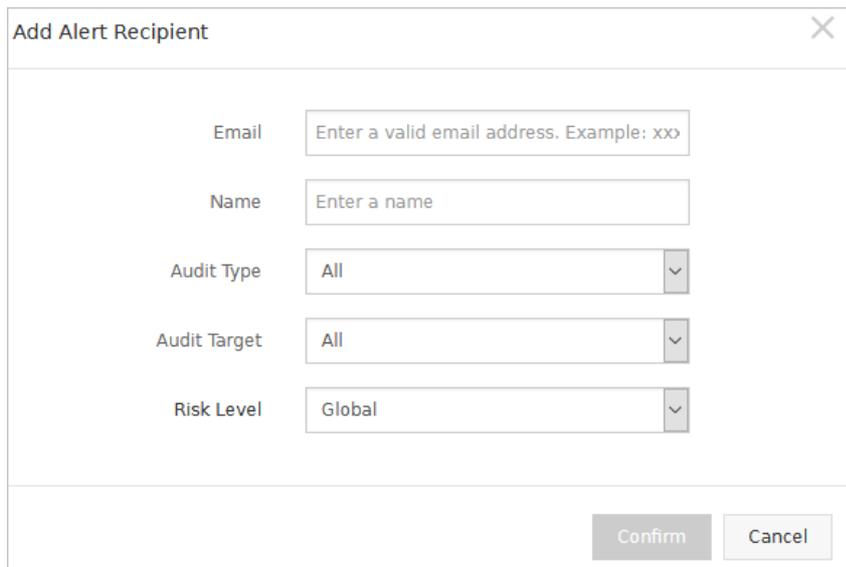
Context

You can add an alert recipient by entering an email address that can be used to receive alerts on audit events.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Security Audit**. On the page that appears, click the **Policies** tab.
3. On the **Policies** tab, click the **Alert Settings** sub-tab.
4. Create an alert recipient.
 - i. Click **New**.

- ii. In the **Add Alert Recipient** dialog box, configure parameters.



- iii. Click **Confirm**.
5. Manage alert recipients.
- o Search for alert recipients
 - Specify **Audit Type**, **Audit Target**, and **Risk Level**, enter the keyword of the email address, and then click **Search**.
 - o Delete alert recipients
 - Find the email address that you want to delete and click **Delete** in the **Actions** column.

21.8.4.4.6.3. Manage archives of events and logs

This topic describes how to query and download the archives of audit events and raw logs.

Context

You can download the archives of events and logs to analyze audit events. This ensures the security of the Apsara Stack environment.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Security Audit**. On the page that appears, click the **Policies** tab.
3. On the **Policies** tab, click the **Archiving** sub-tab.
4. Query the archives of events and logs.
 - i. Specify **Audit Type** and **Archiving Type**.
 - ii. Specify a time range to query.
 - iii. Click **Search**.
5. Find the file where the archive information is stored and click **Download** in the **Actions** column to save the archive file to your on-premises machine.

21.8.4.4.6.4. Manage export tasks

This topic describes how to download or delete exported audit events and logs.

Context

You can export audit events or logs on the **Audit Query** or **Raw Log** tab of the Security Audit page. After you export audit events or logs, you can manage the export tasks on the Exporting sub-tab.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Security Audit**. On the page that appears, click the **Policies** tab.
3. On the **Policies** tab, click the **Exporting** sub-tab.
4. View the created export tasks.

| Created At | Export Task ID | Task Type | Filter Condition | Task Status | Format | Actions |
|------------|----------------|-----------|------------------|-------------|--------|---------|
|------------|----------------|-----------|------------------|-------------|--------|---------|

5. Click **Download** to download audit events or logs to your on-premises server.
6. Click **Delete** to delete the export task.

21.8.4.4.6.5. Modify system settings

This topic describes how to configure system parameters for security audit.

Context

You can configure system parameters to specify the maximum number of system alerts per day and the maximum number of audits per day for raw logs.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Log Analysis > Security Audit**. On the page that appears, click the **Policies** tab.
3. On the **Policies** tab, click the **System Settings** sub-tab.
4. Find the configuration item that you want to modify and click **Edit** in the **Actions** column.

| ID | Description | Updated At | Value | Actions |
|----|--|------------------------|-------|----------------------|
| 1 | Maximum Alerts per Day | Nov 20, 2019, 00:44:19 | 1000 | Edit |
| 2 | Total Logs Audited per Day (GB/day) | Nov 20, 2019, 00:44:19 | 500 | Edit |
| 3 | Database Logs Audited per Day (GB/day) | Nov 20, 2019, 00:44:19 | -1 | Edit |
| 4 | Server Logs Audited per Day (GB/day) | Nov 20, 2019, 00:44:19 | -1 | Edit |
| 5 | Network Device Logs Audited per Day (GB/day) | Nov 20, 2019, 00:44:19 | -1 | Edit |
| 6 | User Operation Logs Audited per Day (GB/day) | Nov 20, 2019, 00:44:19 | -1 | Edit |
| 7 | Administration Logs Audited per Day (GB/day) | Nov 20, 2019, 00:44:19 | -1 | Edit |

5. Enter a required value in the Value column and click **Confirm** in the Actions column.

21.8.5. Rules

21.8.5.1. Create an IPS rule for traffic monitoring

This topic describes how to create an intrusion prevention system (IPS) rule for traffic monitoring in Cloud Firewall. Cloud Firewall has built-in IPS rules. This topic describes how to customize IPS rules based on your business requirements and network environment.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Rules**.
3. On the Rules page, click the **Cloud Firewall IPS Rules** tab.
4. Click **Create Rule**.
5. In the **Create Rule** panel, configure parameters.

| Parameter | Description |
|-------------------------|--|
| Rule Name | The name of the IPS rule. We recommend that you enter an informative name for easy management. |
| Rules Engine | The rules engine that you want to use. Valid values: Basic Policies and Virtual Patches . |
| Attack Type | The type of attack to be detected by the rule. |
| Severity | The level of the severity. Valid values: Low , Medium , and High . |
| CVE | The Common Vulnerabilities and Exposures (CVE) ID of the vulnerability to be listed in the rule. <div style="background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note CVE provides a list of the public security vulnerabilities. CVE IDs are allocated by a CVE Numbering Authority (CNA).</p> </div> |
| Application | The name of the attacked application. |
| Rule Mode | The mode of the rule. Valid values: Packet and Traffic . |
| Direction | The direction of traffic to be monitored by the rule. Valid values: Inbound and Outbound , Inbound , and Outbound . |
| Rule Content | The content of the rule, which must be specified by using the Snort syntax. <div style="background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note To prevent negative impact on your business, make sure that you enter valid content for the rule.</p> </div> |
| Rule Description | The description of the rule. We recommend that you enter information such as the purpose or impact of the rule. |
| Description | The remarks for the rule. We recommend that you enter information such as the purpose or impact of the rule. |

6. Click **OK**.

21.8.5.2. Manage IPS rules of Cloud Firewall

This topic describes how to view, enable, and disable the intrusion prevention system (IPS) rules of Cloud Firewall.

Context

On the **Cloud Firewall IPS Rules** tab of the Rules page in Apsara Stack Security Center, you can view the built-in and custom IPS rules, and enable or disable the rules based on your business requirements.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Rules**.
3. On the Rules page, click the **Cloud Firewall IPS Rules** tab.
4. Manage IPS rules of Cloud Firewall.

In the list of IPS rules, you can view rule details, enable rules, and disable rules.

- o View rule details

Find the rule whose details you want to view and click **Details** in the **Actions** column to view the rule details.

- o Enable a rule

Find the rule that you want to enable and turn on the switch in the **Enable or not** column to change the status of the rule from **Disable** to **Enable**.

- o Disable a rule

If a rule is not suitable for your business, you can disable the rule.

Find the rule that you want to disable and turn off the switch in the **Enable or not** column to change the status of the rule from **Enable** to **Disable**.

21.8.5.3. Create IDS rules for traffic monitoring

This topic describes how to create intrusion detection system (IDS) rules for traffic monitoring.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Rules**.
3. On the Rules page, click the **Traffic Monitoring IDS Rules** tab.
4. Click **Create Rule**.
5. In the **Create Rule** panel, configure parameters.

| Parameter | Description |
|---------------------|---|
| Rule Name | The name of the IPS rule. We recommend that you enter an informative name for easy management. |
| Rules Engine | The rules engine that you want to use. Valid values: Basic Policies and Virtual Patches . |
| Attack Type | The type of attack to be detected by the rule. |
| Severity | The level of the severity. Valid values: Low , Medium , and High . |
| CVE | The Common Vulnerabilities and Exposures (CVE) ID of the vulnerability to be listed in the rule. <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e6f2ff;"> <p> Note CVE provides a list of the public security vulnerabilities. CVE IDs are allocated by a CVE Numbering Authority (CNA).</p> </div> |

| Parameter | Description |
|-------------------------|--|
| Application | The name of the attacked application. |
| Rule Mode | The mode of the rule. Valid values: Packet and Traffic . |
| Direction | The direction of traffic to be monitored by the rule. Valid values: Inbound and Outbound , Inbound , and Outbound . |
| Rule Content | The content of the rule, which must be specified by using the Snort syntax. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note To prevent negative impact on your business, make sure that you enter valid content for the rule. </div> |
| Rule Description | The description of the rule. We recommend that you enter information such as the purpose or impact of the rule. |
| Description | The remarks for the rule. We recommend that you enter information such as the purpose or impact of the rule. |

6. Click OK.

21.8.5.4. Manage IDS rules for traffic monitoring

This topic describes how to view, enable, and disable intrusion detection system (IDS) rules for traffic monitoring.

Context

On the **Traffic Monitoring IDS Rules** tab, you can view the built-in and custom IDS rules. It can also be used to enable or disable the rules based on your business requirements.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Rules**.
3. On the Rules page, click the **Traffic Monitoring IDS Rules** tab.
4. Manage IDS rules for traffic monitoring.

In the IDS rule list, you can view rule details, enable rules, and disable rules.

- o View rule details

Find the rule whose details you want to view and click **Details** in the **Actions** column to view the rule details.

- o Enable a rule

Find the rule that you want to enable and turn on the switch in the **Enable or not** column to change the status of the rule from **Disable** to **Enable**.

- o Disable a rule

If a rule is not suitable for your business, you can disable the rule.

Find the rule that you want to disable and turn off the switch in the **Enable or not** column to change the status of the rule from **Enable** to **Disable**.

21.8.5.5. Specify custom thresholds for DDoS traffic scrubbing policies and traffic redirection

This topic describes how to specify custom thresholds for DDoS traffic scrubbing policies and traffic redirection. Default thresholds are provided. If you want to specify custom thresholds, perform the following steps:

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Rules**.
3. Specify a custom threshold for a DDoS traffic scrubbing policy.
 - i. Click the **AliGuard Rules** tab to open the **Scrubbing Policy** tab.
 - ii. Find the policy for which you want to specify a custom threshold and click **Modify Threshold** in the **Actions** column.
 - iii. In the **Modify Threshold** dialog box, enter a threshold value.
 - iv. Click **OK**.
4. Specify a custom threshold for traffic redirection.
 - i. Click the **AliGuard Rules** tab and the **Reroute Threshold** tab.
 - ii. Find a specific rule and click **Modify Threshold** in the **Actions** column.
 - iii. In the **Modify Threshold** dialog box, enter threshold values.
 - iv. Click **OK**.

21.8.5.6. View Server Guard rules

This topic describes how to view the operations of Server Guard rules. You can view the list of vulnerabilities, baselines, and host exceptions.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Rules**. On the page that appears, click the **Server Guard Rules** tab.
3. In the overview section, you can view the total number of **vulnerability libraries**, number of **baselines**, and number of **host exceptions** as well as the available engines.
4. View the vulnerability list.
 - i. Click the **Vulnerabilities** tab.
 - ii. In the overview section, you can view the total number of **Linux vulnerabilities**, total number of **Windows vulnerabilities**, total number of **Web-CMS vulnerabilities**, and total number of **urgent vulnerabilities**.
 - iii. Specify search conditions to view the vulnerabilities that meet the search conditions.

 **Note** If you want to view all vulnerabilities, skip this step.

In the vulnerability list, you can view the **vulnerability name**, **CVE ID**, **vulnerability type**, **system**, **update time**, and **status**.

5. View the baseline list.
 - i. Click the **Baselines** tab.
 - ii. In the overview section, you can view the numbers of baseline types and check items.

- iii. Specify search conditions to view the baselines that meet the search conditions

 **Note** If you want to view all baselines, skip this step.

In the baseline list, you can view the **baseline type**, **check item category**, **check item name**, **risk level**, **update time**, and **status**.

6. View the host exception list.

- i. Click the **Server Exceptions** tab.
- ii. In the overview section, you can view the number of **rule alert subcategories**, number of webshells, and number of malicious viruses.
- iii. Specify search conditions to view the host exceptions that meet the search conditions.

 **Note** If you want to view all exceptions, skip this step.

In the host exception list, you can view the **subcategory name**, **rule category**, **risk level**, **update time**, **source**, and **status**.

21.8.6. Threat intelligence

21.8.6.1. Enable the service configuration feature

The threat intelligence module integrates threat monitoring and big data analysis. This can be used to obtain the latest information about developments in the threat intelligence field. After you enable the service configuration feature, the system starts to monitor and collect threat intelligence. This topic describes how to enable the service configuration feature.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Threat Intelligence > Service Configuration**.
3. On the **Service Configuration** page, view the data types and descriptions on the **Situation Awareness** and **Web application firewall** tabs.
4. Click the tab where you want to enable threat monitoring and turn on **Activation status**.

After you turn on **Activation status**, the system starts to monitor and collect threat intelligence for the data types listed on the tab.

What's next

After you enable the service configuration feature, choose **Security Operations Center (SOC) > Threat Intelligence**. On the **Overview** page, you can view the overall situation and statistics of threats during the last 30 days. For more information, see [View the Overview page](#).

21.8.6.2. View the Overview page

The Overview page displays the overall situation and statistics of threats to your assets over the last 30 days.

Prerequisites

The **service configuration** feature is enabled. For more information, see [Enable the service configuration feature](#).

Procedure

1. Log on to [Apsara Stack Security Center](#).

2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Threat Intelligence > Overview**.
3. On the **Overview** page, view the statistics and threats that are detected on Apsara Stack services by the threat intelligence module.

On the **Overview** page, you can perform the following operations:

- View **Total malicious metric intelligence**

In the **Total malicious metric intelligence** section of the **Overview** page, view the information about the detected threats on Apsara Stack services. The information includes the number of malicious IP addresses, malicious domain names, and malicious URLs.

- View **Threat trends in the last 30 days**
- Search for an IP address to check whether the IP address is malicious.

Enter the IP address that you want to check in the search box in the upper-right corner of the IP Report page and click the  icon. To view the details of the IP address, choose **Security Operations Center**

(SOC) > **Threat Intelligence > IP Address Search**. On the page that appears, click the **IP Report** tab. For more information, see [Search for an IP address](#).

- View **Top 10 active IP malicious addresses**

In the **Top 10 active IP malicious addresses** section of the **Overview** page, view the information about the top 10 malicious IP addresses. The information includes **IP address**, **First malicious observation**, **Last malicious observation**, and **Malicious label**.

21.8.6.3. Search for and view the information about a suspicious or malicious IP address

The threat intelligence module allows you to search for threat intelligence. This module helps you handle suspicious or malicious IP addresses at the earliest opportunity.

Prerequisites

The **service configuration** feature is enabled. For more information, see [Enable the service configuration feature](#).

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Threat Intelligence > IP Address Search**.
3. In the search box on the **Search** page, enter the suspicious or malicious IP address that you want to query and click the  icon.
4. On the **IP Report** page, view **Threat Level**, **Basic Information**, **Threat Overview**, **IP Details**, and **Attack Risk Level Analysis** of the IP address.

You can view the following information on the IP Report page:

- **Threat Level:** View the threat level of the suspicious or malicious IP address.

The threat intelligence module classifies IP addresses into three threat levels. The levels are normal, suspicious, and high-risk. If the IP address is identified as high-risk, we recommend that you handle the IP address at the earliest opportunity.

- **Basic Information:** View the basic information about the suspicious or malicious IP address.

The basic information includes the server in a data center, Abstract Syntax Notation One (ASN.1), country and city to which the IP address belongs, and the number of domain names for the IP address.

- View the statistics of the suspicious or malicious IP address.
 You can view **Threat Overview**, **IP Details**, and **Threat Details** of the IP address.
 - The **Threat Overview** tab displays **Top 5 Attack Preference**, **Attack Number**, and **Attack Level Analysis** of the IP address.
 - The **IP Details** tab displays **WHOIS** and **IP reverse check information** of the IP address.
 - The **Threat Details** tab displays the **threat tag** list of the IP address. The list contains specific information. The information includes the intelligence source, time when the IP address is detected for the first time, time when the IP address is active for the last time, and threat tag.

21.8.7. Create a report task

This topic describes how to create a report task. After you create a report task, the system sends reports on a regular basis.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > Report Management**.
3. On the Report Management page, click **Create Report**.
4. In the **Create Report** dialog box, configure parameters.

| Parameter | Description |
|-------------|---|
| Report Name | The name of the report task. We recommend that you enter information such as the report purpose for easier identification and management. |
| Task Type | The type of the task. Valid values: Daily Report , Weekly Report , and Monthly Report . |
| Department | The department related to the report. |
| Email Box | The email address of the report recipient. If you enter more than one email address, separate the email addresses with commas (,). |

5. Click **Confirm**.

Result

In the report task list, you can view, edit, and delete the newly created report tasks.

21.8.8. System Configurations

21.8.8.1. View and manage metrics

Apsara Stack Security Center allows you to monitor security services. This helps find performance bottlenecks at the earliest opportunity. Then, you can scale out, scale up, or downgrade services to prevent system failures. This topic describes how to view the information about security services in Apsara Stack Security Center and how to manage metrics.

View information about overall system monitoring

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Security Monitoring**.
3. On the **Overall System Monitoring** tab, view the overall information about security services in Apsara Stack Security Center and the list of monitored security services.

| Services | Metrics | Description | Monitoring Items | Abnormal Items | State | Actions |
|-----------|---------------------|---|------------------|----------------|--------|-------------------------|
| Aegis | System Performance | CPU utilization, QPS, IOPS, disk usage, and memory usage of a database instance | 5 | 0 | Normal | Details |
| WAF | Service Performance | Processing errors of Log Service for WAF, engines, and connection failures | 8 | 0 | Normal | Details |
| | Availability | New connections per second and occupied ports | 8 | 0 | Normal | Details |
| | System Performance | CPU utilization and memory usage | 4 | 0 | Normal | Details |
| Beaver | System Performance | CPU utilization, QPS, IOPS, disk usage, and memory usage of a database instance | 5 | 0 | Normal | Details |
| YundunWaf | System Performance | CPU utilization, QPS, IOPS, disk usage, and memory usage of a database instance | 5 | 0 | Normal | Details |
| SOC | System Performance | CPU utilization, QPS, IOPS, disk usage, and memory usage of a database instance | 5 | 0 | Normal | Details |
| Newsoc | System Performance | CPU utilization, QPS, IOPS, disk usage, and memory usage of a database instance | 5 | 0 | Normal | Details |
| Audit | System Performance | CPU utilization, QPS, IOPS, disk usage, and memory usage of a database instance | 5 | 0 | Normal | Details |

In the upper-left corner of the **Overall System Monitoring** tab, you can view the overall information about security services.

- **Services**: the total number of security services monitored in Apsara Stack Security Center.
- **Monitoring Items**: the total number of metrics.
- **Abnormal Items**: the number of metrics whose status is abnormal.

In the list of monitored security services, you can view the following information.

| Parameter | Description |
|-------------------------|--|
| Services | The security service monitored in Apsara Stack Security Center. |
| Metrics | The monitoring indicator for the monitored security service. |
| Description | The description of the monitoring indicator. |
| Monitoring Items | The total number of metrics that belong to the monitoring indicator. |

| Parameter | Description |
|----------------|--|
| Abnormal Items | <p>The number of metrics whose status is abnormal, and the number of metrics at each urgency level.</p> <p>The urgency levels are indicated by different colors. The red color indicates a critical exception, the orange color indicates an important exception, and the blue color indicates a moderate exception.</p> <div style="display: flex; gap: 5px;"> P1: Critical P2: Major P3: Minor </div> |
| State | The status of the monitoring indicator. |

- Click **Details** in the **Actions** column of a monitoring indicator. In the **Monitoring details** panel, view the details of metrics that belong to the monitoring indicator.

Monitoring details:Aegis-System Performance ✕

5

Total Monitoring Items

5

Normal Items

0

Abnormal Items

All

All

| Monitoring Items | Adjust Alert Threshold | Duration | Monitoring Level | Status | Alert Notifications | Actions |
|--------------------------------------|------------------------|------------|------------------|--------|--------------------------|---|
| MiniRDS Instance CPU Utilization (%) | 80% | 30 Minutes | 2 | Normal | <input type="checkbox"/> | Modify Threshold Handle Adjust Alert Duration |
| MiniRDS Instance QPS | -1 | 10 Minutes | 2 | Normal | <input type="checkbox"/> | Modify Threshold Handle Adjust Alert Duration |

In the upper-left corner of the **Monitoring details** panel, you can view the overall information about the metrics.

- **Total Monitoring Items**: the total number of metrics that belong to the monitoring indicator.
- **Normal Items**: the number of metrics in the normal state.
- **Abnormal Items**: the number of metrics in the abnormal state.

In the metric list, you can view the following information.

| Parameter | Description |
|------------------------|---|
| Monitoring Items | The name of the metric. |
| Adjust Alert Threshold | The threshold for the metric. If the value of the metric reaches the threshold and lasts for the specified period of time, the status of the metric becomes abnormal. |
| Duration | The period of time. If the value of the metric reaches the threshold and lasts for the specified period of time, the status of the metric becomes abnormal. |
| Monitoring Level | The urgency level displayed when the status of the metric becomes abnormal. |
| Status | The status of the metric. |

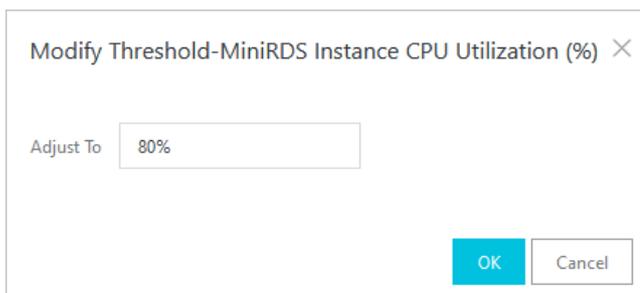
| Parameter | Description |
|---------------------|---|
| Alert Notifications | <p>The switch of the alert notification feature. You can turn on or off the switch in the Alert Notifications column based on your business requirements.</p> <p>If the  icon appears in the Alert Notifications column, a notification is sent when the status of the metric becomes abnormal.</p> |

Manage metrics

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Security Monitoring**.
3. On the **Overall System Monitoring** tab, find a monitoring indicator and click **Details** in the **Actions** column.
4. In the **Monitoring details** panel, find a metric and click **Modify Threshold**, **Handle**, or **Adjust Alert Duration** in the **Actions** column to manage the metric.

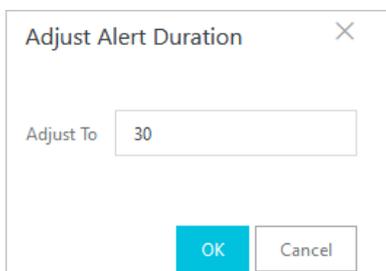
You can perform the following operations on the metric:

- **Modify Threshold:** In the **Modify Threshold** dialog box, modify the threshold and click **OK**.



After the threshold is modified, an alert is generated when the value of the metric reaches the new threshold.

- **Handle:** You can configure the status of the metric based on your business requirements.
 - If you do not want to handle the alert generated from the metric, select **Ignore** from the drop-down list. The status of the metric becomes **Ignored**.
 - After you handle the alert generated from the metric, select **Handled** from the drop-down list. The status of the metric becomes **handled**.
- **Adjust Alert Duration:** In the **Adjust Alert Duration** dialog box, modify the period of time and click **OK**.



When the value of the metric reaches the threshold and lasts for the specified period of time, the status of the metric becomes abnormal.

21.8.8.2. Alert settings

21.8.8.2.1. Configure alert contacts

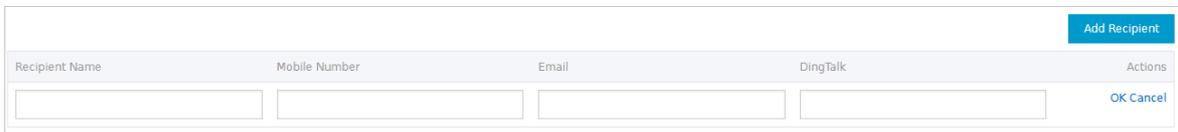
This topic describes how to configure and manage alert contacts.

Context

Apsara Stack Security sends alert notifications to alert contacts by text message, email, or DingTalk. When the detected information matches an alert rule, Apsara Stack Security sends an alert notification to the alert contacts.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Alert Settings**. On the page that appears, click the **Alert Recipient** tab.
3. Click **Add Recipient**.
4. Enter the contact information and click **OK**.



| Recipient Name | Mobile Number | Email | DingTalk | Actions |
|----------------------|----------------------|----------------------|----------------------|---|
| <input type="text"/> | <input type="text"/> | <input type="text"/> | <input type="text"/> | <input type="button" value="OK"/> <input type="button" value="Cancel"/> |

5. **Manage alert contacts.**
In the contact list, find a contact and click **Edit** in the Actions column to edit the contact information.

21.8.8.2.2. Configure alert notifications

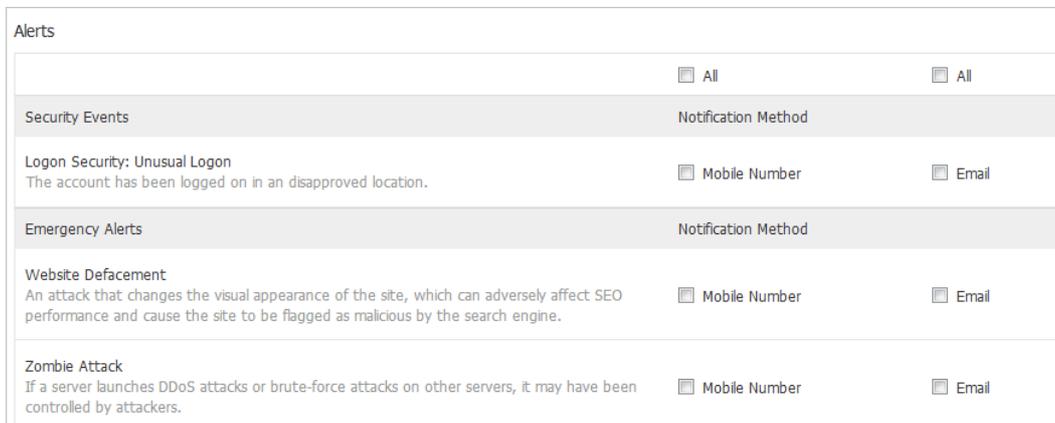
This topic describes how to configure the alert notification method for security events on tenants or platforms.

Context

In the **Alerts** section, security administrators can configure the alert notification method for security events. When a security event occurs, the system notifies the alert contacts by email, text message, or DingTalk. For more information about how to configure alert contacts, see [Set alert recipients](#).

Alerts on tenants

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Alert Settings**. On the page that appears, click the **Tenant Alerts** tab.
3. In the **Alerts** section, select notification methods for each security event.



| Alerts | | All | All |
|--|--|--|--------------------------------|
| Security Events | | Notification Method | |
| Logon Security: Unusual Logon The account has been logged on in an disapproved location. | | <input type="checkbox"/> Mobile Number | <input type="checkbox"/> Email |
| Emergency Alerts | | Notification Method | |
| Website Defacement An attack that changes the visual appearance of the site, which can adversely affect SEO performance and cause the site to be flagged as malicious by the search engine. | | <input type="checkbox"/> Mobile Number | <input type="checkbox"/> Email |
| Zombie Attack If a server launches DDoS attacks or brute-force attacks on other servers, it may have been controlled by attackers. | | <input type="checkbox"/> Mobile Number | <input type="checkbox"/> Email |

4. Click **Confirm**.

Alerts on the platform

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Alert Settings**. On the page that appears, click the **Platform Alerts** tab.
3. In the **Alerts** section, select notification methods for each security event.
4. Click **Confirm**.

21.8.8.3. Updates

21.8.8.3.1. Overview of the system updates feature

The system updates feature allows you to manually or automatically update the Apsara Stack Security and rule libraries for up-to-date protection.

The supported package import method depends on the Apsara Stack network environment.

- If Apsara Stack is connected to the Internet, you can choose **Automatically Download Update Packages**.
- If Apsara Stack is not connected to the Internet, you can choose **Manually Import Update Packages**.

The following table lists the update statuses of a rule library.

Update statuses of a rule library

| Status | Description |
|---------------|--|
| To Be Updated | Indicates that a new version of the rule library is available for update. |
| Updating | Indicates that the rule library is being downloaded from Alibaba Cloud for update. |
| Updated | Indicates that the rule library has been updated. |
| Update Failed | Indicates that the rule library failed to be updated. |

21.8.8.3.2. Enable automatic update check and update rule libraries

This topic describes how to enable automatic download of update packages and update rule libraries.

Context

If the Apsara Stack environment can connect to the Internet, you can enable automatic download of update packages to update the rule libraries.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Updates**.
3. Turn on **Auto Update Configuration** to enable automatic download of update packages.
After this switch is turned on, the system automatically downloads update packages on a regular basis.
4. Update a rule library.
You can update one or more rule libraries at a time.

- Update multiple rule libraries at a time
 - Click **Batch Update** in the upper-right corner to update all rule libraries.
- Update a single rule library
 - a. Click the tab of the rule type that you want to update. For example, click **Server Security**.
 - b. In the **Actions** column, click **Update**.

21.8.8.3.3. Manually import an update package and update your service

This topic describes how to manually import an update package and update your service.

Prerequisites

The security administrator has obtained the offline update package.

Context

If the Apsara Stack environment cannot connect to the Internet, you can update a rule library after you import an offline update package.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Updates**.
3. Manually import an update package.
 - i. Click **Import Update Package** next to **Manual Update** in the upper-left corner.
 - ii. In the **Import Update Package** dialog box, click **Browse** to select an offline update package that is downloaded to your on-premises server.
 - iii. Click **Confirm**.

4. Update a rule library.

You can update one or more rule libraries at a time.

- Update multiple rule libraries at a time
 - Click **Batch Update** in the upper-right corner to update all rule libraries.
- Update a single rule library
 - a. Click the tab of the rule type that you want to update. For example, click **Server Security**.
 - b. Click **Update** in the **Actions** column.

21.8.8.3.4. Roll back a rule library

This topic describes how to roll a rule library back to a previous version.

Context

If an error occurs with an updated rule library, you can roll the library back to a previous version to avoid service interruption.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations >**

Updates.

3. Click the tab of the rule library that you want to roll back. Example: **Server Security**
4. In the Actions column for the rule library, choose **More > Roll Back**.
5. In the **Version Rollback** dialog box, click **Confirm**.

21.8.8.3.5. View the update history of a rule library

This topic describes how to view the update history of a rule library.

Context

You can view the update history of a rule library. If an error occurs with the latest version, you can locate the issue and roll back the rule library to an earlier version.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Updates**.
3. Click the tab of the specific rule library. Example: **Server Security**.
4. In the **Actions** column of a rule library, click **History**.

On the **Previous Updates** page, you can view the update history of the rule library. Click **Details** to view the details of an update package.

21.8.8.4. Global configuration

21.8.8.4.1. Set CIDR blocks for traffic monitoring

21.8.8.4.1.1. Add a CIDR block for traffic monitoring

This topic describes how to add a Classless Inter-Domain Routing (CIDR) block for traffic monitoring. Network Traffic Monitoring System of Apsara Stack Security monitors the traffic of a specific CIDR block.

Context

CIDR blocks are configured for Network Traffic Monitoring System. Security administrators can change the CIDR blocks for monitoring based on business requirements. The settings of CIDR blocks apply only to a data center that is deployed in the region to which the specific CIDR block belongs.

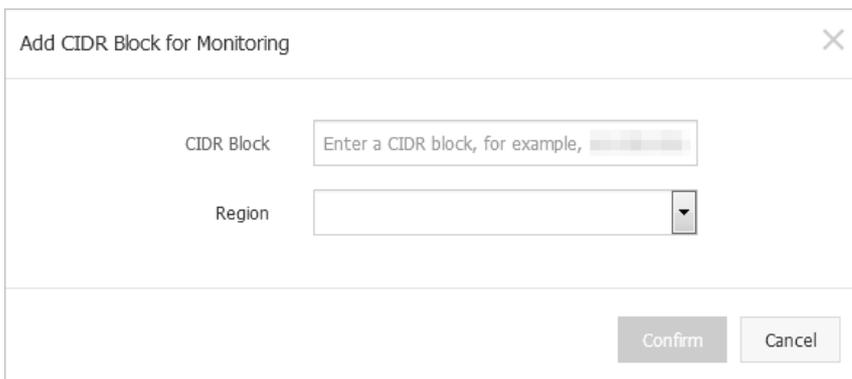
Note

Changes to CIDR block settings immediately take effect without the intervention of security administrators.

If you add the same CIDR block on the traffic collection CIDR block setting page and region setting page, make sure that you select the same region on both pages.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Global Settings**. On the page that appears, click the **Traffic Collection IP Range** tab.
3. Click **Add**.
4. In the **Add CIDR Block for Monitoring** dialog box, configure parameters.



- **CIDR Block:** Enter a CIDR block for traffic monitoring.

Note Take note that the CIDR block that you entered must be valid and unique.

- **Region:** Select the region of the data center.

5. Click **Confirm**.

21.8.8.4.1.2. Manage CIDR blocks for traffic monitoring

This topic describes how to modify or delete Classless Inter-Domain Routing (CIDR) blocks for traffic monitoring.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Global Settings**. On the page that appears, click the **Traffic Collection IP Range** tab.
3. Select a region, enter the CIDR block that you want to query, and then click **Search**.

View the information about the CIDR block for traffic monitoring and the region in the search result.

4. In the **Actions** column, manage a CIDR block for traffic monitoring.
 - **Modify the CIDR block for traffic monitoring**
Click **Modify** to modify the region of the CIDR block for traffic monitoring.
 - **Delete the CIDR block for traffic monitoring**
Click **Delete** to delete the CIDR block for traffic monitoring.

21.8.8.4.2. Region settings

21.8.8.4.2.1. Add a CIDR block for a region

This topic describes how to add Classless Inter-Domain Routing (CIDR) blocks for regions that are detected and reported by using Server Guard.

Context

Region settings are used for region detection of the Server Guard agent. Server Guard servers automatically detect and match the regions of servers based on the IP address information that is reported by the Server Guard agent.

Note You can change the region of a CIDR block. After the region is modified, you must also modify the region for all assets in the CIDR block on the **Asset Overview** page.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Global Settings**. On the page that appears, click the **Region** tab.
3. Click **Add**.
4. In the **Add CIDR Block** dialog box, configure parameters.

- **CIDR Block:** Enter a CIDR block for the region.

Note Enter a valid CIDR block. You cannot enter a CIDR block that has been configured for the region.

- **Region:** Select a region.
5. Click **Confirm**.

21.8.8.4.2.2. Manage CIDR blocks for a region

This topic describes how to modify or delete Classless Inter-Domain Routing (CIDR) blocks for a region.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Global Settings**. On the page that appears, click the **Region** tab.
3. Select a region, enter the CIDR block that you want to modify or delete, and then click **Search**.
You can view the information about the CIDR block for the region in the search result.
4. In the **Actions** column, click **Modify** or **Delete** to manage the CIDR block for the region.
 - **Modify** the CIDR block for the region
Click **Modify** to modify the CIDR block for the region.
 - **Delete** the CIDR block for the region
Click **Delete** to delete the CIDR block for the region.

21.8.8.4.3. Configure whitelists

This topic describes how to configure the whitelist for the feature that blocks brute-force attacks in Server Guard and the following whitelists in Threat Detection Service (TDS). The whitelists contain IP addresses allowed by server brute-force attack blocking, IP addresses allowed by application attack blocking, and IP addresses allowed by web attack blocking.

Context

If a normal request is regarded as an attack by the attack blocking feature of TDS or the unusual logon detection

feature of Server Guard, you can add the source IP address of the request to a whitelist to avoid further false positives.

 **Note** Make sure that the IP addresses in the whitelist are trusted.

Procedure

1. Log on to **Apsara Stack Security Center**.
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Global Settings**. On the page that appears, click the **Whitelist** tab.
3. Click **Add**.
4. In the **Add to Whitelist** dialog box, configure the parameters.

| Parameter | Description |
|-----------|---|
| Source IP | Enter a source IP address or Classless Inter-Domain Routing (CIDR) block. |
| Username | Enter the name of the user who creates the whitelist. |
| Type | <ul style="list-style-type: none"> ◦ Brute-Force Attack Blocking Whitelist: Server Guard does not generate alerts for brute-force attacks or unusual logons from the IP addresses that are contained in this whitelist. ◦ BWAF Whitelist: The attack blocking feature does not generate alerts for the web attacks from the IP addresses that are contained in this whitelist. ◦ Servers with Brute-Force Attack Permissions: The attack blocking feature does not generate alerts for the brute-force attacks from the IP addresses that are contained in this whitelist. ◦ IPs with Application Attack Permissions: The traffic from the IP addresses in this whitelist is not detected as suspicious application attack traffic. |

5. Click **OK**.
If you want to delete an existing whitelist, click **Delete** in the Actions column. In the **Delete Whitelist** message, click **Confirm**.

21.8.8.4.4. Configure attack blocking policies

This topic describes how to enable web attack blocking and brute-force attack blocking.

Context

The attack blocking features protect your servers against web attacks and brute-force attacks.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Global Settings > Region**.
3. Turn on or off the switches in the Actions column to enable or disable **Web Attack Blocking** or **Brute-Force Attack Blocking**.

| Category | Status | Description | Actions |
|-----------------------------|----------|---|---|
| Web Attack Blocking | Disabled |  Web attack blocking is disabled. Only the warning function is provided. |  |
| Brute-Force Attack Blocking | Disabled |  Brute-Force attack blocking is disabled. Only the warning function is provided. |  |

Note

In the Actions column, a red switch indicates a disabled feature and a green switch indicates an enabled feature.

After you disable the blocking feature for an attack type, Apsara Stack Security Center generates only alerts for this type of attacks.

21.8.8.4.5. Block IP addresses

This topic describes how to manually block requests for a specific IP address with a few clicks.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Global Settings**.
3. On the **Global Settings** page, click the **Block IP Addresses** tab.
4. In the upper-right corner of the tab, click **Add**.
5. In the **Add** dialog box, configure parameters.

Add
✕

Source IP

Destination IP

Destination Port

Blocking Duration

--Select--

Type

Blacklist

Note: The whitelist mechanism has precedence over the blacklist.

Confirm

Cancel

| Parameter | Description |
|-------------------|--|
| IP protocol | Specify the protocol type. Valid values: IPv4 and IPv6 . |
| Source IP | Enter the source IP address that you want to block. |
| Destination IP | Enter the destination IP address that you want to block. |
| Destination Port | Enter the destination port that is used with the specified destination IP address. |
| Blocking Duration | Select a time range during which you want to block requests. Valid values: 1 Day , 7 Days , and 30 Days . |
| Type | Select the blocking mode. Valid values: Whitelist and Blacklist . |
| Remarks | Enter the reason for blocking. |

6. Click **Confirm**.

21.8.8.4.6. Configure custom IP addresses and locations

21.8.8.4.6.1. Add custom IP addresses and locations

This topic describes how to add custom IP addresses and locations. You can customize internal IP addresses based on your network plan. After you configure the internal IP addresses, IP addresses from the public address library do not match the addresses outside China.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Global Settings**. On the page that appears, click the **Custom IP Location** tab.
3. Click **Add**.

If you want to add multiple IP addresses and locations at a time, click **Batch Upload (.txt)**. This can be used to import multiple IP addresses and locations as a template.

4. In the **Add** dialog box, configure parameters.
5. Click **OK**.

21.8.8.4.6.2. Manage custom IP addresses and locations

This topic describes how to modify and delete custom IP addresses and locations.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Global Settings**. On the page that appears, click the **Custom IP Location** tab.
3. In the **Actions** column, manage custom IP addresses and locations.
 - To modify a custom IP address and a location:
Click **Modify**. In the **Modify** dialog box, modify the custom geographic location.
 - To delete a custom IP address and a location:
Click **Delete**. In the **Delete** message, click **OK**.

21.8.8.5. System Monitoring

21.8.8.5.1. Configure CIDR blocks for traffic redirection in Cloud Firewall

Before you can use Cloud Firewall, you must configure CIDR blocks for traffic redirection.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > System Monitoring**.
3. On the **ICFW** tab in the **Service Monitoring** section of the **System Monitoring** page, click **Add** to the right of **Traffic Forwarding Configuration - Service CIDR Block**.
4. In the **Add** dialog box, configure the **Device Cluster ID**, **CIDR Block for Traffic Diversion**, **Type**, and **Service Description** parameters.
5. Click **OK**.
After you configure the CIDR blocks, you can view related information in **Traffic Forwarding Configuration - Service CIDR Block**, **Service Check Status**, and **Interface Status**.

21.8.8.6. Inspect services

This topic describes how to inspect services such as Cloud Firewall and Network Traffic Monitoring System in Apsara Stack Security Center. You can monitor the status and features of the services.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > System Monitoring**.
3. In the **System Inspection** section of the **Network Security** tab, inspect the services in the inspection list.
To inspect a single service or multiple services at a time, perform the following operations:

- Inspect multiple services at a time: In the **System Inspection** section, click **One-click Inspection** to inspect all services in the inspection list.
- Inspect a single service: In the **System Inspection** section, click **Inspect Now** in the **Actions** column of the service that you want to inspect.

After the services are inspected, the status of the services changes to **Complete** in the **Inspection Status** column.

4. View the inspection results.

You can view the following information about a service:

- In the inspection list, view the service name, last inspection time, number of inspection items, number of inspection items whose status is normal, number of inspection items whose status is abnormal, and inspection status.
- Click **Details** in the **Actions** column of a service. In the **Inspection Result Details** panel, view the number of inspection items whose status is normal, number of inspection items whose status is abnormal, and details of each item.
- Click **Download** in the **Actions** column of a service. Download the inspection results to your computer as prompted for backup and reference.

21.8.8.7. Remote operations

21.8.8.7.1. Enable remote operations

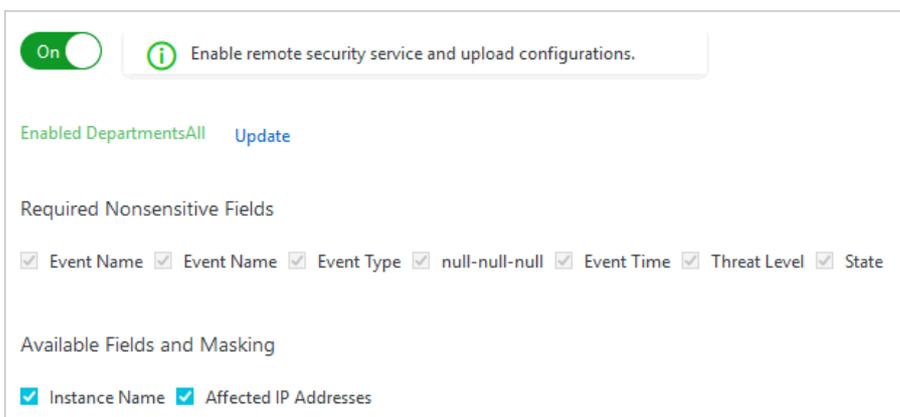
This topic describes how to enable remote operations.

Context

The **remote operations** feature provides remote security operations and rules operations.

Procedure

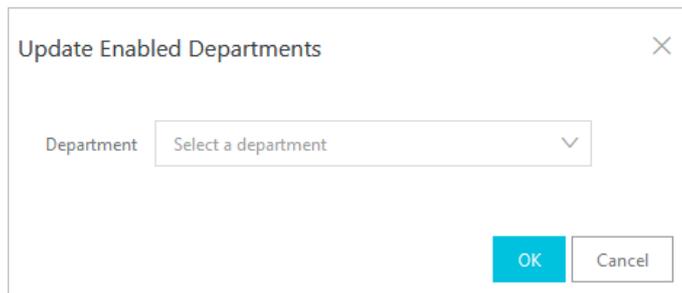
1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Remote O&M**.
3. On the **Remote O&M** page, turn on the switch to enable remote operations.



Note If the switch turns green, the feature is enabled.

4. Select the departments for which you want to enable remote operations.
 - i. Click **Update** to the right of **Enabled Departments**.

- ii. In the **Update Enabled Departments** dialog box, click the **Department** drop-down list to select the departments for which you want to enable remote operations.



- iii. Click **OK**.

Note After you perform these operations, the security logs of departments specified by **Enabled Departments** are encrypted and uploaded to Apsara Stack Security Center.

5. Select fields to encrypt and upload for remote operations.
 - o If you select **Required Nonsensitive Fields**, the system encrypts and uploads the data.
 - o If you select **Available Fields and Masking**, the system masks the data before it encrypts and uploads the data.

21.8.8.8. Account management

21.8.8.8.1. View and modify an Apsara Stack tenant account

This topic describes how to view and modify the information about your Apsara Stack tenant account that is bound to the system.

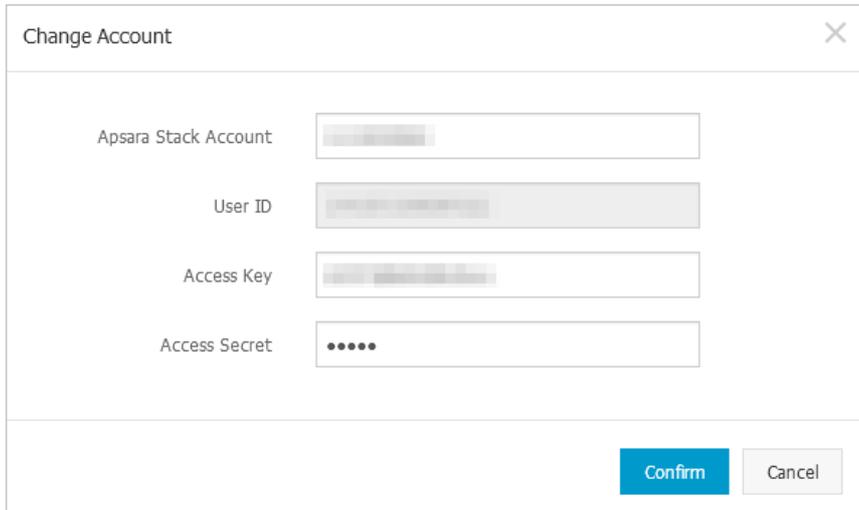
Context

Note All assets in Apsara Stack Security are bound to your Apsara Stack tenant account. You can modify the account information. Proceed with caution.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Accounts**. On the page that appears, click the **Apsara Stack Account** tab.
3. Modify the information about your Apsara Stack tenant account.
 - i. In the Actions column, click **Modify**.

ii. In the **Change Account** dialog box, modify the account information.

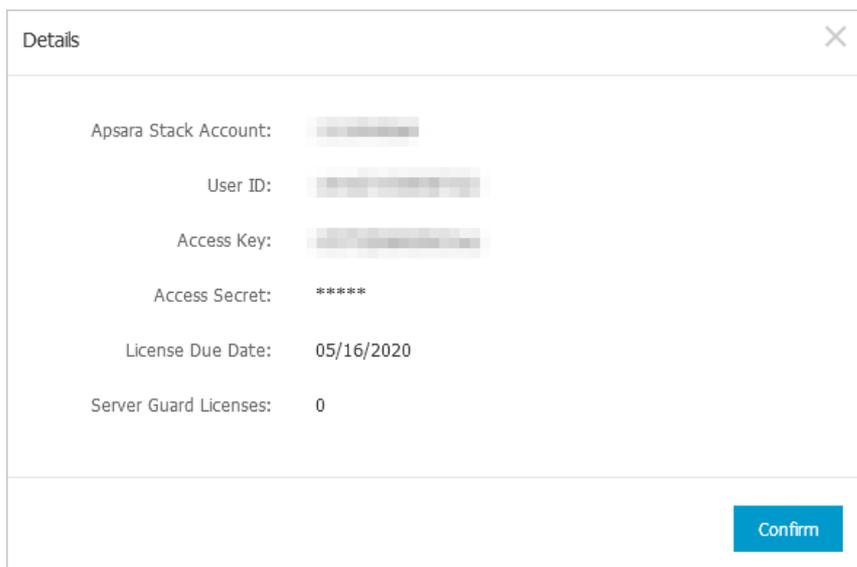


The 'Change Account' dialog box contains four input fields: 'Apsara Stack Account', 'User ID', 'Access Key', and 'Access Secret'. The 'Access Secret' field is masked with dots. At the bottom right, there are 'Confirm' and 'Cancel' buttons.

iii. Click **Confirm**.

4. View the details of your Apsara Stack tenant account.

Click **Details** to view the details of your Apsara Stack tenant account.



The 'Details' dialog box displays the following information: Apsara Stack Account, User ID, Access Key, Access Secret (masked with asterisks), License Due Date (05/16/2020), and Server Guard Licenses (0). A 'Confirm' button is located at the bottom right.

21.8.8.8.2. Add an Alibaba Cloud account

This topic describes how to add an Alibaba Cloud account in Apsara Stack Security Center. After you add the account, you can use features in a hybrid cloud.

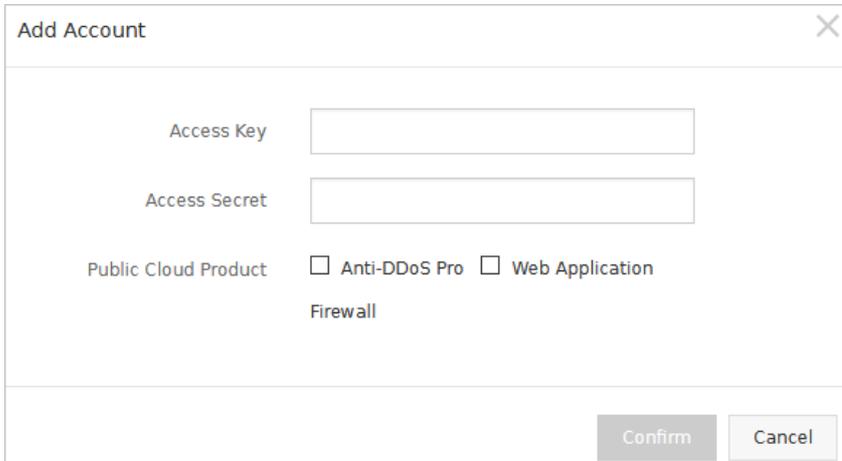
Context

After you add an Alibaba Cloud account in Apsara Stack Security Center, you can manage the Anti-DDoS Pro, Anti-DDoS Premium, and Web Application Firewall (WAF) instances that belong to the Alibaba Cloud account in Apsara Stack Security Center. This way, you can use features in a hybrid cloud.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Security Operations Center (SOC) > System Configurations > Accounts**. On the page that appears, click the **Public Cloud Account** tab.

3. On the **Feature Integration** tab, click **Add**.
4. In the **Add Account** dialog box, enter the information about your Alibaba Cloud account and select Alibaba Cloud services to use.



- Enter the **AccessKey ID** and **AccessKey secret** of your Alibaba Cloud account.
 - Select Alibaba Cloud services to use. Valid values: **Anti-DDoS Pro**, **Web Application Firewall**, and both.
5. Click **Confirm**.

Result

After the account is added, it is displayed on the **Public Cloud Account** tab. To modify or delete the account, you can click **Modify** or **Delete** in the Actions column.

21.9. Optional security products

21.9.1. Anti-DDoS settings

21.9.1.1. Overview

In Distributed Denial of Service (DDoS) attacks, attackers exploit the client-server model to combine multiple computers into a platform that can launch attacks on one or more targets. This greatly increases the threat of attacks.

Common DDoS attack types include:

- **Network-layer attacks:** A typical example is UDP reflection attacks, such as NTP flood. These attacks use heavy traffic to congest the network of the victim, disabling proper responses to user requests.
- **Transport-layer attacks:** Typical examples include SYN flood and connection flood. These attacks consume a large number of connection resources of a server to cause denial of service.
- **Session-layer attacks:** A typical example is SSL flood. These attacks consume the SSL session resources of a server to cause denial of service.
- **Application-layer attacks:** Typical attack types include DNS flood, HTTP flood, and game zombie attacks. These attacks consume a large amount of application processing resources of a server to cause denial of service.

Apsara Stack Security can redirect, scrub, and re-inject attack traffic to protect your server against DDoS attacks and ensure normal business operations.

 **Note** Apsara Stack Security cannot scrub the traffic between internal networks.

21.9.1.2. View and configure DDoS mitigation policies

This topic describes how to view and configure DDoS mitigation policies. Anti-DDoS provides default DDoS mitigation policies and DDoS traffic scrubbing policies.

Context

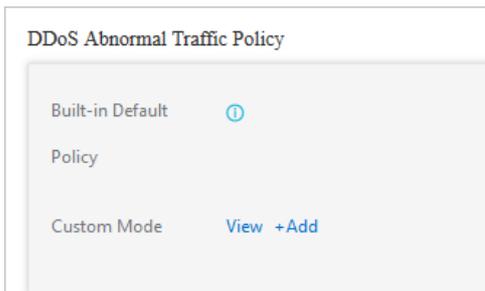
After an alert threshold of DDoS traffic for an IP address is set, an alert is triggered when traffic to the IP address reaches the threshold. The alert threshold for an IP address must be set based on the traffic volume. An abnormally large traffic volume may indicate DDoS attacks. We recommend that you set an alert threshold to a value slightly higher than the peak traffic volume.

Apsara Stack Security supports a global alert threshold, alert threshold for a specific CIDR block, and alert threshold for an IP address.

- Global alert threshold: You cannot set a global alert threshold. It is automatically set when Apsara Stack Security is initialized.
- Alert threshold for a specific CIDR block: You can set an alert threshold for a specific CIDR block based on its traffic volume. CIDR block-specific alert thresholds allow you to control the traffic to each CIDR block.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Network Security > Policy Configuration > DDoS Defense Policy**.
3. View and customize DDoS mitigation policies.



| Operation | Description |
|-------------------------|--|
| View the default policy | Move the pointer over the icon next to Built-in Default Policy in the preceding figure to view the default DDoS mitigation policy. |
| Customize a policy | Click View to view CIDR block-specific policies. Click +Add to customize a DDoS mitigation policy for a CIDR block. |

To customize a policy for a CIDR block, perform the following steps:

- i. Click **+Add** next to **Custom Mode**.

- ii. In the **Set Thresholds for Alerts** dialog box, configure the parameters.

| Parameter | Description |
|---------------------|--|
| CIDR Block | The CIDR block for which the alert thresholds are used. |
| Bandwidth Threshold | The alert threshold for bandwidth usage in a data center. When the sum of inbound and outbound traffic reaches this threshold, DDoS detection is triggered. Set this parameter to a value slightly higher than the peak traffic volume. We recommend that you set the value to 100 or higher. Unit: Mbit/s. |
| Packets Threshold | The alert threshold for the packet rate in a data center. When the sum of inbound and outbound packet rates reaches this threshold, DDoS detection is triggered. Set this parameter to a value slightly higher than the peak packet rate. We recommend that you set the value to 20000 or higher. Unit: packets per second (PPS). |

- iii. Click **OK**.

4. In the **DDoS Scrubbing Defense Strategy** section, click **View** to view DDoS traffic scrubbing policies.



21.9.1.3. View DDoS events

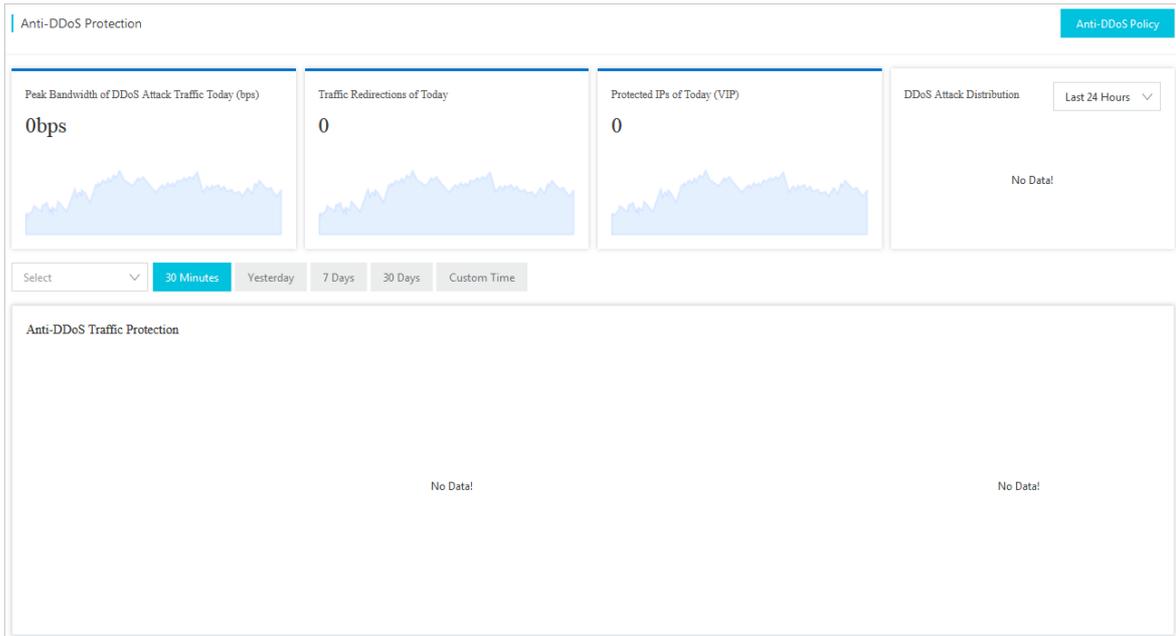
This topic describes how to view DDoS events.

Context

During or after traffic scrubbing, Apsara Stack Security reports security events to Apsara Stack Security Center.

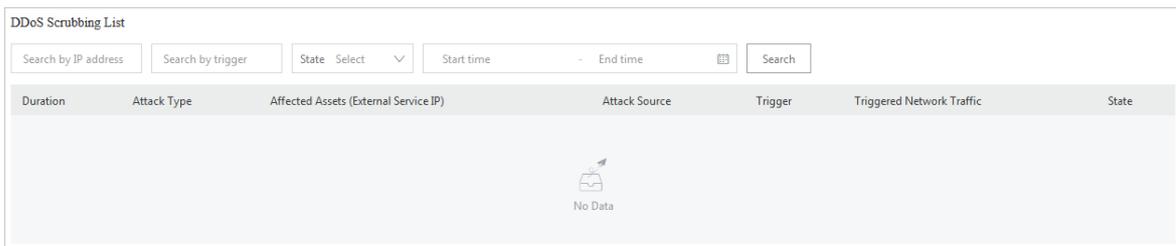
Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Network Security > Network Protection**.
3. View anti-DDoS statistics.



4. (Optional) In the **DDoS Scrubbing List** section, specify search conditions and click **Search**.

Note If you need to view all traffic scrubbing events, skip this step.



| Search condition | Description |
|-------------------------|--|
| IP address | The IP address that was under a DDoS attack. |
| Trigger | The metric that exceeds the configured alert threshold in the DDoS attack traffic. |
| State | <ul style="list-style-type: none"> Scrubbing: indicates that traffic scrubbing is in progress. Scrubbing Complete: indicates that traffic scrubbing is complete. |
| Start time and End time | The start time and end time of DDoS traffic scrubbing. |

5. In the DDoS Scrubbing List section, view details about DDoS traffic scrubbing events.

21.9.2. Cloud Firewall

21.9.2.1. Policy configuration

21.9.2.1.1. Synchronize assets for the Internet firewall

If new IP addresses are not in the IP address list of the Internet firewall, you can manually synchronize the IP address assets for the Internet firewall. This topic describes how to manually synchronize assets for the Internet firewall.

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Network Security > Policy Configuration > Firewall Switch Policy**.
3. Click the **Internet Firewall** tab. In the upper-right corner of the page, click **Update Assets**.
4. In the **Update Assets** message, click **OK**.
After the assets are synchronized, the system updates the IP address list.

21.9.2.1.2. Create a VPC firewall

A virtual private cloud (VPC) firewall is a distributed firewall that can detect and control traffic between VPCs. Cloud Firewall can be used to analyze and control traffic between two VPCs only after a VPC firewall is created and enabled. This topic describes how to create a VPC firewall.

Context

A VPC firewall can be created only between two VPCs that are connected. VPCs can be connected by using Express Connect or Cloud Enterprise Network (CEN).

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Network Security > Policy Configuration > Firewall Switch Policy**.
3. On the **Firewall Switches** page, click the **VPC-VPC** tab. Find the required instance and click **Create** in the Actions column.

| Parameter | Description |
|-------------------------------|---|
| Instance Name | Enter a name for the VPC firewall. We recommend that you enter a unique and informative name that indicates specific business to help identify the VPC firewall. The name can contain letters, digits, underscores (_), and hyphens (-). However, the name cannot contain only digits. |
| Route Table | When you create a VPC, the system automatically creates a default route table. You can add system routes to the route table to manage VPC traffic. VPC allows you to create multiple route tables. When you create a VPC firewall in the Cloud Firewall console, Cloud Firewall automatically reads your VPC route tables. Express Connect supports multiple route tables. If you create a VPC firewall for Express Connect, you can view multiple VPC route tables and can select specific route tables for protection. |
| Destination CIDR Block | After you select a route table from the Route Table drop-down list, the default destination Classless Inter-Domain Routing (CIDR) block of the route table appears in the Destination CIDR Block field. If you want to protect traffic destined for other CIDR blocks, you can manually modify the destination CIDR block. You can enter multiple CIDR blocks and separate them with commas (,). |
| Peer Route Table | Confirm the region and name of the peer VPC, and select the route table for protection. |

| Parameter | Description |
|-------------------------------------|---|
| Peer Destination CIDR Blocks | Confirm the destination CIDR block of the peer VPC. |
| Firewall Mode | <ul style="list-style-type: none"> ◦ Test: In this mode, the VPC firewall tests the health status of CIDR blocks or IP addresses to ensure normal links. ◦ Active: In this mode, the VPC firewall redirects and protects traffic. ◦ Bypass: In this mode, the VPC firewall does not redirect traffic. If a self-test or a health test fails, the VPC firewall automatically changes to the Bypass mode. <p>When you create a VPC firewall, you can set Firewall Mode only to Test. After the firewall is created, you can change its mode to Active or Bypass. You cannot directly change the mode from Bypass to Active. You must change the mode from Bypass to Test first, and then to Active.</p> <p>In the Health Test section, enter 32-bit test IP addresses that belong to the local and peer VPCs.</p> |
| IPS Mode | <p>Select the working mode of the intrusion prevention system (IPS). Valid values:</p> <ul style="list-style-type: none"> ◦ Monitoring Mode: If Cloud Firewall detects malicious traffic, it monitors the traffic and sends alerts. ◦ Traffic Control Mode: Cloud Firewall intercepts malicious traffic and blocks intrusion attempts. |
| IPS Capabilities | <p>Select the intrusion prevention policies that you want to enable. Valid values:</p> <ul style="list-style-type: none"> ◦ Basic Policies: This feature provides basic intrusion prevention capabilities, such as protection against brute-force attacks and attacks that exploit command execution vulnerabilities. This feature also allows you to manage and control the connections from infected hosts to a command and control (C&C) server. ◦ Virtual Patches: This feature defends against the most common and high-risk application vulnerabilities in real time. |
| Enable VPC Firewall | If you turn on Enable VPC Firewall, the VPC firewall is automatically enabled after it is created. If you do not require the VPC firewall to be automatically enabled, turn the switch off. |

4. Click **Submit**.

If Firewall Status of the VPC firewall changes to **Enabled**, the VPC firewall takes effect.

21.9.2.1.3. Create an IDC-VPC firewall

An IDC-VPC firewall can detect the traffic between a data center and a virtual private cloud (VPC). Cloud Firewall allows you to control traffic through IDC-VPC firewalls. This topic describes how to create an IDC-VPC firewall.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Network Security > Policy Configuration > Firewall Switch Policy**.
3. On the **Firewall Switches** page, click the **IDC-VPC** tab. In the upper-right corner of the tab, click **Create**.

| Parameter | Description |
|---------------------------------------|--|
| Instance Name | <p>The name of the IDC-VPC firewall. We recommend that you enter a unique and informative name that indicates specific business to help identify the IDC-VPC firewall.</p> <p>The name can contain letters, digits, underscores (_), and hyphens (-). However, the name cannot contain only digits.</p> |
| VPC Instance | <p>The ID of the VPC. The ID is the unique identifier of the VPC. Cloud Firewall automatically synchronizes the VPCs that are connected to the data center.</p> |
| VPC Route Table | <p>When you create a VPC, the system automatically creates a default route table. You can add system routes to the route table to manage VPC traffic. VPC allows you to create multiple route tables.</p> <p>When you create a VPC firewall in the Cloud Firewall console, Cloud Firewall automatically reads your VPC route tables. Express Connect supports multiple route tables. If you create a VPC firewall for Express Connect, you can view multiple VPC route tables and can select specific route tables for protection.</p> |
| VPC Destination CIDR Block | <p>After you select a route table from the VPC Route Table drop-down list, the default destination CIDR block of the route table is automatically displayed for the VPC Destination CIDR Block field. If you need to protect traffic to other CIDR blocks, you can manually modify destination CIDR blocks.</p> <p>You can add multiple CIDR blocks. Separate the CIDR blocks with commas (,).</p> |
| IDC Express Connect Circuit (Primary) | <p>The ID of the Express Connect circuit that you create when you connect your data center to Apsara Stack.</p> <p>When the customer edge (CE) in your data center connects to the primary and secondary vSwitches, you must specify a primary Express Connect circuit. The IDC-VPC firewall automatically synchronizes the primary Express Connect circuit that you specify. You must specify this parameter when you create the IDC-VPC firewall.</p> |

| Parameter | Description |
|---|---|
| VBR(Primary) | The virtual border router (VBR) that is bound to the primary Express Connect circuit in your data center. The VBR facilitates communication between the VPC and your data center. The IDC-VPC firewall automatically synchronizes the VBR that you specify. You must specify this parameter when you create the IDC-VPC firewall. |
| IDC Express Connect Circuit (Secondary) | The ID of the Express Connect circuit that you create when you connect your data center to Apsara Stack. You must specify a secondary Express Connect circuit when the CE in your data center connects to the primary and secondary vSwitches. The value of this parameter cannot be the same as that of the IDC Express Connect Circuit (Primary) parameter. The IDC-VPC firewall automatically synchronizes the secondary Express Connect circuit that you specify. You must specify this parameter when you create the IDC-VPC firewall. |
| VBR(Secondary) | The VBR that is bound to the secondary Express Connect circuit in your data center. The VBR facilitates communication between the VPC and your data center. The IDC-VPC firewall automatically synchronizes the VBR that you specify. You must specify this parameter when you create the IDC-VPC firewall. |
| IDC Destination CIDR Block | The destination CIDR block of the peer VPC. |
| Configure Firewall Mode | <ul style="list-style-type: none"> ◦ Test: This mode is used to test the health status of the CIDR block or the IP address to ensure that the link is normal. ◦ Active: This mode is used to redirect and protect traffic. ◦ Bypass: The firewall does not redirect traffic in this mode. If the self-check or the health test fails, the firewall is automatically changed to the Bypass mode. <p>When you create an IDC-VPC firewall, you can set Configure Firewall Mode only to Test. After the firewall is created, you can change its mode to Active or Bypass. You cannot directly change the mode from Bypass to Active. You must change the mode from Bypass to Test and then to Active.</p> <p>In the Health Test section, enter a 32-bit test IP address that is created in the local VPC.</p> |
| IPS Mode | The working mode of the intrusion prevention system (IPS). Valid values: <ul style="list-style-type: none"> ◦ Monitoring Mode: In this mode, Cloud Firewall monitors malicious traffic and generates alerts on the traffic. ◦ Traffic Control Mode: In this mode, Cloud Firewall intercepts malicious traffic and blocks intrusion attempts. |
| IPS Capabilities | The intrusion prevention policies that you want to enable. Valid values: <ul style="list-style-type: none"> ◦ Basic Policies: This feature provides basic intrusion prevention capabilities, such as protection against brute-force attacks and attacks that exploit command execution vulnerabilities. It also allows you to manage and control the connections from compromised hosts to a command and control (C&C) server. ◦ Virtual Patches: This feature protects against common high-risk application vulnerabilities in real time. |
| Enable VPC Firewall | After you turn on Enable VPC Firewall, an IDC-VPC firewall is automatically enabled after it is created. If you do not require the IDC-VPC firewall to be automatically enabled, turn off the switch. |

4. Click **Submit**.

If the status of the IDC-VPC firewall changes to **Enabled**, the IDC-VPC firewall is in effect.

21.9.2.2. Access control

21.9.2.2.1. Manage address books

This topic describes how to create and manage IP address books and port address books. You can use address books to store one or more CIDR blocks or ports.

Context

You can store frequently used IP addresses and ports in address books to facilitate configurations of the Internet firewall.

Create an IP address book

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Network Security > Access Control**.
3. On the page that appears, click the **Internet Firewall** tab.
4. In the upper-right corner, click **Manage Address Books**. In the dialog box that appears, click the **IP Address Books** tab.
5. Click **+ Create Address Book** to configure the parameters.

| Parameter | Description |
|-------------------|--|
| Address Book Type | The type of the address book. Set the value to IP Addresses . |
| Address Book Name | The name of the address book. The name must be unique. The name can contain letters, digits, and underscores (_). However, the name cannot contain only digits. |
| IP Address | The CIDR block that you want to add. Separate multiple CIDR blocks with commas (,). |
| Description | The content and use scenarios of the address book. The description must be 2 to 512 characters in length. |

6. Click **Submit**.
After the address book is created, you can view the address book on the **IP Address Books** tab. You can click **Modify** or **Delete** in the Actions column of an address book to manage the address book.

Create a port address book

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Network Security > Access Control**.
3. On the page that appears, click the **Internet Firewall** tab.
4. In the upper-right corner, click **Manage Address Books**. In the dialog box that appears, click the **Port Address Books** tab.
5. Click **+ Create Address Book** to configure the parameters.

| Parameter | Description |
|-------------------|--|
| Address Book Name | The name of the address book. The name must be unique. The name can contain letters, digits, and underscores (_). However, the name cannot contain only digits. |
| Ports | The port number or port range that you want to add. Separate multiple port numbers or ranges with commas (,). |
| Description | The content and use scenarios of the address book. The description must be 2 to 512 characters in length. |

6. Click **Submit**.
 After the address book is created, you can view the address book on the **Port Address Books** tab. You can click **Modify** or **Delete** in the Actions column of an address book to manage the address book.

21.9.2.2.2. Configure access control policies on the Internet firewall

The access control feature allows you to configure access control policies on the Internet firewall. You can configure inbound and outbound policies on the Internet Firewall tab to forbid unauthorized access between the Internet and your servers.

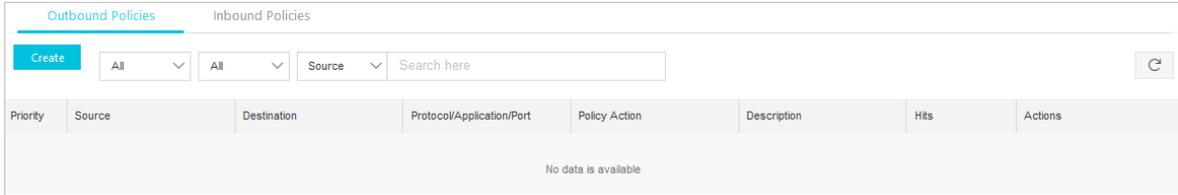
Context

The Internet firewall supports **Outbound Policies** and **Inbound Policies**.

In this topic, IP address books, port address books, and domain address books are used. For more information about how to create these address books, see [Manage address books](#).

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Network Security > Access Control**.
3. Click the **Internet Firewall** tab.
4. Click the **Outbound Policies** tab. You can also click the **Inbound Policies** tab if required.



- o **Outbound Policies** tab: You can configure policies for traffic from your internal network to the Internet.
 - o **Inbound Policies** tab: You can configure policies for traffic from the Internet to your internal network.
5. Click **Create**.

In the **Create Outbound Policy** dialog box, configure the parameters.

Create Outbound Policy ✕

Source Type : IP Address Book

Source : * ?

Destination Type : IP Address Book Domain Name Region

Destination : * ?

Protocol : * ▼

Port Type : Ports Address Book

Ports : * ?

Application : * ▼

Policy Action : * ▼

Description : *

| Parameter | Description |
|-------------|--|
| Source Type | The type of the traffic source. Valid values: IP and Address Book . <ul style="list-style-type: none"> o IP: If you select this option, enter only one CIDR block in the Source field. o Address Book: If you select this option, select a pre-configured IP address book for the Source field. An IP address book contains multiple CIDR blocks. This allows you to control traffic from multiple IP addresses. |
| Source | The source address that is allowed to access the Internet. <ul style="list-style-type: none"> o If you set Source Type to IP, enter a CIDR block. Example: 1.XX.XX.1/32. o If you set Source Type to Address Book, click Select Address Book. In the Select Address Book as Source dialog box, select an IP address book. |

| Parameter | Description |
|------------------|---|
| Destination Type | The type of the traffic destination. Valid values: IP , Address Book , Domain Name , and Region . |
| Destination | <p>The destination address that can be accessed. You must set Destination to an address on the Internet.</p> <ul style="list-style-type: none"> ◦ If you set Destination Type to IP, enter a CIDR block. Example: 1.XX.XX.1/32. ◦ If you set Destination Type to Address Book, click Select Address Book. In the Select Address Book as Destination dialog box, select an IP address book. ◦ If you set Destination Type to Domain Name, enter a domain name. Example: www.example.com. ◦ If you set Destination Type to Region, select one or more destination regions from the drop-down list. |
| Protocol | The protocol of outbound traffic. Valid values: TCP , UDP , ICMP , and ANY . If you do not know which protocol is used, select ANY . |
| Port Type | <p>The type of the port that is used for the selected protocol. Valid values: Ports and Address Book.</p> <ul style="list-style-type: none"> ◦ Ports: If you select this option, enter a port range in the Ports field. ◦ Address Book: If you select this option, select a pre-configured port address book for the Ports field. A port address book contains multiple ports. This allows you to control traffic on multiple ports. |
| Ports | <p>The ports on which you want to control traffic. Enter a port range or select a port address book based on the value of the Port Type parameter.</p> <ul style="list-style-type: none"> ◦ If you set Port Type to Ports, enter a port range. ◦ If you set Port Type to Address Book, click Select Address Book. In the Select Ports dialog box, select a port address book. |
| Application | The application for the traffic. |
| Policy Action | <p>The action on the traffic.</p> <ul style="list-style-type: none"> ◦ Allow: If outbound traffic meets the preceding conditions that you specify for the policy, the traffic is allowed. ◦ Monitor: If outbound traffic meets the preceding conditions that you specify for the policy, the traffic is recorded and allowed. ◦ Deny: If outbound traffic meets the preceding conditions that you specify for the policy, the traffic is blocked. |
| Description | <p>The description of the policy. Enter an informative description to help identify the policy.</p> <p>The description can contain digits, letters, and underscores (_).</p> |

6. Click **Submit**.

Result

After the policy is created, it is displayed in the policy list. In the **Actions** column that corresponds to the policy, you can click **Modify**, **Delete**, **Insert**, or **Move** to manage the policy.

21.9.2.2.3. Create a policy group

This topic describes how to create a policy group for an internal firewall. You can configure access control policies to forbid unauthorized access between Elastic Compute Service (ECS) instances.

Context

An internal firewall is implemented by leveraging the security group module of ECS. The access control policies that you configure on the **Internal Firewall** tab are automatically synchronized to the security group module of ECS.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Network Security > Access Control**.
3. Click the **Internal Firewall** tab.
4. In the upper-right corner, click **Create Policy Group**.
5. In the **Create Policy Group** dialog box, configure parameters based on your business requirements.

| Parameter | Description | Configuration method |
|--------------------|---|--|
| Name | The name of the policy group. The name must be 2 to 128 characters in length. | Enter an informative name to help identify the policy group. |
| VPC | The virtual private cloud (VPC) to which the policy group is applied. | Select a VPC from the VPC drop-down list.  Note You can select only one VPC. |
| Instance ID | The ID of the ECS instance in the selected VPC. | Select an ECS instance ID from the Instance ID drop-down list.  Note You can select multiple instance IDs. |
| Description | The description of the policy group. The description must be 2 to 256 characters in length. | Enter an informative description to help identify the policy group. |
| Template | The template of the policy group. | Select a template from the Template drop-down list. Valid values: <ul style="list-style-type: none"> ◦ default-accept-login: allows all inbound traffic on ports 22 and 3389. ◦ default-drop-all: blocks all traffic in the policy group. ◦ default-accept-all: allows all traffic in the policy group. |

6. Click **Submit**.

21.9.2.2.4. Configure access control policies on an internal firewall

This topic describes how to view the policy groups of an internal firewall, configure access control policies in the policy groups, and synchronize the policies to the security group module of Elastic Compute Service (ECS).

Context

On the **Internal Firewall** tab, you can view custom policy groups and security groups that are synchronized from ECS.

For more information about how to create a custom policy group, see [Create a policy group](#).

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Network Security > Access Control**.
3. Click the **Internal Firewall** tab.
4. (Optional) Specify filter conditions and click **Search** to search for the required policy group.

 **Note** If you want to view all policy groups, skip this step.

The required policy group appears in the policy group list.

5. Find the required policy group and configure an access control policy in this group.
 - i. In the **Actions** column, click **Configure Policy**.
 - ii. On the **Policies** page, click **Create Policy**.

 **Note** If you want to modify or delete an access control policy, you can perform the following steps: Click the **Inbound** or **Outbound** tab, find the policy, and then click **Modify** or **Delete** in the **Actions** column.

- iii. In the **Create Policy** dialog box, configure the parameters.

| Parameter | Description |
|----------------------|--|
| Network Type | The type of the network to which the policy is applied. Default value: Internal , which indicates that the policy is applied to an internal network. |
| Direction | The direction of traffic that is controlled by the policy. Valid values: <ul style="list-style-type: none"> ▪ Inbound: Traffic is from other ECS instances to the specified ECS instance. ▪ Outbound: Traffic is from the specified ECS instance to other ECS instances. |
| Policy Type | The action on traffic that is controlled by the policy. Valid values: <ul style="list-style-type: none"> ▪ Allow: Traffic is allowed in the internal network. ▪ Deny: Traffic is blocked in the internal network. |
| Protocol Type | The protocol of traffic that is controlled by the policy. Select a protocol from the Protocol Type drop-down list. Valid values: <ul style="list-style-type: none"> ▪ TCP ▪ UDP ▪ ICMP ▪ ANY: If you do not know which protocol is used, select ANY. |
| Port Range | The port range of traffic that is controlled by the policy. The traffic is destined for ports in this range. Enter a port range. Example: 22/22. |

| Parameter | Description |
|-------------|---|
| Priority | <p>The priority of the policy.</p> <p>Note The priority number must be an integer from 1 to 100. Different policies can have the same priority. If an Allow policy and a Deny policy have the same priority, the Deny policy takes effect. If two Allow policies have the same priority, both policies take effect.</p> |
| Source Type | <p>The type of the traffic source. Valid values:</p> <ul style="list-style-type: none"> ▪ CIDR Block: The traffic source is a CIDR block. ▪ Policy Group: The traffic source is an ECS instance in the policy group. <p>Note If you create a policy in the security group module of ECS, you cannot set Source Type to Policy Group.</p> |
| Source | <p>The source of traffic that is controlled by the policy.</p> <p>Enter a CIDR block or select a policy group based on how you set the Source Type parameter.</p> <ul style="list-style-type: none"> ▪ If you set Source Type to CIDR Block, enter only one CIDR block. ▪ If you set Source Type to Policy Group, select a policy group from the Source drop-down list. In this case, the traffic source is an ECS instance in the policy group. <p>Note You can select only one policy group from the Source drop-down list.</p> |
| Destination | <p>The destination of traffic that is controlled by the policy. Valid values:</p> <ul style="list-style-type: none"> ▪ All ECS Instances: The policy is applied to traffic destined for your ECS instances. ▪ CIDR Block: The policy is applied to traffic destined for the specified CIDR block. |
| Description | <p>The description of the policy. Enter an informative description to help identify the policy.</p> <p>The description must be 2 to 256 characters in length.</p> |

iv. Click **Submit**.

6. Find the required policy group and click **Publish** in the **Actions** column to apply the policy and synchronize it to the security group module of ECS.

21.9.2.2.5. Configure access control policies on a VPC firewall

A virtual private cloud (VPC) firewall detects and controls the traffic between two VPCs. This topic describes how to configure access control policies on a VPC firewall.

Prerequisites

VPC firewalls are not automatically created. Before you configure access control policies for VPCs, you must create and enable a VPC firewall. For more information, see [Create a VPC firewall](#).

Context

Access control policies of a VPC firewall take effect only after you enable the VPC firewall.

By default, a VPC firewall allows all traffic. If you want to control traffic between VPCs, you can configure access control policies to block traffic from untrusted sources. You can also allow traffic from trusted sources and block traffic from all other sources.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Network Security > Access Control**.
3. Click the **VPC Firewall** tab.
4. Click **Create**.
5. In the **Create VPC Firewall Policy** dialog box, configure an access control policy.

| Parameter | Description |
|-------------|--|
| Source Type | The type of the traffic source. Valid values: IP and Address Book. <ul style="list-style-type: none"> ◦ IP: If you select this option, you must enter a CIDR block in the Source field. ◦ Address Book: If you select this option, you must select a pre-configured address book for the Source field. You can add multiple CIDR blocks to an address book to simplify policy configuration. |
| Source | The source of the traffic. <div style="background-color: #e6f2ff; padding: 5px; margin: 5px 0;"> ? Note You can enter only one CIDR block. Example: 1.*.*.1/32. </div> If you set Source Type to Address Book, select a pre-configured address book. |

| Parameter | Description |
|-------------------------|---|
| Destination Type | <p>The type of the traffic destination. Valid values:</p> <ul style="list-style-type: none"> ◦ IP: If you select this option, you must enter an IP address in the Destination field. ◦ Address Book: If you select this option, you must select an address book for the Destination field. ◦ Domain Name: If you select this option, you must enter a domain name in the Destination field. Wildcard domain names are supported. Example: <i>*.aliyun.com</i>. <p> Note By default, if an HTTP header does not contain the Host field or an HTTPS request does not contain the Server Name Indication (SNI), Cloud Firewall allows the traffic.</p> |
| Destination | <p>The destination of the traffic. You can enter only one CIDR block.</p> <p>If you set Destination Type to Domain Name, enter a domain name. Wildcard domain names are supported. Example: <i>*.aliyun.com</i>.</p> |
| Protocol | <p>The protocol of the traffic. Valid values:</p> <ul style="list-style-type: none"> ◦ ANY: all protocols ◦ TCP ◦ UDP ◦ ICMP |
| Ports | <p>The port range of the traffic. The value 0/0 indicates all ports.</p> <p> Note If you set Protocol to ICMP, the port configuration does not take effect. If you set Protocol to ANY, the port configuration does not take effect in controlling ICMP traffic.</p> |
| Application | <p>The application for the traffic. Valid values:</p> <p>ANY, HTTP, HTTPS, Memcache, MongoDB, MQTT, MySQL, RDP, Redis, SMTP, SMTPS, SSH, and VNC</p> <p>If you set Protocol to TCP, multiple applications are supported. If you set Protocol to another value, only ANY is supported.</p> <p> Note Cloud Firewall identifies applications based on packet characteristics, instead of port numbers. If Cloud Firewall fails to identify the application for a packet, it allows the packet.</p> |
| Policy Action | <p>The action on the traffic. Valid values:</p> <ul style="list-style-type: none"> ◦ Allow: If traffic meets the preceding conditions that you specify for the policy, the traffic is allowed. ◦ Deny: If traffic meets the preceding conditions that you specify for the policy, the traffic is blocked. ◦ Monitor: If traffic meets the preceding conditions that you specify for the policy, the traffic is recorded and allowed. After you observe the traffic for a period of time, you can change the policy action to Allow or Deny. |

| Parameter | Description |
|-------------|--|
| Description | The description of the policy. Enter an informative description to help identify the policy. |

6. Click **Submit**.

21.9.2.2.6. Configure access control policies on an IDC-VPC firewall

Cloud Firewall detects and controls access traffic between Internet data centers (IDCs) and virtual private clouds (VPCs). This topic describes how to create and manage access control policies on an IDC-VPC firewall.

Prerequisites

IDC-VPC firewalls are not automatically created. Before you configure access control policies between an IDC and a VPC, you must create and enable an IDC-VPC firewall.

Access control policies of an IDC-VPC firewall take effect only after you enable the IDC-VPC firewall.

Context

By default, an IDC-VPC firewall allows all traffic. If you want to control traffic between an IDC and a VPC, you can configure access control policies to block traffic from untrusted sources. You can also allow traffic from trusted sources and block traffic from all other sources.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Network Security > Access Control**.
3. On the page that appears, click the **IDC-VPC Firewall** tab.
4. Click **Create**.
5. In the **Create IDC-VPC Firewall Policy** dialog box, configure an access control policy.

Create VPC Firewall Policy ✕

Source Type: IP Address Book

* Source: ⓘ

Destination Type: IP Address Book Domain Name

* Destination: ⓘ

* Protocol: ▾

Port Type: Ports Address Book

* Ports: ⓘ

* Application: ▾

* Policy Action: ▾

* Description:

| Parameter | Description |
|-------------------------|---|
| Source Type | <p>The type of the traffic source. Valid values:</p> <ul style="list-style-type: none"> IP: If you select this option, you must enter an IP address or a Classless Inter-Domain Routing (CIDR) block in the Source field. Address Book: If you select this option, you must select a pre-configured address book for the Source field. <p>You can add multiple CIDR blocks to an address book to simplify policy configuration.</p> |
| Source | <p>The source of the traffic.</p> <div style="background-color: #e6f2ff; padding: 5px; border: 1px solid #d9e1f2;"> <p>ⓘ Note You can enter only one CIDR block. Example: 1.*.*.1/32.</p> </div> <p>If you set Source Type to Address Book, select a pre-configured address book.</p> |
| Destination Type | <p>The type of the traffic destination. Valid values:</p> <ul style="list-style-type: none"> IP: If you select this option, you must enter an IP address or a CIDR block in the Destination field. Address Book: If you select this option, you must select an address book for the Destination field. Domain Name: If you select this option, you must enter a domain name in the Destination field. Wildcard domain names are supported. Example: *.aliyun.com. <div style="background-color: #e6f2ff; padding: 5px; border: 1px solid #d9e1f2;"> <p>ⓘ Note By default, if an HTTP header does not contain the Host field or an HTTPS request does not contain the Server Name Indication (SNI), Cloud Firewall allows the traffic.</p> </div> |

| Parameter | Description |
|----------------------|--|
| Destination | The destination of the traffic. You can enter only one CIDR block. If you set Destination Type to Domain Name, enter a domain name. Wildcard domain names are supported. Example: <i>*.aliyun.com</i> . |
| Protocol | The protocol of the traffic. Valid values: <ul style="list-style-type: none"> ◦ ANY: all protocols ◦ TCP ◦ UDP ◦ ICMP |
| Ports | The port range of the traffic. The value 0/0 indicates all ports. <div style="background-color: #e6f2ff; padding: 5px;"> <p> Note If you set Protocol to ICMP, the port configuration does not take effect. If you set Protocol to ANY, the port configuration does not take effect in controlling ICMP traffic.</p> </div> |
| Application | The application for the traffic. Valid values: ANY, HTTP, HTTPS, Memcache, MongoDB, MQTT, MySQL, RDP, Redis, SMTP, SMTPS, SSH, SSL, and VNC If you set Protocol to TCP, multiple applications are supported. If you set Protocol to another value, only ANY is supported. <div style="background-color: #e6f2ff; padding: 5px;"> <p> Note Cloud Firewall identifies applications based on packet characteristics, instead of port numbers. If Cloud Firewall fails to identify the application for a packet, it allows the packet.</p> </div> |
| Policy Action | The action on the traffic. Valid values: <ul style="list-style-type: none"> ◦ Allow: If traffic meets the preceding conditions that you specify for the policy, the traffic is allowed. ◦ Deny: If traffic meets the preceding conditions that you specify for the policy, the traffic is blocked. ◦ Monitor: If traffic meets the preceding conditions that you specify for the policy, the traffic is recorded and allowed. After you observe the traffic for a period of time, you can change the policy action to Allow or Deny. |
| Description | The description of the policy. Enter an informative description to help identify the policy. |

6. Click **Submit**.

21.9.2.3. Intrusion prevention

21.9.2.3.1. Configure intrusion prevention policies

Cloud Firewall uses a built-in threat detection engine to defend against intrusions and common cyberattacks. It provides virtual patches to protect against vulnerabilities and intelligently block intrusion attempts.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Network Security > Policy Configuration > Intrusion Prevention Policies**.
3. Select a running mode for the threat engine.



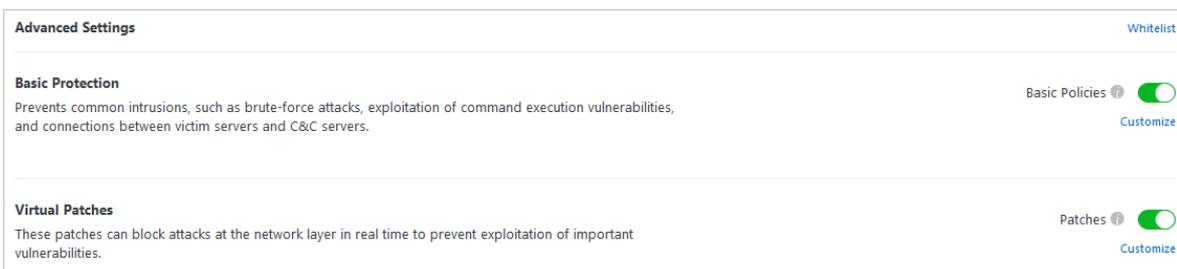
Valid values: **Monitoring Mode** and **Traffic Control Mode**.

- o **Monitoring Mode:** In this mode, the system generates alerts on malicious traffic instead of blocking the traffic.
- o **Traffic Control Mode:** In this mode, the system automatically blocks malicious traffic and generates alerts.

To configure **Traffic Control Mode**, perform the following steps:

- i. Select **Traffic Control Mode**.
- ii. In the message that appears, click **OK**.

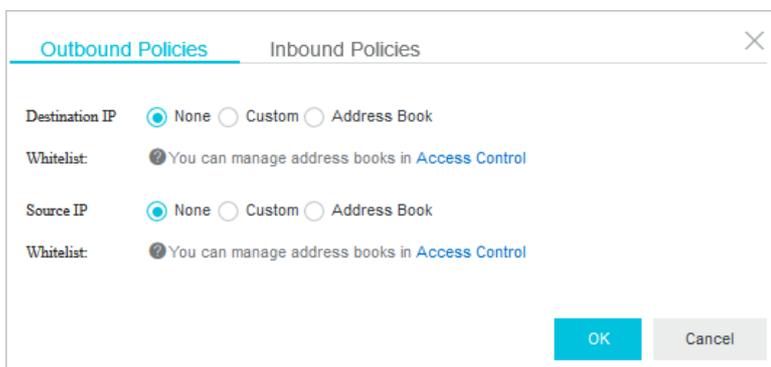
4. Configure the advanced features.



- o Cloud Firewall allows traffic from IP addresses in the whitelists that you configured.

In the **Advanced Settings** section, click **Whitelist** to add trusted IP addresses to a whitelist.

You can add trusted source IP addresses, destination IP addresses, or address books of both inbound and outbound traffic to the whitelist.



- o The basic protection feature defends against common intrusions, such as brute-force attacks and attacks that exploit command execution vulnerabilities. The basic protection feature also manages connections from compromised hosts to a command and control (C&C) server. After you enable the basic protection feature, it provides basic protection for your assets.
 - a. Turn on **Basic Policies** to enable the basic protection feature.

- b. In the Basic Protection section, click **Customize**. In the **Customize Basic Protection Policies** dialog box, customize one or more basic protection policies.

| Rule ID | Policy Name | Description | Risk Level | CVE ID | Attack Type | Victim | Default Action | Current Action |
|----------|----------------------------------|-------------|------------|--------|-------------|----------|----------------|----------------|
| 10003... | Acunetix WVS scan | ... | Medium | - | Scan | Acunetix | Block | Block |
| 10003... | Acunetix WVS scan | ... | Medium | - | Scan | Acunetix | Block | Block |
| 10003... | Acunetix RFI vulnerability sc... | ... | Medium | - | Scan | Acunetix | Block | Block |
| 10003... | Acunetix URI Injection Vuln... | ... | Medium | - | Scan | Acunetix | Block | Block |
| 10003... | Acunetix XSS vulnerability s... | ... | Medium | - | Scan | Acunetix | Block | Block |
| 10003... | Acunetix Certified Vulnerabi... | ... | Medium | - | Scan | Acunetix | Block | Block |
| 10003... | Acunetix XSS vulnerability s... | ... | Medium | - | Scan | Acunetix | Block | Block |

- o Virtual patches are installation-free. You can use them to protect against common high-risk application vulnerabilities.
 - a. Turn on **Patches** to enable the virtual patching feature.
 - b. In the **Virtual Patches** section, click **Customize**. In the **Customize Virtual Patches Policies** dialog box, customize one or more basic virtual patching policies.

| Rule ID | Policy Name | Description | Risk Level | CVE ID | Attack Type | Victim | Default Action | Current Action |
|----------|--------------------------------|-------------|------------|---------------|-------------------|------------------|----------------|----------------|
| 10000... | Adobe ColdFusion remote c... | ... | Medium | CVE-2017-3066 | Command Execution | Adobe ColdFusion | Block | Block |
| 10000... | Apache ActiveMQ arbitrary c... | ... | Medium | CVE-2015-5254 | Command Execution | Apache ActiveMQ | Block | Block |
| 10003... | Apache Axis freemarker com... | ... | Medium | - | Command Execution | Apache Axis | Block | Block |
| 10000... | Apache CunchDB remote co... | ... | Medium | - | Command Execution | Apache CunchDB | Block | Block |
| 10000... | Apache CunchDB Elevation ... | ... | Medium | CVE-2017-1263 | Others | Apache CunchDB | Block | Block |
| 20000... | Apache Jmeter RMI deserial | ... | Medium | CVE-2018-1287 | Command | Apache | Block | Block |

21.9.2.3.2. View the traffic blocked by IPS

This topic describes how to view the details about Internet and VPC attacks that are blocked by the intrusion prevention system (IPS) feature of Cloud Firewall.

Procedure

1. Log on to [Apsara Stack Security Center](#).

2. In the left-side navigation pane, choose **Network Security > Intrusion Prevention**.
3. Click the **Internet Traffic Blocking** and **VPC Traffic Blocking** tabs to view the details of the IPS-blocked traffic.

- o **Internet Traffic Blocking**

On the **Internet Traffic Blocking** tab, you can view the blocking events of inbound and outbound traffic from the last one hour, one day, or seven days.

The **Internet Traffic Blocking** tab contains the following sections:

- **Defended Attacks:** displays the number and trend of attacks whose inbound or outbound traffic is blocked by IPS.
- **Attack Types:** displays the distribution of attacks whose inbound or outbound traffic is blocked by IPS.
- **Frequently Attacked Apps:** displays the applications whose inbound or outbound traffic is blocked by IPS.
- **Frequently Attacked Destinations:** displays the traffic distribution of attack targets whose inbound or outbound traffic is blocked by IPS.
- **Frequently Attacked Sources:** displays the traffic distribution of attack sources whose inbound or outbound traffic is blocked by IPS.
- **Detailed Data:** displays the details of each traffic blocking event. The details include the risk level, number of times the event occurred, source IP address, and destination IP address.

In the **Detailed Data** section, you can perform the following operations:

- Specify a risk level, module, traffic direction, or time range to search for events.
- Find a traffic blocking event and click **View Details** in the **Action** column to view the details. The details include the event description.

- o **VPC Traffic Blocking**

The **VPC Traffic Blocking** tab displays the suspicious traffic blocked by IPS between VPCs. You can view the details of a traffic blocking event over a specific time range. The details include the event name, risk level, and attack type.

On the **VPC Traffic Blocking** tab, you can perform the following operations:

- Specify a risk level, defense mode, attack type, or time range to search for events. After you specify the filter conditions, you must click **Search**.
- Find a traffic blocking event and click **View Details** in the **Action** column to view the details. The details include the event description and defense mode.

21.9.2.4. View security groups

This topic describes how to view the details about and relationships between security groups and Elastic Compute Service (ECS) instances in a virtual private cloud (VPC).

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Network Security > Security Groups**.
3. On the **Visual Security Groups** page, view information about security groups and ECS instances in VPCs. Select a VPC and view the numbers of security groups, vSwitches, and ECS instances in the VPC.
4. Click the VPC of the security groups that you want to view. Then, you can view details about the security groups and ECS instances in the VPC.

You can specify a time range to view the details. The time range can be one hour, one day, seven days, or one month.

- o View details about all security groups and ECS instances in the VPC.

- On the right side of the **Visual Security Group** page, you can view the total numbers of security groups and ECS instances in the VPC, the number of first visits, and the list of security groups in the **Security Group** view.

 **Note** **First Visits** indicates the total number of first visits between the ECS instances that have dependency relationships in a security group.

By default, the **Visual Security Groups** page displays the icons of all security groups. In **Security Group List**, you can find a security group and click the  icon in the **Actions** column to hide the icon of the security group on the **Visual Security Groups** page.

- In the upper-right corner of the **Visual Security Groups** page, click **Data Details**. On the **Service Nodes** tab, you can view all the security groups in the VPC in the **Security Group** view.
- View details about a single security group and the ECS instances in the security group.
 - On the **Visual Security Group** page, click the  icon in the **Actions** column in **Security Group List**. On the **Service Nodes** tab, you can view details about the ECS instances in the security group in the **ECS** view.
 - On the **Visual Security Groups** page, click the icon of a security group. On the right side of the page, you can view the total number of ECS instances in the security group, the number of dependency security groups, the number of dependent security groups, the number of first visits, and the lists of the two types of security groups in the **Security Group** view. Dependency security groups indicate the security groups on which the current security group depends, and dependent security groups indicate the security groups that depend on the current security group. The popover of the security group displays **data details**, including the following items:
 - **ECS instance details**: You can click the  icon to view details about the ECS instances in the security group.
 - **Dependency security groups**: You can click the  icon to view details about the dependency security groups.
 - **Dependent security groups**: You can click the  icon to view details about the dependent security groups.
 - **First visits**: You can click the  icon to view details about the security groups or ECS instances that have dependency relationships.
- View details about a single ECS instance.
 - On the **Visual Security Groups** page, double-click the icon of the security group that contains the ECS instance. Then, all ECS instances in the security group appear. On the right side of the page, you can view the total number of ECS instances in the security group, the numbers of dependency and dependent security groups, the number of first visits, and the ECS instance list in the **ECS** view.

By default, the **Visual Security Groups** page displays the icons of all ECS instances in a security group. In **ECS Instance List**, you can find an ECS instance and click the  icon in the **Actions** column to hide the icon of the ECS instance on the **Visual Security Groups** page.

- On the **Visual Security Groups** page, click the icon of an ECS instance. On the right side of the page, you can view the total numbers of dependency and dependent ECS instances, the number of first visits, and the lists of the two types of security groups in the ECS view. The popover of the ECS instance displays **data details**, including the following items:
 - Dependency ECS instances: You can click the  icon to view details about the dependency ECS instances.
 - Dependent ECS instances: You can click the  icon to view details about the dependent ECS instances.
 - First visits: You can click the  icon to view details about the ECS instances that have dependency relationships.

21.9.2.5. Log audit

21.9.2.5.1. View event logs

All traffic that passes through Cloud Firewall is recorded on the Log Audit page. The logs are classified into traffic logs and event logs. You can use the logs to audit your network traffic in real time and take specific actions. This topic describes how to view event logs.

Procedure

1. Log on to [Apsara Stack Security Center](#).
2. In the left-side navigation pane, choose **Network Security > Log Audit > Event Log**.
3. On the **Event Logs** tab, click **Internet Firewall**, **VPC-VPC Firewall**, or **IDC-VPC Firewall**.
4. Specify the search conditions and click **Search**.

 **Note** If you want to view all event logs, skip this step.

| Search condition | Description |
|------------------|--|
| Source IP | The source IP address of the event. |
| Destination IP | The destination IP address of the event. |
| Type | The type of the event. |
| Action | The action for the event. Valid values: <i>All</i> , <i>Monitor</i> , and <i>Discard</i> . |
| Time | The time range to query. You can specify a time range within the last seven days. |

5. View the required information in the event list. The information includes the detection time, threat type, traffic direction (inbound or outbound), source IP address, destination IP address, application type, severity, and policy action.

21.9.2.5.2. View traffic logs

All traffic that passes through Cloud Firewall is recorded on the Log Audit page. The logs are classified into traffic logs and event logs. You can use the logs to audit your network traffic in real time and take specific actions. This topic describes how to view traffic logs.

Procedure

1. Log on to [Apsara Stack Security Center](#).

2. In the left-side navigation pane, choose **Network Security > Log Audit > Traffic Log**.
3. On the **Event Logs** tab, click **Internet Firewall**, **VPC-VPC Firewall**, or **IDC-VPC Firewall**.
4. Specify the search condition and click **Search**.

Note If you want to view all traffic logs, skip this step.

| Search condition | Description |
|------------------|---|
| Source IP | The source IP address of the traffic. |
| Destination IP | The destination IP address of the traffic. |
| Application | The application type of the traffic. |
| Time | The time range to query. You can specify a time range within the last seven days. |

To enable the advanced search, perform the following steps:

- **Show Advanced Search**
Click **Show Advanced Search** to configure more search conditions.
 - **List Configuration**
Click **List Configuration** to select items for the traffic list.
5. View the required information of access traffic. The information includes the start time and end time, traffic direction (inbound or outbound), source IP address, destination IP address, application type, supported protocol, action, bytes, and packets.

21.9.3. Data Encryption Service

21.9.3.1. Data Encryption Service overview

Data Encryption Service is an encryption solution in the cloud. The underlying layer of the service uses hardware cryptographic devices that are tested and certified by the Office of the State Commercial Cryptography Administration (OSCCA). The service uses virtualization technologies to help you meet compliance and regulatory requirements in data security. This helps ensure the confidentiality of business data in the cloud.

Data Encryption Service allows you to manage keys in a secure and reliable manner and ensures reliable data encryption and decryption by using various encryption algorithms.

The following figure shows the communication mode used by Data Encryption Service. Elastic Compute Service (ECS) instances and hardware security modules (HSMs) must be in the same virtual private cloud (VPC).

Features

Data Encryption Service provides the following features:

- **Key management:** allows you to store and use keys in a secure manner, generate sub-keys by using multiple components, and import and export keys by using secure packets.
- **Data encryption:** supports all cryptographic algorithms that are approved for use in mainland China and some globally used algorithms.
- **Message authentication code (MAC) calculation:** supports multiple MAC algorithms that are defined in the China Financial Integrated Circuit (IC) Card Specifications. PBOC for short.
- **Transaction authentication:** verifies the authorization request cryptogram (ARQC) and generates the

authorization response cryptogram (ARPC) based on PBOC 2.0 and 3.0.

21.9.3.2. Management of Data Encryption Service instances

21.9.3.2.1. Create a Data Encryption Service instance

This topic describes how to create a Data Encryption Service instance.

Context

Before you can use Data Encryption Service, you must create an instance.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the top navigation bar, choose **Security > Data Encryption Service**. On the page that appears, click **Data Encryption Service**.
3. On the **Data Encryption Service** page, click **Create Instance**.
4. On the page that appears, specify **Region**, **Network Type**, **Manufacturer**, and **Device Model**.
5. Click **Create Instance**.

21.9.3.2.2. Configure a Data Encryption Service instance

This topic describes how to configure a Data Encryption Service instance.

Context

You can use a Data Encryption Service instance only in a virtual private cloud (VPC). After you create a Data Encryption Service instance, you must configure a VPC and a private IP address for the instance.

Procedure

1. [Log on to Apsara Stack Security Center.](#)
2. In the top navigation bar, choose **Security > Data Encryption Service**. On the page that appears, click **Data Encryption Service**.
3. On the **Data Encryption Service** page, find the instance that you create and click **Configure** in the **Operation** column.
4. In the **Configure IP** dialog box, configure the following parameters.

| Parameter | Description |
|-----------------------------|---|
| VPC network ID | Select a VPC. The instance and your Elastic Compute Service (ECS) instances must reside in the same VPC. |
| VPC subnet | Select a VPC subnet. |
| Allocate private IP address | Specify a private IP address. The private IP address must be within the CIDR block that is specified by VPC subnet . |

5. Click **OK**.

21.9.3.2.3. Release a Data Encryption Service instance

This topic describes how to release a Data Encryption Service instance.

Context

The number of Data Encryption Service instances that you can create in Data Encryption Service is limited. If you no longer need an instance, you can release the instance.

Procedure

1. [Log on to Apsara Stack Security Center](#).
2. In the top navigation bar, choose **Security > Data Encryption Service**. On the page that appears, click Data Encryption Service.
3. On the **Data Encryption Service** page, find the instance that you want to release and click **Release** in the **Operation** column.
4. In the message that appears, click **OK**.

21.10. Restrictions

Before logging on to Apsara Stack Security Center, make sure that your local PC meets the requirements.

Configuration requirements

| Item | Requirements |
|------------------|---|
| Browser | <ul style="list-style-type: none">• Internet Explorer: 11 or later• Google Chrome (recommended): 42.0.0 or later• Mozilla Firefox: 30 or later• Safari: 9.0.2 or later |
| Operating system | <ul style="list-style-type: none">• Windows XP, Windows 7, or later• Mac |

21.11. Log on to Cloud Security Operations Center

This topic describes how to log on to Cloud Security Operations Center.

Prerequisites

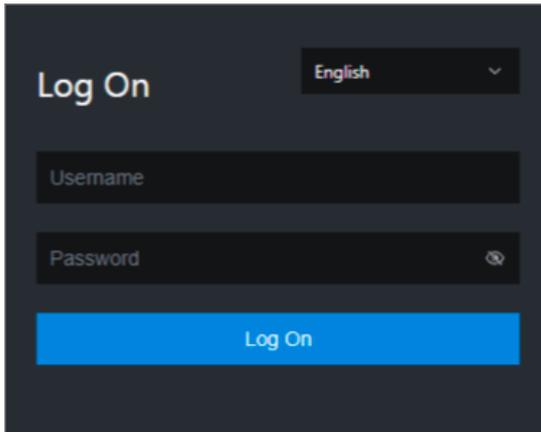
- The endpoint of the Apsara Uni-manager Operations Console and the username and password used to log on to the console are obtained from the deployment personnel or an administrator.

The endpoint of the Apsara Uni-manager Operations Console is in the following format: *region-id.ops.console.intranet-domain-id*.

- A browser is available. We recommend that you use Google Chrome.

Procedure

1. Open your browser.
2. In the address bar, enter the endpoint of the Apsara Uni-manager Operations Console. Press the Enter key.



Note You can select a language from the drop-down list in the upper-right corner of the page.

3. Enter your username and password.

Note Obtain the username and password used to log on to the Apsara Uni-manager Operations Console from the deployment personnel or an administrator.

When you log on to the Apsara Uni-manager Operations Console for the first time, you must change the password of your username.

For security reasons, your password must meet the following requirements:

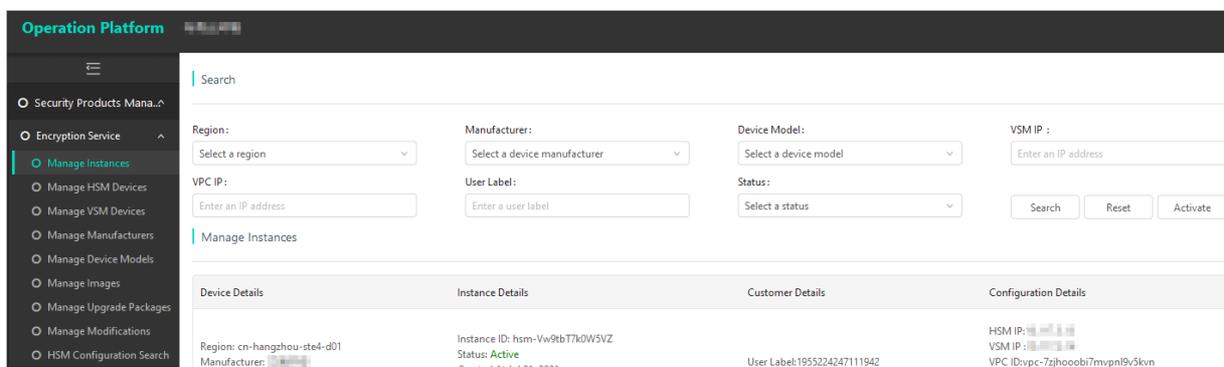
- The password contains uppercase and lowercase letters.
- The password contains digits.
- The password contains special characters such as exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%).
- The password must be 10 to 20 characters in length.

4. Click **Log On**.

5. In the top navigation bar of the Apsara Uni-manager Operations Console, click **O&M**. In the left-side navigation pane, choose **Product Management > Products**. On the page that appears, click **Cloud Security Operation Center** in the **Security Services** section.

Result

On the **Operations Platform** page, you can view the following information.



21.12. Services

21.12.1. Data Encryption Service

21.12.1.1. Manage Data Encryption Service instances

21.12.1.1.1. Create an instance

This topic describes how to create a Data Encryption Service instance.

Context

After you create an instance, you can view the instance in the **Manage Instances** section. The instance is in the *Not Configured* state.

Note Before you use Data Encryption Service, you must configure a virtual private cloud (VPC) for the instance. For more information, see [Configure VPC](#).

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage Instances**.
3. On the page that appears, click **Activate** in the **Search** section.
4. In the **Activate VSM Device** dialog box, configure the following parameters.

The screenshot shows a dialog box titled "Activate VSM Device" with a close button (X) in the top right corner. The dialog contains the following fields:

- * User Label:** A text input field with the placeholder text "Enter a user label".
- * Region:** A dropdown menu with the placeholder text "Select a region".
- * Manufacturer:** A dropdown menu with the placeholder text "Select a device manufacturer".
- * Device Model:** A dropdown menu with the placeholder text "Select a device model".
- vsmNumber:** A text input field containing the value "0".

At the bottom right of the dialog, there are two buttons: "Cancel" and "Confirm".

Note You can view the user ID in Apsara Stack Security Center. To view the user ID, log on to Apsara Stack Security Center, choose **System Configuration > Accounts**. On the page that appears, view the user ID in the **User ID** column.

5. Click **Yes**.

21.12.1.1.2. Configure a VPC

This topic describes how to configure a virtual private cloud (VPC) for a Data Encryption Service instance.

Prerequisites

A VPC is created.

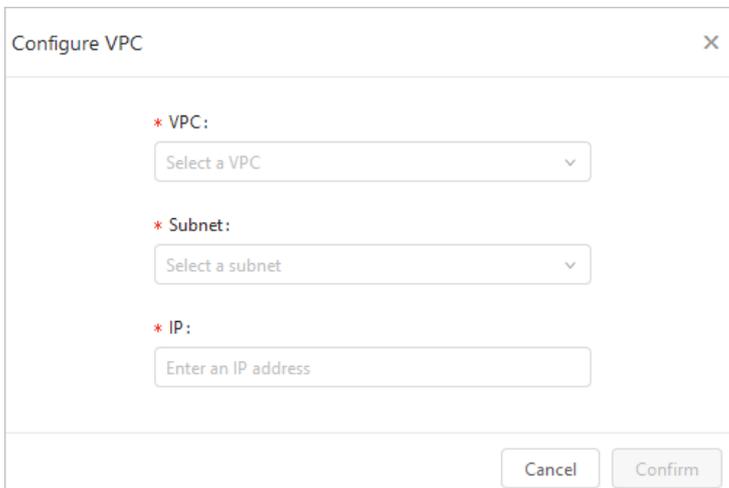
Context

Before you can use a Data Encryption Service instance, you must configure a VPC for the instance.

 **Note** After you create an instance, you must configure a VPC for the instance.

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage Instances**.
3. On the page that appears, find the instance for which you want to configure a VPC in the **Manage Instances** section and click **Configure VPC** in the Actions column.
In the **Search** section, you can specify the search conditions to find an instance.
4. In the Configure VPC dialog box, configure the following parameters for a VPC.



| Parameter | Description |
|-----------|--|
| vpc | The VPC that you want to bind with the instance. |
| Subnet | The subnet of the VPC that you want to bind with the instance. |
| IP | The IP address of the instance. The IP address must be within the subnet of the VPC. |

5. Click **Confirm**.

21.12.1.1.3. Manage an instance

This topic describes how to unbind a virtual private cloud (VPC) from a Data Encryption Service instance. This topic also describes how to release and delete the instance.

Context

You can perform the following operations on an existing instance.

| Operation | Description |
|---------------------------------|--|
| Unbind a VPC from the instance. | If the VPC or the subnet of the VPC to which the instance belongs is changed, you can unbind the VPC by performing the steps described in this topic. After you unbind the VPC, the state of the instance is changed to <i>Not Configured</i> . If you want to use the instance again, you must configure a VPC for the instance. |
| Release the instance. | If you no longer need the instance, you can release the instance. You can activate and reuse the instance that you released. |
| Delete the instance. | If the instance is in the <i>Released state</i> , you can delete the instance. |

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage Instances**.
3. On the page that appears, find the instance that you want to manage in the **Manage Instances** section and perform the following operations based on your business requirements.
 - Click **Unbind VPC** to unbind the VPC from the instance.
 - Click **Release** to release the instance.
 - Click **Delete** to delete the instance that you released.
4. In the dialog box that appears, click **OK**.

21.12.1.2. Manage HSMs

21.12.1.2.1. Add an HSM

This topic describes how to add a hardware security module (HSM).

Prerequisites

- The HSM can communicate with Cloud Security Operations Center.
- The HSM and all the virtual security modules (VSMs) on the HSM are in the running state.

Context

An HSM is a physical cryptographic device provided by a manufacturer. An HSM can host multiple VSMs.

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage HSM Devices**.
3. On the page that appears, click **Add** in the **Search** section.
4. In the Add HSM Device dialog box, configure the following parameters.

Add HSM Device
✕

*** Region :**

Select a region

*** Manufacturer:**

Select a device manufacturer

*** Device Model:**

Select a device model

*** IP :**

Enter an IP address

*** Mask:**

Enter a mask

*** Gateway:**

Enter a gateway

*** Reserve:**

Yes No

*** Control Level:**

Automatic Manual

Cancel

Confirm

| Parameter | Description |
|---------------|---|
| Region | Select the region of the HSM. |
| Manufacturer | Select the manufacturer of the HSM. |
| Device Model | Select the model of the HSM. |
| IP | Specify the IP address of the HSM. |
| Mask | Specify the mask of the HSM. |
| Gateway | Specify the IP address of the gateway. |
| Reserve | Determine whether to reserve the HSM. All the VSMS on the reserved HSM are allocated and used only when you want to migrate or update an HSM. |
| Control Level | The control level of the HSM. Valid values: <ul style="list-style-type: none"> ○ Automatic ○ Manual |

5. Click **Confirm**.

The task of adding the HSM requires some time to complete. You can go to the **Manage Modifications** page to view the task.

21.12.1.2.2. Configure the network information for an HSM

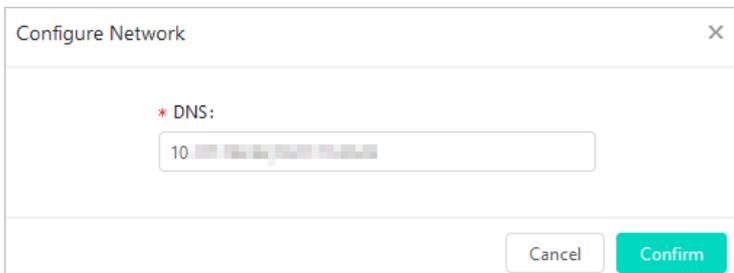
This topic describes how to configure the Domain Name System (DNS) information for a hardware security module (HSM).

Context

The network information for an HSM is configured. For more information about how to configure the network information, see [Add an HSM](#). If you want to modify the network information of the HSM, find the HSM and click **Configure Network** in the Actions column. Then, modify the DNS information.

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage HSM Devices**.
3. In the **Manage HSM Devices** section, find the HSM whose network information you want to modify and click **Configure Network**.
4. In the Configure Network dialog box, configure the DNS parameter.



5. Click **Confirm**.

21.12.1.2.3. Migrate an HSM

This topic describes how to migrate a source hardware security module (HSM) to a destination HSM.

Context

If a source HSM requires to be maintained or fails, you can migrate the source HSM to a destination HSM to ensure service continuity. After the source HSM recovers, the service is switched back to the source HSM.

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage HSM Devices**.
3. On the page that appears, find the source HSM in the **Manage HSM Devices** section and click **Migrate**.
4. In the Migrate dialog box, configure the following parameters.

Migrate
✕

* label.destDevice:

Select a %sination
 ▼

* Migration Type:

Switchover
 Failover

* Operation Type:

Automatic
 Manual

Cancel
Confirm

| Parameter | Description |
|-----------------|---|
| Destination HSM | Specify the destination HSM. |
| Migration Type | Select a scenario in which you want to migrate the source HSM to the destination HSM. Valid values: <ul style="list-style-type: none"> ◦ Switchover: <p>Select Switchover for daily maintenance. For example, you can select Switchover when you want to update the source HSM. After the update is complete, you can migrate the service back to the source HSM.</p> ◦ Failover: <p>Select Failover if the source HSM fails.</p> |
| Operation Type | Select a migration type. Valid values: <ul style="list-style-type: none"> ◦ Automatic <p>The source HSM is automatically migrated without manual operations. In most cases, we recommend that you select Automatic.</p> ◦ Manual <p>The source HSM must be manually migrated. If you select Manual, you are redirected to the Job Details page. On this page, you must manually migrate the source HSM.</p> |

5. Click **Confirm**.
6. Determine whether to view the migration process.
 - If you do not want to view the migration process, click **Cancel** to stay in the **Manage HSM Devices** section.
 - If you want to view the migration process, click **Confirm** and click OK in the message that appears to go to the **Job Details** page. On this page, you can view the migration process.

21.12.1.2.4. Update an HSM

This topic describes how to update a hardware security module (HSM).

Context

Before you update an HSM, you must migrate the HSM to another HSM that is of the same model and is running.

You must update or roll back an HSM and the VSMs in the following order:

- If you want to update an HSM and the VSMs at the same time, you must first update the HSM and then the VSMs.
- If you want to roll back an HSM and the VSMs at the same time, you must first roll back the VSMs and then the HSM.

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage HSM Devices**.
3. On the page that appears, find the HSM that you want to update in the **Manage HSM Devices** section and click **Upgrade**.
4. In the **Upgrade Device** dialog box, configure the Upgrade Package and Operation Type parameters.

5. Click **Confirm**.

21.12.1.2.5. Manage an HSM

This topic describes how to manage a hardware security module (HSM), such as view details about the HSM, reload the HSM, and configure Network Time Protocol (NTP) for the HSM.

Context

You can perform the following operations on an existing HSM.

| Operation | Description |
|---|---|
| Refresh the URL that is used to export snapshots. A snapshot contains the keys and configuration data on the HSM. | If the snapshots fail to be exported, the specified export URL may be invalid. Click Refresh Export Image URL to refresh the URL. Then, export the snapshots again. |
| View the details about the HSM. | View the details about the HSM. |
| Reserve the HSM. | A reserved HSM is allocated and used only when you want to migrate or update an HSM. If you want to update an HSM or migrate an HSM, the HSM is migrated to the reserved HSM first. |
| Refresh the information about Log Service. | Refresh the information about Log Service. |
| Reload the information about the HSM. | If the information in the dialog box that appears after you click Device Details is different from the information in the database that is connected to Data Encryption Service, click Reload to synchronize the information. |
| Enable or disable monitoring on the HSM. | Enable or disable monitoring on the HSM. |

| Operation | Description |
|----------------------------|--|
| Configure NTP for the HSM. | Configure NTP to synchronize the clock between the HSM and the NTP server. |

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage HSM Devices**.
3. On the page that appears, find the HSM that you want to manage in the **Manage HSM Devices** section and perform the following operations based on your business requirements:
 - Click **Refresh Image Export URL** to refresh the URL that is used to export snapshots.
 - Click **Device Details** to view the details about the HSM.
 - Click **Reserve** to reserve the HSM.
 - Click **Refresh SLS** to refresh the information about Log Service.
 - Click **Reload** to reload the information about the HSM.
 - Click **Modify** to enable or disable monitoring on the HSM.
 - Click **Configure NTP** to configure the NTP server and synchronization frequency.
4. In the message or dialog box that appears, click **OK** or **Confirm**.

21.12.1.3. Manage VSMs

21.12.1.3.1. Configure the network information for a VSM

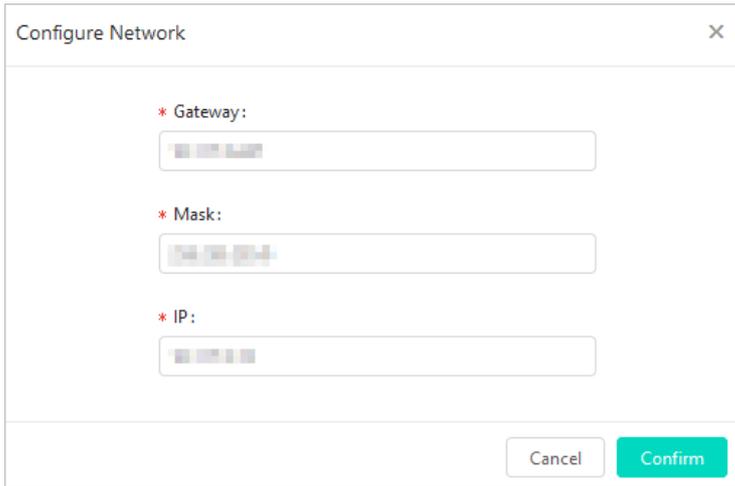
This topic describes how to configure the gateway, mask, and IP address of a virtual security module (VSM).

Context

If you want to modify the network information of the VSM, find the VSM and click **Configure Network** in the Actions column. Then, change the gateway, the mask, or the IP address of the VSM.

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage VSM Devices**.
3. In the **Manage VSM Devices** section, find the VSM whose network information you want to modify and click **Configure Network**.
4. In the **Configure Network** dialog box, configure the Gateway, Mask, and IP parameters.



5. Click **Confirm**.

21.12.1.3.2. Update a VSM

This topic describes how to update a virtual security module (VSM).

Context

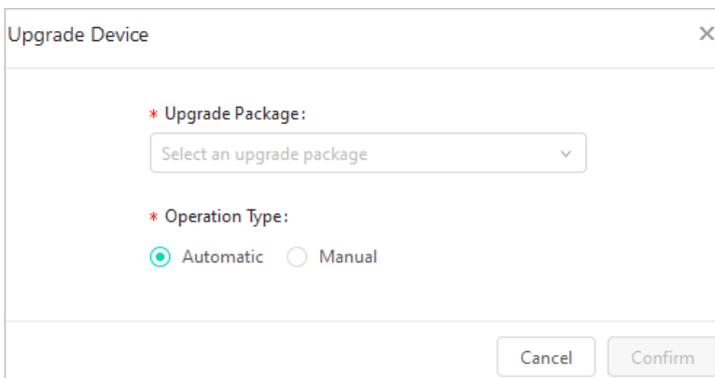
Before you update a VSM, you can migrate the VSM to another VSM that is of the same type and is running.

You must update and roll back an HSM and the VSMs in the following order:

- If you want to update an HSM and the VSMs at the same time, you must first update the HSM and then the VSMs.
- If you want to roll back an HSM and the VSMs at the same time, you must first roll back the VSMs and then the HSM.

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage VSM Devices**.
3. On the page that appears, find the VSM that you want to update in the **Manage VSM Devices** section and click **Upgrade**.
4. In the **Upgrade Device** dialog box, configure the Upgrade Package and Operation Type parameters.



5. Click **Confirm**.

21.12.1.3.3. Export snapshots

This topic describes how to export the snapshots of a virtual security module (VSM). A snapshot contains keys and configuration data on the VSM.

Context

The exported snapshots can be used to restore the VSM.

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage VSM Devices**.
3. On the page that appears, find the VSM whose snapshots you want to export in the **Manage VSM Devices** section and click **Export Snapshot**.
4. In the message that appears, click **OK**.

After the snapshots are exported, you can view the information about the snapshots. To view the information, click **Manage Images** in the left-side navigation pane. On the page that appears, view the information in the **Manage Images** section.

21.12.1.3.4. Manage a VSM

This topic describes how to manage a virtual security module (VSM), such as view details about the VSM, reserve the VSM, and restart the VSM.

Context

You can perform the following operations on the VSM.

| Operation | Description |
|--|--|
| View the details about the VSM. | View the details about the VSM. |
| Reserve the VSM. | A reserved VSM is allocated and used only when you want to migrate or update a VSM. If you want to update or migrate a VSM, the VSM is migrated to the reserved VSM first. |
| Reset the VSM. | If the VSM is released but data on the VSM is not deleted, reset the VSM. |
| Stop the VSM. | Stop the running VSM. |
| Restart the VSM. | Restart the VSM. |
| Enable or disable monitoring on the VSM. | Enable or disable monitoring on the VSM. |
| Delete the VSM. | Delete the VSM that you no longer need. |
| Change the state of the VSM to Running. | After the VSM that was in the <i>System Unavailable</i> state becomes available, click Activate to change the state of the VSM to <i>Running</i> . Then, the system can allocate the VSM. |

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage VSM Devices**.
3. On the page that appears, find the VSM that you want to manage in the **Manage VSM Devices** section and perform the following operations based on your business requirements:
 - Click **Device Details** to view the details about the VSM.

- Click **Reserve** to reserve the VSM.
 - Click **Reset** to reset the VSM.
 - Click **Stop** to stop the running VSM.
 - Click **Restart** to restart the VSM.
 - Click **Modify** to enable or disable monitoring on the VSM.
 - Click **Delete** to delete the VSM.
 - Click **Activate** to change the state of the VSM from *System Unavailable* to *Running*.
4. In the message or dialog box that appears, click **OK** or **Confirm**.

21.12.1.4. Manage manufacturers

21.12.1.4.1. Add a manufacturer

This topic describes how to add a manufacturer.

Context

Accurate manufacturer information is required for Data Encryption Service instances and hardware security modules (HSMs). Before you use an HSM that you purchased, you must configure the manufacturer information.

The following manufacturer is supported.

| Manufacturer | Code |
|--------------|------|
| TASS | jnta |

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage Manufacturers**.
3. In the upper-right corner of the page that appears, click **Add**.
4. In the Manufacturer dialog box, configure the parameters.

The screenshot shows a dialog box titled "Manufacturer" with a close button (X) in the top right corner. Inside the dialog, there are three input fields:

- A field labeled "* Manufacturer Name:" with a placeholder text "Enter a manufacturer name".
- A field labeled "* Manufacturer Code:" with a placeholder text "Enter a manufacturer code".
- A field labeled "Comments:".

At the bottom right of the dialog, there are two buttons: "Cancel" and "Confirm".

5. Click **Confirm**.

21.12.1.4.2. Manage a manufacturer

This topic describes how to modify the information about a manufacturer and delete a manufacturer.

Context

Accurate manufacturer information is required for Data Encryption Service instances and hardware security modules (HSMs). Before you use an HSM that you purchased, you must configure the manufacturer information.

The following manufacturer is supported.

| Manufacturer | Code |
|--------------|------|
| TASS | jnta |

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage Manufacturers**.
3. On the page that appears, delete a manufacturer or modify the information about a manufacturer.
 - o Click **Delete** to delete a manufacturer.
 - o Click **Modify** to modify the information about a manufacturer.
4. Click **OK** or **Confirm**.

21.12.1.5. Manage HSM models

21.12.1.5.1. Add an HSM model

This topic describes how to add a hardware security module (HSM) model.

Context

Accurate information about the HSM model is required for Data Encryption Service instances and HSMs. Before you use an HSM that you purchased, you must configure the HSM model information.

The following HSM models are supported.

| Manufacturer | HSM model | HSM model number |
|--------------|----------------------------|------------------|
| TASS | Financial HSM | jnta.EVSM |
| TASS | Server HSM | jnta.GVSM |
| TASS | Signature verification HSM | jnta.SVSM |

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage Device Models**.
3. In the upper-right corner of the page that appears, click **Add**.
4. In the Add Device Model dialog box, configure the parameters.

The screenshot shows a dialog box titled "Add Device Model" with a close button (X) in the top right corner. The form contains the following fields:

- * Manufacturer:** A dropdown menu with the text "Select a device manufacturer".
- * Device Model Name:** A text input field with the placeholder "Enter a device model name".
- * Device Model Number:** A text input field with the placeholder "Enter a device model number".
- * Port:** A text input field with the placeholder "Enter a port number. Use commas to separate multiple".
- Authentication Algorithm:** A text input field with the placeholder "Enter the authentication algorithm".
- Comments:** A text input field.

At the bottom right of the dialog box, there are two buttons: "Cancel" and "Confirm".

5. Click **Confirm**.

21.12.1.5.2. Manage an HSM model

This topic describes how to modify the information about a hardware security module (HSM) model or delete an HSM model.

Context

Accurate information about the HSM model is required for Data Encryption Service instances and HSMs. Before you use an HSM that you purchased, you must configure the HSM model information.

The following HSM models are supported.

| Manufacturer | HSM model | HSM model number |
|--------------|----------------------------|------------------|
| TASS | Financial HSM | jnta.EVSM |
| TASS | Server HSM | jnta.GVSM |
| TASS | Signature verification HSM | jnta.SVSM |

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage Device Models**.
3. On the page that appears, find the HSM model that you want to manage and perform the following operations based on your business requirements:
 - Click **Delete** to delete the HSM model.
 - Click **Modify** to modify the information about the HSM model.
4. Click **Confirm**.

21.12.1.6. View the information about snapshots

This topic describes how to view the snapshots of a virtual security module (VSM). A snapshot contains keys and configuration data on the VSM.

Context

You can use the snapshots that are displayed in the **Manage Images** section to restore VSMs. For more information about how to export snapshots, see [Export snapshots](#).

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage Images**.
3. View the information about the snapshots.

21.12.1.7. Manage update files

21.12.1.7.1. Upload an update file

This topic describes how to upload an update file for a hardware security module (HSM) or a virtual security module (VSM).

Context

You can upload an update file based on your business requirements. If an update file fails to be uploaded, you must import the required HTTPS certificate on your browser. To import the required HTTPS certificate, you can open a new browser window, import the certificate, enter the upload URL in the address bar, and then re-upload the file.

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage Upgrade Packages**.
3. On the page that appears, click **Add** in the **Search** section.
4. In the Add Upgrade Package dialog box, configure the parameters.

The screenshot shows a dialog box titled "Add Upgrade Package" with a close button (X) in the top right corner. The dialog contains the following fields and controls:

- * Manufacturer:** A dropdown menu with the placeholder text "Select a device manufacturer".
- * Device Model:** A dropdown menu with the placeholder text "Select a device model".
- * Device Type:** Two radio buttons: "HSM" (unselected) and "VSM" (selected).
- * Device Version:** A text input field with the placeholder text "Enter the device version".
- * Package Version:** A text input field with the placeholder text "Enter the upgrade package version".
- * Snapshot Digest:** A text input field with the placeholder text "Enter a snapshot digest".
- * Digest Algorithm:** A text input field with the placeholder text "Enter the snapshot digest algorithm".
- * Attachment:** A button with an upload icon and the text "Uploading files".

At the bottom of the dialog, there are two buttons: "Cancel" and "Confirm".

5. Click **Uploading files**. Then, select and upload a update file from your computer.
6. Click **Confirm**.

21.12.1.7.2. Delete an update file

This topic describes how to delete an update file of a hardware security module (HSM) or a virtual security module (VSM).

Context

You can delete update files that you no longer need to free up storage space.

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage Upgrade Packages**.
3. On the page that appears, find the update file that you want to delete in the Manage Upgrade Packages section and click **Delete** in the Operation column.

21.12.1.8. Manage tasks

21.12.1.8.1. View task details

This topic describes how to view task details.

Context

If you perform an operation in Data Encryption Service, a task is generated and displayed on the Modification Details page. The task details show how the task is executed. If the task fails, you can view the task details to identify the issues.

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage Modifications**.
3. On the page that appears, find the task whose details you want to view and click **Job Details** in the Operation column.
4. View the details about the task.

21.12.1.8.2. Terminate a task

This topic describes how to terminate a running task.

Context

You can terminate running tasks that are performed by mistake or no longer needed.

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > Manage Modifications**.
3. On the page that appears, find the running task that you want to terminate and click **Stop**.

21.12.1.9. Query the configurations of an HSM

This topic describes how to query the configurations of a hardware security module (HSM).

Context

You can query the configurations of an HSM. The configurations include the manufacturer, status, and network information about an HSM, and the virtual security modules (VSMs) on the HSM.

Procedure

1. [Log on to Cloud Security Operations Center](#).
2. Choose **Security Products Management > Encryption Service > HSM Configuration Search**.
3. On the page that appears, configure **Manufacturer** and **HSM IP** in the **Search** section and click **Search**.
4. View the configurations of the HSM and VSMs.

22.Key Management Service (KMS)

22.1. Manage keys in the KMS console

22.1.1. Log on to the KMS console

This topic describes how to log on to the Key Management Service (KMS) console.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the operations administrator before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use Google Chrome.

Procedure

1. In the address bar, enter the URL used to log on to the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that are used to log on to the console from the operations administrator.

 **Note** When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username as prompted. For security purposes, your password must meet the minimum complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters.
- Digits.
- Special characters including exclamation points (!), at signs (@), number signs (#), dollar signs (\$), and percent signs (%).

3. Click **Login**.
4. In the top navigation bar, choose **Products > Security > Key Management Service**.

22.1.2. Create a CMK

This topic describes how to create a CMK for subsequent encryption and decryption operations.

Procedure

1. [Log on to the KMS console](#).
2. On the Keys page, click **Create Key**.
3. On the **Create Key** page, select an organization from the **Organization** drop-down list. Then, the **Resource Set** and **Region** parameters are automatically set.
4. In the **Basic Settings** section, set the **Key Type**, **Key Purpose**, and **Protection Level** parameters.

You can select one of the following key types:

- Symmetric keys:
 - Aliyun_AES_256
 - Aliyun_SM4
- Asymmetric keys:

- RSA_2048
- EC_P256
- EC_P256K
- EC_SM2

 **Note** Aliyun_SM4 and EC_SM2 types are supported only in mainland China regions where managed HSMs are available.

5. In the **Basic Settings** section, set the **Alias** and **Description** parameters.
6. In the **Advanced Settings** section, set the **Rotation Period** parameter.
 - **Disable**: The key is not automatically rotated.
 - **Enable**: The key is automatically rotated. You can select or customize an interval for rotation.

 **Note**

- You can set this parameter only if the Key Type parameter is set to Aliyun_AES_256 or Aliyun_SM4.
- If the Key Material Source parameter is set to External, automatic rotation is not supported.

7. In the **Advanced Settings** section, set the **Key Material Source** parameter.
 - **Key Management Service**: Use KMS to generate key material.
 - **External**: Import key material from an external source.

 **Note** You can set this parameter only if the Key Type parameter is set to Aliyun_AES_256 or Aliyun_SM4.

8. Click **Submit**.

22.1.3. View the details of a CMK

After you create a CMK, you can view the details of the CMK, such as the key ID, status, purpose, and creation time.

Procedure

1. Log on to the KMS console. For more information, see [Log on to the KMS console](#).
2. On the Keys page, find the CMK that you want to view and click its alias or choose **More>Key Details** in the Actions column.
3. In the **Key Details** section, view the details of the CMK, such as the key ID, status, purpose, and creation time.

22.1.4. Enable a CMK

This topic describes how to enable a CMK that is in the disabled state.

Procedure

1. Log on to the KMS console. For more information, see [Log on to the KMS console](#).
2. Find a CMK in the **Disabled** state and click **Enable** in the **Actions** column.
3. In the **Enable Key** message, click **OK**.
After the CMK is enabled, the status of the CMK is changed from **Disabled** to **Enabled**.

22.1.5. Disable a CMK

This topic describes how to disable a CMK.

Context

If a CMK is disabled, the CMK cannot be used for encryption and decryption. The ciphertext encrypted by using this CMK cannot be decrypted until you re-enable the CMK.

Procedure

1. Log on to the KMS console. For more information, see [Log on to the KMS console](#).
2. Find a CMK in the **Enabled** state and click **Disable** in the **Actions** column.
3. In the **Disable Key** message, click **OK**.

After the CMK is disabled, the status of the CMK is changed from **Enabled** to **Disabled**.

22.1.6. Schedule the deletion of a CMK

You can schedule the deletion of a CMK by specifying a waiting period. After the period ends, the CMK is automatically deleted.

Context

You can specify a waiting period of 7 to 30 days.

Warning

- During the waiting period, the CMK is in the Pending Deletion state. It cannot be used to encrypt data, decrypt data, or generate data keys.
- Deleting a CMK may have a severe impact on data availability. Therefore, in most cases, we recommend that you disable a CMK instead of deleting it.
- After a CMK is deleted, it cannot be recovered. The data encrypted by using this CMK and the ciphertext data keys generated by using this CMK cannot be decrypted. Therefore, KMS only allows you to schedule the deletion of a CMK. You cannot immediately delete a CMK.
- KMS deletes the CMK within 24 hours after the specified waiting period ends.

For example, if you submit a deletion application at 14:00, September 10, 2017 and specify a waiting period of seven days, KMS deletes the CMK within 24 hours after 14:00, September 17, 2017.

Procedure

1. [Log on to the KMS console](#).
2. Find the CMK for which you want to schedule the deletion and choose **More > Schedule Key Deletion** in the **Actions** column.
3. In the **Schedule Key Deletion** dialog box, specify **Delete In (7-30 days)**.
4. Click **OK**.

The status of the CMK becomes **Pending Deletion**.

Before the specified period ends, you can cancel the scheduled deletion by performing the following operations: Choose **More > Cancel Key Deletion** and click **OK** in the **Cancel Key Deletion** message.

22.1.7. Configure the rotation policy of a CMK

This topic describes how to configure the rotation policy of a CMK.

Procedure

1. [Log on to the KMS console](#).

2. On the Keys page, find a symmetric CMK whose **Key Spec** is `Aliyun_AES_256` or `Aliyun_SM4` and choose **More > Key Details** in the **Actions** column.
3. In the **Key Version** section, click **Set Rotation Policy**.
4. In the **Set Rotation Policy** dialog box, set the **Rotation Period** parameter.
 - **Disable**: The automatic rotation is disabled for the CMK.
 - **30 Days, 90 Days, 180 Days, and 365 Days**: Select an interval for automatic CMK rotation.
 - **Customize**: Customize the interval for automatic CMK rotation.
5. Click **OK**.

Related information

- [Automatic key rotation](#)
- [Manual CMK rotation](#)

22.1.8. Generate a data key

Data keys can be used to encrypt and decrypt on-premises data. This topic describes how to generate a data key.

Procedure

1. Log on to the KMS console. For more information, see [Log on to the KMS console](#).
2. Find a symmetric CMK for which the **Key Purpose** parameter is set to **Encrypt / Decrypt** and click its alias or choose **More > Key Details** in the **Actions** column.

 **Note** The type of symmetric CMKs is `Aliyun_AES_256` or `Aliyun_SM4`.

3. In the **DataKey management** section, click **DataKey management**.
4. On the **Generate data key** tab, set the parameters as required and click **Submit**.
 - i. In the **Generate dataKey parameters** section, set the **Key Spec** and **Encryption Context** parameters.
 - ii. In the **Generate dataKey parameters** section, set the **Export using public key encryption** parameter. Valid values:
 - **Yes**: Encrypt a data key by using a specific public key and export the data key. In the **Specify the export publicKey** section, set the **PublicKey Type**, **Public Key**, and **Key Algorithm** parameters.
 - **No**: Export a data key without encrypting the data key by using a specific public key. This way, you can directly generate a data key.

 **Note** After the data key is generated, you can view the ciphertext of the data key in the **Data key export results** dialog box.

5. Optional. Click the **Data key transfer protection** tab, set the parameters as required, and then click **Submit**.

| Transfer protection scenario | Step |
|------------------------------|---|
| Use public key to protect | <ol style="list-style-type: none"> i. In the Parameters to export dataKey section, set the Data Key Ciphertext and Encryption Context parameters. ii. In the Specify the export publicKey section, set the PublicKey Type, Public Key, and Key Algorithm parameters. |

| Transfer protection scenario | Step |
|---|---|
| Data key transfer protection in the same region | <ol style="list-style-type: none"> i. In the Parameters to export dataKey section, set the Data Key Ciphertext and Source EncryptionContext parameters. ii. In the Target key parameters section, set the Target KeyId and Target EncryptionContext parameters. |
| Data key transfer protection across regions | <ol style="list-style-type: none"> i. In the Parameters to export dataKey section, set the Data Key Ciphertext and Source EncryptionContext parameters. ii. In the Target key parameters section, set the Target RegionId, Target Asymmetric KeyID, Target Asymmetric KeyVersionID, Target KeyId, and Target EncryptionContext parameters. |

 **Note** After transfer protection is configured for the data key, you can view the ciphertext of the data key in the **Data key transfer protection result** dialog box.

22.1.9. Generate a CSR

To generate a certificate, you must submit a certificate signing request (CSR) to a certificate authority (CA). This topic describes how to generate a CSR.

Procedure

1. Log on to the KMS console. For more information, see [Log on to the KMS console](#).
2. Find an asymmetric CMK for which the Key Purpose parameter is set to **Sign/Verify** and click its alias or choose **More > Key Details** in the Actions column.

 **Note** The type of asymmetric CMKs is RSA_2048, EC_P256, EC_P256K, or EC_SM2.

3. In the **Key Version** section, click **Generate CSR** in the **Actions** column.
4. In the **Generate CSR** dialog box, set the parameters as required.

Generate CSR
✕

* Common Name: 7/255 (Example: xxx company)

Name Of Organization: 7/255

Department(OU): 8/255 (Example: IT Dept)

* State(S): 7/255 (Example: Shanghai)

* Locality(L): 7/255 (Example: Shanghai)

* Country(C): ▼
(International Standard Organization ISO country code 2-digit country code. For China, please fill in CN)

Key Algorithm: ▼

Email(E): 0/255

5. Click **Generate CSR**.
6. Click **Download CSR file** to save the CSR file to your on-premises computers.

22.2. Use RAM for access control

Key Management Service (KMS) uses Resource Access Management (RAM) to control access to resources. This topic describes the resource types, actions, and policy conditions in KMS.

Apsara Stack tenant accounts have full operation permissions on their own resources. RAM users and roles are granted only operation permissions on specified resources.

Resource types in KMS

The following table describes all resource types and corresponding Alibaba Cloud Resource Names (ARNs) in KMS. They can be used in the Resource parameter of a RAM policy.

| Resource type | ARN |
|--------------------------|---|
| Abstract key container | acs:kms:\${region}:\${account}:key |
| Abstract alias container | acs:kms:\${region}:\${account}:alias |
| Key | acs:kms:\${region}:\${account}:key/\${key-id} |
| Alias | acs:kms:\${region}:\${account}:alias/\${alias-name} |

Actions in KMS

KMS defines actions used in RAM policies for each API operation that requires access control. Actions must be in the `kms:${api-name}` format.

 **Note** The DescribeRegions operation requires no access control. It can be called by Apsara Stack tenant accounts, RAM users, or RAM roles after they pass authentication.

The following table describes the RAM actions and resource types that correspond to each KMS API operation.

- Key service operations

| Operation | Action | Resource type |
|------------------------|--|--------------------------|
| ListKeys | kms:ListKeys | Abstract key container |
| CreateKey | kms:CreateKey | Abstract key container |
| DescribeKey | kms:DescribeKey | Key |
| UpdateKeyDescription | kms:UpdateKeyDescription | Key |
| EnableKey | kms:EnableKey | Key |
| DisableKey | kms:DisableKey | Key |
| ScheduleKeyDeletion | kms:ScheduleKeyDeletion | Key |
| CancelKeyDeletion | kms:CancelKeyDeletion | Key |
| GetParametersForImport | kms:GetParametersForImport | Key |
| ImportKeyMaterial | kms:ImportKeyMaterial | Key |
| DeleteKeyMaterial | kms>DeleteKeyMaterial | Key |
| ListAliases | kms:ListAliases | Abstract alias container |
| CreateAlias | kms:CreateAlias | Alias and key |
| UpdateAlias | kms:UpdateAlias | Alias and key |
| DeleteAlias | kms>DeleteAlias | Alias and key |
| ListAliasesByKeyId | kms:ListAliasesByKeyId | Key |
| CreateKeyVersion | kms:CreateKeyVersion | Key |
| DescribeKeyVersion | kms:DescribeKeyVersion | Key |
| ListKeyVersions | kms:ListKeyVersions | Key |
| UpdateRotationPolicy | kms:UpdateRotationPolicy | Key |
| Encrypt | kms:Encrypt | Key |
| Decrypt | kms:Decrypt | Key |
| ReEncrypt | <ul style="list-style-type: none"> ◦ kms:ReEncryptFrom ◦ kms:ReEncryptTo ◦ kms:ReEncrypt* | Key |
| GenerateDataKey | kms:GenerateDataKey | Key |

| Operation | Action | Resource type |
|---------------------------------|-------------------------------------|---------------|
| GenerateDataKeyWithoutPlaintext | kms:GenerateDataKeyWithoutPlaintext | Key |
| ExportDataKey | kms:ExportDataKey | Key |
| GenerateAndExportDataKey | kms:GenerateAndExportDataKey | Key |
| AsymmetricSign | kms:AsymmetricSign | Key |
| AsymmetricVerify | kms:AsymmetricVerify | Key |
| AsymmetricEncrypt | kms:AsymmetricEncrypt | Key |
| AsymmetricDecrypt | kms:AsymmetricDecrypt | Key |
| GetPublicKey | kms:GetPublicKey | Key |

- Tag management operations

| Operation | Action | Resource type |
|------------------|----------------------|---------------|
| ListResourceTags | kms:ListResourceTags | Key |
| UntagResource | kms:UntagResource | Key |
| TagResource | kms:TagResource | Key |

Policy conditions in KMS

You can add conditions to RAM policies to control access to KMS. RAM authentication will be successful only when the specified conditions are met. For example, you can use `acs:CurrentTime` to control the time period when a RAM policy is valid.

In addition to global conditions, you can use tags as filters to restrict the use of cryptographic operations such as Encrypt, Decrypt, and GenerateDataKey. Filters must be in the `kms:tag/${tag-key}` format.

RAM policy examples

- A RAM policy that allows users to access all KMS resources

```
{
  "Version": "1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms:*"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

- A RAM policy that allows users only to query keys, aliases, and key usage permissions

```
{
  "Version": "1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms:List*", "kms:Describe*",
        "kms:Encrypt", "kms:Decrypt", "kms:GenerateDataKey"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

- A RAM policy that allows users to use keys that contain the following tag to perform cryptographic operations:
 - Tag key: `Project`
 - Tag value: `Apollo`

```
{
  "Version": "1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms:Encrypt", "kms:Decrypt", "kms:GenerateDataKey"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringEqualsIgnoreCase": {
          "kms:tag/Project": [
            "Apollo"
          ]
        }
      }
    }
  ]
}
```

22.3. Use aliases

Aliases are optional to CMKs. Aliases must be unique in a region for each Apsara Stack tenant account. An Alibaba Stack tenant account can use the same alias in different regions. An alias can be bound to only one CMK in a region, but a CMK can have multiple aliases.

Overview

Although each alias is bound to a CMK, aliases are resources independent from the CMKs to which they are bound. Take note of the following points about aliases:

- You can call the `UpdateAlias` operation to bind an alias to a different CMK. This operation does not affect the original CMK.
- If you delete an alias, the CMK to which the alias is bound is not deleted.
- RAM users must be authorized before they can perform operations on an alias. For more information, see [Use RAM for access control](#).

- Aliases cannot be modified. To change the alias of a CMK, you must delete the original alias and create another one for the CMK.

You can replace the CMK ID in the request parameters of the following API operations with an alias that is bound to the CMK:

- DescribeKey
- Encrypt
- GenerateDataKey
- GenerateDataKeyWithoutPlaintext

 **Note** For more information about the API operations, see the *API reference* topic of *Developer Guide*.

To specify an alias instead of a CMK ID in the request parameters of the preceding operations, a RAM user must have the relevant permissions on the CMK. The RAM user does not need to have permissions on the alias.

When you use an alias, you must make sure that the alias is complete. Example:

```
// A complete alias must have the alias/ prefix.
alias/example
```

Create an alias

An alias must contain the `alias/` prefix. An alias, excluding the prefix, can contain letters, digits, underscores (`_`), hyphens (`-`), and forward slashes (`/`). An alias, excluding the prefix, must be 1 to 255 characters in length.

To allow a RAM user to create an alias, you must create a custom policy to grant the RAM user the permissions on the alias and the CMK to which the alias is bound. The following sample policy allows User 123456 to create an alias named `alias/example` for CMK `08ec3bb9-034f-485b-b1cd-3459baa8****`:

```
{
  "Version": "1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms:CreateAlias"
      ],
      "Resource": [
        "acs:kms:cn-hangzhou:123456:key/08ec3bb9-034f-485b-b1cd-3459baa8****",
        "acs:kms:cn-hangzhou:123456:alias/example"
      ]
    }
  ]
}
```

 **Note** A new alias created for a CMK does not affect the existing aliases of the CMK.

You can use one of the following methods to create an alias:

- Create an alias in the KMS console.
 - [Log on to the KMS console](#).
 - In the left-side navigation pane, click **Keys**.
 - Find the CMK for which you want to create an alias and click its alias. In the Aliases section, click **Create Alias**.
 - In the **Create Alias** dialog box, set the Alias Name parameter.

v. Click **OK**.

- Create an alias by calling the `CreateAlias` operation.
- Create an alias by running the `aliyun kms CreateAlias` command in Apsara Stack CLI.

```
aliyun kms CreateAlias --KeyId 08ec3bb9-034f-485b-b1cd-3459baa8**** --AliasName alias/example
```

Bind an alias to a different CMK

You can bind an existing alias to a different CMK.

To allow a RAM user to bind an existing alias to a different CMK, you must create a custom policy to grant the RAM user the permissions on the original CMK, new CMK, and alias. The following sample policy allows User 123456 to bind the `alias/example` alias to CMK 127d2f84-ee5f-4f4d-9d41-dbc1aca2****. The original CMK to which this alias is bound is 08ec3bb9-034f-485b-b1cd-3459baa8****.

```
{
  "Version": "1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms:UpdateAlias"
      ],
      "Resource": [
        "acs:kms:cn-hangzhou:123456:key/08ec3bb9-034f-485b-b1cd-3459baa8****",
        "acs:kms:cn-hangzhou:123456:key/127d2f84-ee5f-4f4d-9d41-dbc1aca2****",
        "acs:kms:cn-hangzhou:123456:alias/example"
      ]
    }
  ]
}
```

You can use one of the following methods to bind an alias to a different CMK:

- Bind an alias to a different CMK by calling the `UpdateAlias` operation.
- Bind an alias to a different CMK by running the `aliyun kms UpdateAlias` command in Apsara Stack CLI.

```
aliyun kms UpdateAlias --AliasName alias/example --KeyId 127d2f84-ee5f-4f4d-9d41-dbc1aca2****
```

Delete an alias

You can delete aliases that are not required. If you delete an alias, the CMK to which the alias is bound is not affected.

To allow a RAM user to delete an alias, you must create a custom policy to grant the RAM user the permissions on the alias and the CMK to which the alias is bound. The following sample policy allows User 123456 to delete the `alias/example` alias that is bound to CMK 127d2f84-ee5f-4f4d-9d41-dbc1aca2****:

```
{
  "Version": "1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms:DeleteAlias"
      ],
      "Resource": [
        "acs:kms:cn-hangzhou:123456:key/127d2f84-ee5f-4f4d-9d41-dbc1aca2****",
        "acs:kms:cn-hangzhou:123456:alias/example"
      ]
    }
  ]
}
```

You can use one of the following methods to delete an alias:

- Delete an alias in the KMS console.
 - i. [Log on to the KMS console](#).
 - ii. In the left-side navigation pane, click **Keys**.
 - iii. Find the CMK for which you want to delete an alias and click its alias.
 - iv. In the **Aliases** section, click **Delete Alias** next to the alias to be deleted.
 - v. In the **Delete Alias** message, click **OK**.
- Delete an alias by calling the `DeleteAlias` operation.
- Delete an alias by running the `aliyun kms DeleteAlias` command in Apsara Stack CLI.

```
aliyun kms DeleteAlias --AliasName alias/example
```

Query aliases

You can query all aliases that belong to your Apsara Stack tenant account in the current region.

To allow a RAM user to query aliases, you must create a custom policy to grant the RAM user the permissions on aliases. The following sample policy allows User 123456 to query aliases:

```
{
  "Version": "1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms:ListAliases"
      ],
      "Resource": [
        "acs:kms:cn-hangzhou:123456:alias"
      ]
    }
  ]
}
```

You can use one of the following methods to query aliases:

- Query aliases by calling the `ListAliases` operation.
- Query aliases by running the `aliyun kms ListAliases` command in Apsara Stack CLI.

aliyun kms ListAliases

Query aliases bound to a specified CMK

If you query the aliases bound to a specified CMK, only the related aliases are returned.

To allow a RAM user to query the aliases bound to a specified CMK, you must create a custom policy to grant the RAM user the permissions on the CMK. The following sample policy allows User 123456 to query the aliases bound to CMK 127d2f84-ee5f-4f4d-9d41-dbc1aca2****:

```
{
  "Version": "1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms:ListAliasesByKeyId"
      ],
      "Resource": [
        "acs:kms:cn-hangzhou:123456:key/127d2f84-ee5f-4f4d-9d41-dbc1aca2****"
      ]
    }
  ]
}
```

You can use one of the following methods to query aliases bound to a specified CMK:

- Query aliases bound to a specified CMK by calling the ListAliasesByKeyId operation.
- Query aliases bound to a specified CMK by running the `aliyun kms ListAliasesByKeyId` command in Apsara Stack CLI.

aliyun kms ListAliasesByKeyId --KeyId 127d2f84-ee5f-4f4d-9d41-dbc1aca2****

22.4. Use CMKs

The key service is a core component of Key Management Service (KMS). The key service provides fully managed keys and key protection features. The key service supports simple data encryption and digital signature management based on cloud-native API operations.

Key-based cryptographic algorithms

The following table describes the encryption algorithms supported by KMS.

| Algorithm class | Algorithm subclass | Encryption and decryption | Signature generation and verification |
|--------------------------|-------------------------|---------------------------|---------------------------------------|
| Symmetric key algorithm | AES | Supported | Not supported |
| Symmetric key algorithm | SM4 <small>Note</small> | Supported | Not supported |
| Asymmetric key algorithm | RSA | Supported | Supported |
| Asymmetric key algorithm | ECC | Not supported | Supported |
| Asymmetric key algorithm | SM2 <small>Note</small> | Supported | Supported |

Symmetric keys are used to encrypt or decrypt data. If you do not specify the `KeySpec` parameter during key creation, KMS creates a symmetric key. You can call the `Encrypt` or `Decrypt` operation to encrypt or decrypt data without the need to obtain the plaintext of a symmetric key. For more information, see [Overview of symmetric encryption](#).

Asymmetric keys can be used to encrypt data, decrypt data, generate a signature, or verify a signature. An asymmetric CMK in KMS consists of a public key and a private key, which are cryptographically related to each other. The public key can be made available for anyone to use, but the private key must be kept secure. To keep private keys secure, KMS does not provide an API operation for you to export the private key of an asymmetric key pair. You can use a private key to decrypt data or generate a signature by calling the related operations. Anyone with a public key can use it to encrypt data or verify the signature generated by the corresponding private key. For more information, see [Overview of asymmetric keys](#).

 **Note**

- Only hardware security modules (HSMs) support SM2 and SM4 keys.
- Only KMS of the Advanced edition supports Rivest-Shamir-Adleman (RSA) and elliptic-curve cryptography (ECC) keys.

Key protection level

KMS provides the managed HSM feature. You can set the protection level of your CMK to HSM to manage the CMK in an HSM. The managed HSM feature allows you to use HSMs as dedicated hardware to safeguard keys. For a CMK whose protection level is HSM, the plaintext of its key material is stored only inside an HSM. KMS calls an HSM-related API operation to perform cryptographic operations. During the operations, no one can have access to the plaintext of the key material. The plaintext of the key material cannot be exported from the HSM.

Key managers

In most cases, you are the manager of your CMKs in KMS, and their `Creator` attributes are set to the ID of your Apsara Stack tenant account.

Apsara Stack services can be integrated with KMS to implement server-side encryption (SSE). In this scenario, KMS allows a service to automatically manage an encryption key that can be used only by the service to encrypt your data. This makes it easier for you to use the entry-level data encryption features and reduces your workload in the management of key lifecycles and permissions. These service-managed keys are called service keys. To facilitate identification, KMS sets the `Creator` attribute of the service key managed by an Apsara Stack service to the code of this service and assigns an alias in the format of `acs/<Code of the Apsara Stack service>` to the service key. For example, the `Creator` attribute of the service key managed by Object Storage Service (OSS) is set to OSS and the `acs/oss` alias is bound to the service key.

22.5. Use symmetric keys

22.5.1. Overview

This topic describes symmetric encryption, which is the most commonly used data encryption method. KMS provides easy-to-use API operations that allow you to encrypt and decrypt data on the cloud.

 **Note** For more information about the API operations, see the *API reference* topic of *Developer Guide*.

If you do not specify the `KeySpec` parameter during key creation, KMS creates a symmetric key. KMS supports popular symmetric key algorithms and provides high-level data security by using strong cryptography.

Types of symmetric keys

The following table describes the types of symmetric keys that KMS supports.

| Algorithm | Key length | Key type | Data encryption mode | Protection level |
|---------------------|------------|----------------|----------------------|---|
| AES | 256 bits | Aliyun_AES_256 | GCM | <ul style="list-style-type: none"> • Software • HSM |
| SM4 ^{Note} | 128 bits | Aliyun_SM4 | GCM | HSM |

Encryption and decryption features

To encrypt data, you need only to specify the ID or alias of an CMK. KMS uses the specified CMK for encryption and returns the ciphertext. When you call the Decrypt operation to decrypt data, you need only to specify the ciphertext that you want to decrypt. You do not need to specify the CMK. The Decrypt operation is available for you to decrypt the ciphertext that is generated by calling the Encrypt, GenerateDataKey, or GenerateDataKeyWithoutPlaintext operation.

AAD

Symmetric keys of KMS use Galois/Counter Mode (GCM) for block ciphers. You can use additional authenticated data (AAD) to provide supplemental protection for the integrity of encrypted data. KMS allows you to customize authentication data by encapsulating the entered ADD. For more information, see [EncryptionContext](#).

Envelope encryption

KMS allows you to generate a two-level key hierarchy to accelerate envelope encryption by calling the GenerateDataKey and GenerateDataKeyWithoutPlaintext operations.

For more information about envelope encryption, see the *Envelope encryption* in the *Features* topic of *Technical White Paper*.

Rotation of symmetric keys

Each symmetric CMK that is generated in KMS supports multiple key versions. KMS automatically rotates CMKs by generating new key versions. You can customize the key rotation policy.

If a CMK has multiple versions, the latest version of the CMK is used in the Encrypt, GenerateDataKey, and GenerateDataKeyWithoutPlaintext operations to encrypt data. To decrypt the data, you do not need to specify the ID or the CMK or key version. KMS automatically identifies the CMK and its key version with which the data is encrypted. Then, KMS uses the key material of the identified key version to decrypt the ciphertext.

KMS rotates a CMK by generating a new version of the CMK. After a rotation is complete, KMS automatically uses the new key version to encrypt data. However, the earlier key version is still available for decrypting the ciphertext generated before the rotation. For more information, see [Automatic key rotation](#).

BYOK feature

KMS allows you to encrypt your data in the cloud by using the Bring Your Own Key (BYOK) feature. This feature helps you meet stringent security and compliance requirements. We recommend that you use a managed hardware security module (HSM) to protect your keys by importing key material into CMKs with HSM-grade security. The keys imported into the managed HSM can only be destroyed, and their plaintext cannot be exported.

22.5.2. EncryptionContext

EncryptionContext is a JSON string that may be used in KMS API operations, such as Encrypt, GenerateDataKey, and Decrypt.

Roles of EncryptionContext

EncryptionContext is a JSON string, and it must be in the string:string format. EncryptionContext is used to ensure data integrity.

If this parameter is specified during encryption, you must specify an equivalent `EncryptionContext` value for decryption. You can call the `Encrypt` or `GenerateDataKey` operation for encryption and call the `Decrypt` operation for decryption. `EncryptionContext` is related to decryption, but `EncryptionContext` is not included in ciphertext that is specified by the `CipherBlob` parameter.

Valid values of EncryptionContext

A valid value of `EncryptionContext` is a JSON string of up to 8,192 characters in the `string:string` format. When you specify `EncryptionContext` for an API operation, take note of the escape characters.

Sample valid `EncryptionContext`:

```
{"ValidKey":"ValidValue"}
{"Key1":"Value1","Key2":"Value2"}
```

Sample invalid `EncryptionContext`:

```
[{"Key":"Value"}] // JSON array
{"Key":12345} //String-int
{"Key":["value1","value2"]} // String:Array
```

Equivalent EncryptionContext

`EncryptionContext` is a map or hash table in the `string:string` format. When `EncryptionContext` is used as a parameter, two values of the `EncryptionContext` parameter are considered to be equivalent if their key-value pairs are consistent. If `EncryptionContext` is specified during encryption, you can specify an equivalent `EncryptionContext` value to decrypt the ciphertext. The sequences of the key-value pairs of two `EncryptionContext` values can be different.

Sample equivalent `EncryptionContext` values:

```
{"Key1":"Value1","Key2":"Value2"} is equivalent to {"Key2":"Value2","Key1":"Value1"}.
```

22.5.3. Import and delete key material

This topic describes how to import and delete external key material for a CMK.

Context

CMKs are basic resources of KMS. A CMK is composed of a key ID, basic metadata such as key status, and key material that is used to encrypt and decrypt data. By default, KMS generates key material when you call the `CreateKey` operation to create a CMK. You can select an external key as the material source. In this case, KMS does not create key material. You must import external key material to the CMK. You can call the `DescribeKey` operation to check the key material source of an existing CMK.

 **Note** For more information about the API operations, see the *API reference* topic of *Developer Guide*.

Keys can be divided into normal keys and external keys based on the source of key material.

- If the value of `Origin` in `KeyMetadata` is `Aliyun_KMS`, the key material is generated by KMS. In this case, the CMK is considered a **common key**.
- If the value of `Origin` is `EXTERNAL`, the key material is imported from an external source. In this case, the CMK is considered an **external key**.

Before you import external key material, take note of the following points:

- Make sure that the source of randomness from which the key material is generated meets security requirements.
- Make sure that the key material is reliable.

- KMS ensures the high availability of imported key material. However, KMS cannot ensure that the imported key material has the same reliability as the key material generated by KMS.
- You can call the `DeleteKeyMaterial` operation to delete the key material that you have imported. You can also set a validity period to enable KMS to automatically delete the key material after the validity period ends. The CMK is not deleted. To delete key material generated by KMS, you can only call the `ScheduleKeyDeletion` operation to specify a waiting period of 7 to 30 days for deleting the CMK. The key material is deleted with the relevant CMK after the waiting period ends.
- After you delete the imported key material, you can re-import the same key material to make the relevant CMK available again. Therefore, we recommend that you save a copy of the key material.
- Key material is unique for each CMK. When you import key material into a CMK, the CMK is associated with that key material. Even after the key material expires or is deleted, you cannot import different key material into that CMK. If you need to rotate a CMK that uses external key material, you must create a CMK and then import new key material.
- CMKs are independent. You cannot use a CMK to decrypt data that is encrypted by using another CMK, even if the two CMKs use the same key material.
- The key material to be imported must be a 256-bit symmetric key.

Import key material

1. Create an external key.

You can use one of the following methods to create an external key:

- Method 1: Log on to the KMS console and click **Create Key**. On the **Create Key** page, set the **Key Material Source** parameter to **External** in the **Advanced Settings** section. For more information, see [Create a CMK](#).
- Method 2: Call the `CreateKey` operation. Set the `Origin` parameter to `EXTERNAL`.

```
./aliyun --skip-secure-verify --region ${please-replace-your-region} --endpoint ${please-replace-your-endpoint} --access-key-id ${please-replace-your-access-key-id} --access-key-secret ${please-replace-your-access-key-secret} kms CreateKey --Origin EXTERNAL --Description "External key"
```

2. Obtain the parameters that are used to import key material.

The parameters include a public key and an import token. The public key is used to encrypt the key material. You can use one of the following methods to obtain the parameters:

- Method 1: Obtain the parameters in the KMS console.
- Method 2: Obtain the parameters by calling the `GetParametersForImport` operation.

```
./aliyun --skip-secure-verify --region ${please-replace-your-region} --endpoint ${please-replace-your-endpoint} --access-key-id ${please-replace-your-access-key-id} --access-key-secret ${please-replace-your-access-key-secret} kms GetParametersForImport --KeyId 1339cb7d-54d3-47e0-b595-c7d3dba8**** --WrappingAlgorithm RSAES_OAEP_SHA_1 --WrappingKeySpec RSA_2048
```

3. Import key material.

Note

- You can import key material into an external key that does not have key material. You can also reset the expiration time of key material or re-import key material that has expired or been deleted.
- Each import token is bound to a public key that is used to encrypt key material. A CMK is specified when an import token is generated. The import token can only be used to import key material into the specified CMK.
- The lifecycle of an import token is 24 hours. It can be used repeatedly within this period. After it expires, you must obtain a new import token and a new public key.

- i. Use the public key to encrypt the key material.

The public key is a 2,048-bit Rivest-Shamir-Adleman (RSA) public key. The encryption algorithm must be the same as that specified when you obtain the parameters that are used to import the key material. The public key returned when you call the `GetParametersForImport` operation is Base64-encoded. You must first decode the public key. KMS supports the following encryption algorithms: `RSAES_OAEP_SHA_1`, `RSAES_OAEP_SHA_256`, and `RSAES_PKCS1_V1_5`.

- ii. Encode the encrypted key material in the Base64 format.
- iii. Call the `ImportKeyMaterial` operation to import the encoded key material and the import token to KMS.

```
./aliyun --skip-secure-verify --region ${please-replace-your-region} --endpoint ${please-replace-your-endpoint} --access-key-id ${please-replace-your-access-key-id} --access-key-secret ${please-replace-your-access-key-secret} kms ImportKeyMaterial --KeyId 1339cb7d-54d3-47e0-b595-c7d3dba8**** --EncryptedKeyMaterial xxx -ImportToken xxxx
```

Delete key material

After you import key material into an external key, you can use the external key in the same way as a normal key. The only difference lies in that the key material of an external key may expire and can be independently deleted. After the key material of an external key expires or is deleted, the external key can no longer be used, and the ciphertext encrypted by using this external key cannot be decrypted. To use the external key and decrypt the related ciphertext again, you must re-import the same key material.

If an external key is in the `PendingDeletion` state when its key material expires or is deleted, the key status does not change after the key material expires or is deleted. Otherwise, the key status becomes `PendingImport` after the key material expires or is deleted.

You can use one of the following methods to delete key material:

- Method 1: Delete the key material in the KMS console.
- Method 2: Delete the key material by calling the `DeleteKeyMaterial` operation.

```
./aliyun --skip-secure-verify --region ${please-replace-your-region} --endpoint ${please-replace-your-endpoint} --access-key-id ${please-replace-your-access-key-id} --access-key-secret ${please-replace-your-access-key-secret} kms DeleteKeyMaterial --KeyId 1339cb7d-54d3-47e0-b595-c7d3dba8****
```

Use OpenSSL to encrypt and upload key material

1. Create an external key.
2. Create key material.

The key material must be a 256-bit symmetric key. In this example, OpenSSL is used to generate a 32-byte random number.

```
openssl rand -out KeyMaterial.bin 32
```

3. Obtain parameters that are used to import the key material.
4. Encrypt the key material.

- i. Decode the public key that is used to encrypt the key material in the Base64 format.
- ii. Use an encryption algorithm such as `RSAES_OAEP_SHA_1` to encrypt the key material.
- iii. Encode the encrypted key material in the Base64 format and save it to a text file.

```
openssl rand -out KeyMaterial.bin 32
openssl enc -d -base64 -A -in PublicKey_base64.txt -out PublicKey.bin
openssl rsautl -encrypt -in KeyMaterial.bin -oaep -inkey PublicKey.bin -keyform DER -pubin -out EncryptedKeyMaterial.bin
openssl enc -e -base64 -A -in EncryptedKeyMaterial.bin -out EncryptedKeyMaterial_base64.txt
```

- iv. Upload the encrypted key material and import token.

Use SDK for Java to encrypt and upload key material

```
// Use the latest KMS SDK for Java.
//KmsClient.java
import com.aliyuncs.DefaultAcsClient;
import com.aliyuncs.exceptions.ClientException;
import com.aliyuncs.http.FormatType;
import com.aliyuncs.kms.model.v20160120.*;
import com.aliyuncs.profile.DefaultProfile;
// Encapsulate KMS API operations
public class KmsClient {
    DefaultAcsClient client;
    public KmsClient(String regionId, String ak, String secret) {
        DefaultProfile profile = DefaultProfile.getProfile(regionId, ak, secret);
        this.client = new DefaultAcsClient(profile);
    }
    // Specify a custom endpoint. In most cases, a VPC endpoint is specified. You can also specify a custom endpoint when
    you cannot obtain the IP address in the Apsara Stack environment based on the regionId parameter.
    public KmsClient(String regionId, String ak, String secret, String endpoint) {
        DefaultProfile.addEndpoint(regionId, "kms", endpoint);
        DefaultProfile profile = DefaultProfile.getProfile(regionId, ak, secret);
        this.client = new DefaultAcsClient(profile);
    }
    public CreateKeyResponse createKey() throws Exception {
        CreateKeyRequest request = new CreateKeyRequest();
        request.setOrigin("EXTERNAL");// Create an external key.
        return this.client.getAcsResponse(request);
    }
    public GetParametersForImportResponse getParametersForImport(String keyId, String wrappingKeySpec, String wrappingAlgorithm) throws ClientException {
        GetParametersForImportRequest request = new GetParametersForImportRequest();
        request.setAcceptFormat(FormatType.JSON);
        request.setKeyId(keyId);
        request.setWrappingKeySpec(wrappingKeySpec);
        request.setWrappingAlgorithm(wrappingAlgorithm);
        return this.client.getAcsResponse(request);
    }
    public ImportKeyMaterialResponse importKeyMaterial(String keyId, String importToken, String encryptedKeyMaterial, Long expireTimestamp) throws ClientException {
        ImportKeyMaterialRequest request = new ImportKeyMaterialRequest();
        request.setAcceptFormat(FormatType.JSON);
        request.setKeyId(keyId);
        request.setImportToken(importToken);
        request.setEncryptedKeyMaterial(encryptedKeyMaterial);
        request.setKeyMaterialExpireUnix(expireTimestamp);
        return this.client.getAcsResponse(request);
    }
    //... Omitted. The remaining operations are encapsulated in the same way as the preceding operations.
}
//example.java
import com.aliyuncs.kms.model.v20160120.*;
import javax.crypto.Cipher;
import javax.crypto.spec.OAEPParameterSpec;
import javax.crypto.spec.PSource.PSpecified;
import javax.xml.bind.DatatypeConverter;
import java.security.KeyFactory;
import java.security.PublicKey;
import java.security.spec.MGF1ParameterSpec;
import java.security.spec.X509EncodedKeySpec;
import java.util.Random;
```

```

public class CreateAndImportExample {
    public static void main(String[] args) {
        String regionId = "cn-hangzhou";
        String accessKeyId = "**** Provide your AccessKeyId ****";
        String accessKeySecret = "**** Provide your AccessKeySecret ****";
        KmsClient kmsclient = new KmsClient(regionId, accessKeyId, accessKeySecret);
        //String endpoint = "**** VPC Endpoint ****";
        //KmsClient kmsclient = new KmsClient(regionId, accessKeyId, accessKeySecret, endpoint);
        //Create External Key
        try {
            CreateKeyResponse keyResponse = kmsclient.createKey();
            String keyId = keyResponse.getKeyMetadata().getKeyId();
            // Generate a 32-byte random number.
            byte[] keyMaterial = new byte[32];
            new Random().nextBytes(keyMaterial);
            // Obtain the parameters that are used to import key material.
            GetParametersForImportResponse paramResponse = kmsclient.getParametersForImport(keyId, "RSA_2048", "RSA
ES_OAEP_SHA_256");
            String importToken = paramResponse.getImportToken();
            String encryptPublicKey = paramResponse.getPublicKey();
            // Decode the public key in the Base64 format.
            byte[] publicKeyDer = DatatypeConverter.parseBase64Binary(encryptPublicKey);
            // Use RSA to parse the public key.
            KeyFactory keyFact = KeyFactory.getInstance("RSA");
            X509EncodedKeySpec spec = new X509EncodedKeySpec(publicKeyDer);
            PublicKey publicKey = keyFact.generatePublic(spec);
            // Encrypt the key material.
            Cipher oaepFromAlgo = Cipher.getInstance("RSA/ECB/OAEPWithSHA-1AndMGF1Padding");
            String hashFunc = "SHA-256";
            OAEPParameterSpec oaepParams = new OAEPParameterSpec(hashFunc, "MGF1", new MGF1ParameterSpec(hashFunc), PSpecified.DEFAULT);
            oaepFromAlgo.init(Cipher.ENCRYPT_MODE, publicKey, oaepParams);
            byte[] cipherDer = oaepFromAlgo.doFinal(keyMaterial);
            // Encode the encrypted key material in the Base64 format.
            String encryptedKeyMaterial = DatatypeConverter.printBase64Binary(cipherDer);
            // Import the key material.
            Long expireTimestamp = 1546272000L; // The UNIX timestamp, which is precise to the second. The value 0 indicates
that the key material does not expire.
            kmsclient.importKeyMaterial(keyId, importToken, encryptedKeyMaterial, expireTimestamp);
        } catch (Exception e) {
            //... Omitted.
        }
    }
}

```

22.6. Use asymmetric keys

22.6.1. Overview

Unlike symmetric keys, asymmetric keys are used to verify digital signatures or encrypt sensitive information between systems with different trust levels.

You can create asymmetric CMKs in KMS. An asymmetric key pair consists of a public key and a private key, which are cryptographically related to each other. The public key is available for all to use, but the private key must be kept secure and used only by trusted users. KMS supports popular asymmetric key algorithms and provides sufficient security strength to ensure the security of encrypted data and digital signatures.

Types of asymmetric keys

The following table describes the types of asymmetric keys that KMS supports.

| Algorithm | Key type | Description | Purpose |
|-----------|---|-----------------------------------|---|
| RSA | RSA_2048 | RSA asymmetric cryptosystem | <ul style="list-style-type: none"> Encrypt or decrypt data. Generate a digital signature. |
| ECC | <ul style="list-style-type: none"> EC_P256: National Institute of Standards and Technology (NIST)-recommended elliptic curve P-256 EC_P256K: Standards for Efficient Cryptography Group (SECG) elliptic curve secp256k1 | Elliptic-curve cryptography (ECC) | Generate a digital signature. |
| SM2 | EC_SM2 | ECC defined by GB/T 32918 | <ul style="list-style-type: none"> Encrypt or decrypt data. Generate a digital signature. |

Data encryption

Asymmetric keys are used to encrypt and transmit sensitive information. The following steps describe a typical scenario:

1. An information receiver distributes a public key to a transmitter.
2. The transmitter uses the public key to encrypt sensitive information.
3. The transmitter sends the ciphertext generated from the sensitive information to the information receiver.
4. The information receiver uses the private key to decrypt the ciphertext.

The private key can be used only by the information receiver. This ensures that the plaintext of sensitive information cannot be intercepted and decrypted by unauthorized parties during transmission. This encryption method is widely used to exchange keys. For example, session keys are exchanged in Transport Layer Security (TLS) handshakes, and encryption keys are exported and imported between different hardware security modules (HSMs).

For more information, see [Encrypt and decrypt data by using an asymmetric CMK](#).

Digital signature

Asymmetric keys are also used to generate digital signatures. Private keys can be used to sign messages or information. Private keys are strictly protected and can be used only by trusted users to generate signatures. After a signature is generated, you can use the corresponding public key to verify the signature to achieve the following purposes:

- Verify data integrity. If the data does not match its signature, the data may be tampered with.
- Verify message authenticity. If a message does not match its signature, the message transmitter does not hold the private key.
- Provide non-repudiation for signatures. If the data matches its signature, the signer cannot deny this signature.

The following operations describe a typical signature verification scenario:

1. A signer sends a public key to a receiver.
2. The signer uses the private key that matches the public key to sign data.
3. The signer sends the data and signature to the message receiver.

4. After the data and signature are received, the receiver uses the public key to verify the signature.

Digital signatures are widely used to defend against data tampering and authenticate identities.

For more information, see [Generate and verify a digital signature by using an asymmetric CMK](#).

Key versions

KMS does not support automatic rotation of asymmetric CMKs due to the limits of public and private keys. You can call the `CreateKeyVersion` operation to create another key version in a specific CMK and generate a new pair of public and private keys. If you use a new key version to generate a digital signature or encrypt data, you must also distribute the new version of the public key.

In addition, unlike symmetric CMKs, asymmetric CMKs do not have a primary key version. Therefore, you must specify the corresponding CMK ID or CMK alias and a key version if you need to call the operations related to asymmetric CMKs in KMS.

Public key operations

In most cases, you can call the `GetPublicKey` operation to obtain a public key and distribute the public key to users for encryption or verification. Then, the users can use cryptographic libraries such as OpenSSL and Java Cryptography Extension (JCE) to perform calculation in their business systems.

You can also call the `AsymmetricEncrypt` or `AsymmetricVerify` operation to perform public key operations. KMS records the logs of the calls and allows you to use RAM to limit the use of public keys. Compared with calculation in business systems, KMS offers flexible features that better suit your needs.

Private key operations

You can only call the `AsymmetricDecrypt` or `AsymmetricSign` operation to decrypt data or generate a digital signature with a private key.

22.6.2. Encrypt and decrypt data by using an asymmetric CMK

This topic describes how to use an asymmetric customer master key (CMK) to encrypt and decrypt data in Alibaba Cloud CLI.

Asymmetric encryption and decryption generally include the following steps:

1. An information receiver distributes a public key to a transmitter.
2. The transmitter uses the public key to encrypt sensitive information.
3. The transmitter sends the ciphertext generated from the sensitive information to the information receiver.
4. The receiver uses the private key to decrypt the ciphertext.

Before you start

Call the `CreateKey` operation to create an asymmetric CMK in KMS. Set the `KeySpec` parameter to a desired key type and the `KeyUsage` parameter to `ENCRYPT/DECRYPT`.

The following code demonstrates how to create an RSA encryption key:

```
$ aliyun kms CreateKey --KeySpec=RSA_2048 --KeyUsage=ENCRYPT/DECRYPT --ProtectionLevel=HSM
```

Obtain the public key

1. Call the `GetPublicKey` operation to obtain the public key of the asymmetric key pair.

```
$ aliyun kms GetPublicKey --KeyId=**** --KeyVersionId=****
```

Sample success responses:

```
{
  "RequestId": "82c383eb-c377-4mf6-bxx8-81hkc1g5g7ab",
  "KeyId": "****",
  "KeyVersionId": "****",
  "PublicKey": "PublicKey-DataBlob"
}
```

2. Save the public key to the `rsa_publickey.pub` file. `PublicKey-DataBlob` is a placeholder. You must replace it with the obtained public key.

```
$ echo PublicKey-DataBlob > rsa_publickey.pub
```

Use the public key to encrypt data

1. Create a sample plaintext file `plaintext-file.txt` that contains "this is plaintext".

```
echo "this is plaintext" > plaintext-file.txt
```

2. Use OpenSSL to encrypt the file and write the obtained binary ciphertext into the `plaintext-file.enc` file.

```
openssl pkeyutl -encrypt -in plaintext-file.txt \
-inkey rsa_publickey.pub -pubin \
-pkeyopt rsa_padding_mode:oaep \
-pkeyopt rsa_oaep_md:sha256 \
-pkeyopt rsa_mgf1_md:sha256 \
-out plaintext-file.enc
```

Call the KMS API to decrypt data

You must call the KMS API and use the private key to decrypt data.

1. Before you transmit the encrypted data over the network, encode it in Base64.

```
$ openssl base64 -in plaintext-file.enc
```

The following Base64-encoded ciphertext is returned:

```
5kdCB06HHeAwgfH9ARY4/9Nv5vlpQ94GXZcmaC9FE59Aw8v8RYdozT6ggSbyZbi+
8STKVq9402MEfmUDmwJLuu0qgAZsCe5wU4JWHh1y84Qn6HT068j0qOy5X2Hllrjs
fCdetgtMtVorSgb3bbERK2RV67nHWrdkecNbUaz+6ik4AlZxv2uWrV62eQ9yUBYm
Jb956LbqnfWdCFxUSHH/qB5QCnLpizvPmfNlZr653H4nF08gpZjnmlF4FJTU3i2
mGLzK4J3Rh/l7PQHivMdc4hSnXosg68QmMvdZBGLK9/cD9SYngPDiirU7z0q7Git
dleoyCAUDFyuQC6a+SqzA==
```

2. Pass the Base64-encoded ciphertext to KMS to decrypt data.

```
aliyun kms AsymmetricDecrypt \
--KeyId **** \
--KeyVersionId **** \
--Algorithm RSAES_OAEP_SHA_256 \
--CiphertextBlob 5kdCB06HHeAwgfH9ARY4/9Nv5vlpQ94GXZcmaC9FE59Aw8v8RYdozT6ggSbyZbi+8STKVq9402MEfmUD
mwJLuu0qgAZsCe5wU4JWHh1y84Qn6HT068j0qOy5X2HllrjsfCdetgtMtVorSgb3bbERK2RV67nHWrdkecNbUaz+6ik4AlZxv2
uWrV62eQ9yUBYmJb956LbqnfWdCFxUSHH/qB5QCnLpizvPmfNlZr653H4nF08gpZjnmlF4FJTU3i2mGLzK4J3Rh/l7PQHivMdc
4hSnXosg68QmMvdZBGLK9/cD9SYngPDiirU7z0q7GitdleoyCAUDFyuQC6a+SqzA==
```

The following results are returned:

```
{
  "KeyId": "*****",
  "KeyVersionId": "*****",
  "Plaintext": "dGhpcyBpcyBwbGFpbnRleHQgDQo=",
  "RequestId": "6be7a8e4-35b9-4549-ad05-c5b1b535a22c"
}
```

3. Decode the returned Base64-encoded plaintext in Base64.

```
echo dGhpcyBpcyBwbGFpbnRleHQgDQo= | openssl base64 -d
```

The following decrypted plaintext is returned:

```
this is plaintext
```

22.6.3. Generate and verify a digital signature by using an asymmetric CMK

This topic uses Alibaba Cloud CLI as an example to describe how to use an asymmetric customer master key (CMK) to generate and verify a digital signature. You can also perform this operation by using the KMS SDK.

Asymmetric encryption generally includes the following steps:

1. A signer sends a public key to a message receiver.
2. The signer uses the private key to sign data.
3. The signer sends the data and signature to the message receiver.
4. After receiving the data and signature, the message receiver uses the public key to verify the signature.

Before you start

Call the `CreateKey` operation to create an asymmetric CMK in KMS. Set the `KeySpec` parameter to a desired key type and the `KeyUsage` parameter to `SIGN/VERIFY`.

- Create an RSA signature key:

```
aliyun kms CreateKey --KeySpec=RSA_2048 --KeyUsage=SIGN/VERIFY --ProtectionLevel=HSM
```

- Create a NIST P-256 signature key:

```
aliyun kms CreateKey --KeySpec=EC_P256 --KeyUsage=SIGN/VERIFY --ProtectionLevel=HSM
```

- Create a secp256k1 signature key:

```
aliyun kms CreateKey --KeySpec=EC_P256K --KeyUsage=SIGN/VERIFY --ProtectionLevel=HSM
```

Preprocess signature: compute a message digest

Both RSA and ECC signature operations involve first computing the digest of an unsigned message and then signing the digest.

Note The algorithm used to obtain a message digest must match the algorithm used to call KMS to compute a signature. For example, the `ECDSA_SHA_256` signature algorithm must be used in conjunction with the SHA-256 digest algorithm. It does not support the SHA-384 digest algorithm.

The following example uses the SHA-256 digest algorithm.

1. Save the message "this is message" that needs to be signed into the file message-file.txt:

```
echo "this is message" > message-file.txt
```

2. Compute the SHA-256 digest of the message and save the binary digest to the file message-sha256.bin:

```
openssl dgst -sha256 -binary -out message-sha256.bin message-file.txt
```

Call KMS to compute the signature

You must call the KMS API to compute the signature of a message with the private key.

1. Before you transmit the message digest over the network, encode it in Base64.

```
openssl base64 -in message-sha256.bin
```

The following Base64 encoded digest is returned:

```
hRP2cuRFSIfEoUXCGuPyi7kZr18VCTZeVOTw0jbUB6w=
```

2. Pass the Base64 encoded digest to KMS to generate a signature.

Note The parameters passed and the results generated vary depending on key types and signature algorithms. Each signature result generated in the example is stored in a different file.

• RSASSA-PSS

For RSA keys, you can use the RSASSA-PSS signature algorithm and the SHA-256 digest algorithm to create a signature. Run the following command:

```
aliyun kms AsymmetricSign --KeyId=**** --KeyVersionId=**** \  
--Algorithm=RSA_PSS_SHA_256 --Digest=hRP2cu...  
{  
  "KeyId": "****",  
  "KeyVersionId": "****",  
  "Value": "J7xmdnZ...",  
  "RequestId": "70f78da9-c1b6-4119-9635-0ce4427cd424"  
}
```

Decode the signature value in Base64 and generate a binary signature. This signature is saved in the file rsa_pss_signature.bin:

```
echo J7xmdnZ... | openssl base64 -d -out rsa_pss_signature.bin
```

• RSASSA_PKCS1_V1_5

For RSA keys, you can use the RSASSA_PKCS1_V1_5 signature algorithm and the SHA-256 digest algorithm to create a signature. Run the following command:

```
aliyun kms AsymmetricSign --KeyId=**** --KeyVersionId=**** \  
--Algorithm=RSA_PKCS1_SHA_256 --Digest=hRP2cu...  
{  
  "KeyId": "****",  
  "KeyVersionId": "****",  
  "Value": "qreBkH/u...",  
  "RequestId": "4be57288-f477-4ecd-b7be-ad8688390fbc"  
}
```

Decode the signature value in Base64 and generate a binary signature. This signature is saved in the file `rsa_pkcs1_signature.bin`:

```
echo qreBkH/u... | openssl base64 -d -out rsa_pkcs1_signature.bin
```

- **NIST P-256**

For NIST curve P-256, you can use the ECDSA signature algorithm and the SHA-256 digest signature to create a signature. Run the following command:

```
aliyun kms AsymmetricSign --KeyId=**** --KeyVersionId=**** \
--Algorithm=ECDSA_SHA_256 --Digest=hrP2cu...
{
  "KeyId": "****",
  "KeyVersionId": "****",
  "Value": "MEYCIQD33Y98...",
  "RequestId": "472d789c-d4be-4271-96bb-367f7f0f8ec3"
}
```

Decode the signature value in Base64 and generate a binary signature. This signature is saved in the file `ec_p256_signature.bin`:

```
echo MEYCIQD33Y98... | openssl base64 -d -out ec_p256_signature.bin
```

- **secp256k1**

For SECG curve secp256k1, you can use the ECDSA signature algorithm and the SHA-256 digest algorithm to create a signature. Run the following command:

```
aliyun kms AsymmetricSign --KeyId=**** --KeyVersionId=**** \
--Algorithm=ECDSA_SHA_256 --Digest=hrP2cu...
{
  "KeyId": "****",
  "KeyVersionId": "****",
  "Value": "MEYCIQDWuul...",
  "RequestId": "fe41abed-91e7-4069-9f6b-0048f5bf4de5"
}
```

Decode the signature Value in Base64 and generate a binary signature. This signature is saved in the file `ec_p256k_signature.bin`:

```
echo MEYCIQDWuul... | openssl base64 -d -out ec_p256k_signature.bin
```

Obtain the public key

Obtain the public key of the created asymmetric key pair from the KMS. The preceding example assumes that:

- The public key of the RSA key pair is saved to the file `rsa_publickey.pub`.
- The public key of the NIST P-256 key pair is saved to the file `ec_p256_publickey.pub`.
- The public key of the secp256k1 key pair is saved to the file `ec_p256k_publickey.pub`.

Use the public key to verify the signature

Run the following command lines to verify the signature (the command varies depending on the algorithm used to generate the public key):

- **RSASSA-PSS**

```
openssl dgst \  
-verify rsa_publickey.pub \  
-sha256 \  
-sigopt rsa_padding_mode:pss \  
-sigopt rsa_pss_saltlen:-1 \  
-signature rsa_pss_signature.bin \  
message-file.txt
```

- **RSASSA_PKCS1_V1_5**

```
openssl dgst \  
-verify rsa_publickey.pub \  
-sha256 \  
-signature rsa_pkcs1_signature.bin \  
message-file.txt
```

- **NIST P-256**

```
openssl dgst \  
-verify ec_p256_publickey.pub \  
-sha256 \  
-signature ec_p256_signature.bin \  
message-file.txt
```

- **secp256k1**

```
openssl dgst \  
-verify ec_p256k_publickey.pub \  
-sha256 \  
-signature ec_p256k_signature.bin \  
message-file.txt
```

If the verification succeeds, the system displays the following message:

```
Verified OK
```

22.7. Use managed HSMs

22.7.1. Overview

Managed HSM is an important feature of KMS to enable easy access to certified HSMs.

An HSM is a highly secure hardware device that performs cryptographic operations and generates and stores keys. You can store the keys for your most sensitive Apsara Stack workloads and assets in HSMs that are managed on Apsara Stack.

High security assurance

- **Hardware protection**

Managed HSMs use secure hardware mechanisms to help you protect keys in KMS. The plaintext key material of CMKs is processed only inside HSMs for cryptographic operations. The material is kept within the hardware security boundary of HSMs.

- **Secure key generation**

Randomness is crucial to the encryption strength of keys. Managed HSMs use a random number generation algorithm to generate key material. The algorithm is secure and licensed and has high system entropy seeds. This protects keys from being recovered or predicted by attackers.

Ease of operation

HSM hardware is fully managed by Apsara Stack. This eliminates the costs otherwise incurred by the following hardware management operations:

- Hardware lifecycle management
- HSM cluster management
- High availability and scalability management
- System patching
- Most disaster recovery operations

Ease of integration

Native key management capabilities allow you to use the following features:

- Key version management
- Automatic key rotation
- Resource tag management
- Controlled authorization

These features enable rapid integration of your applications with managed HSMs, and the integration of other services such as ECS and ApsaraDB RDS with managed HSMs. This way, you can implement static cloud data encryption without R&D investment.

Key control

Managed HSMs allow you to better control encryption keys on the cloud and move the most sensitive computing tasks and assets to the cloud.

If you use both managed HSMs and Bring Your Own Key (BYOK), you can have full control over the following items:

- Generation modes of key material.
- Processing of key material: The key material that you import to a managed HSM can be destroyed but cannot be exported.
- Lifecycle of keys.
- Persistence of keys.

22.7.2. Use managed HSMs to create and use keys

This topic describes how to use managed HSMs to create and use keys.

Use a managed HSM to create a key

1. [Log on to the KMS console.](#)
2. On the Keys page, click **Create Key**.
3. On the **Create Key** page, select an organization from the **Organization** drop-down list. Then, the **Resource Set** and **Region** parameters are automatically set.
4. In the **Basic Settings** section, set the **Key Type** and **Key Purpose** parameters, and set the **Protection Level** parameter to **HSM**.
5. In the **Basic Settings** section, set the **Alias** and **Description** parameters.
6. In the **Advanced Settings** section, set the **Rotation Period** parameter.
 - **Disable**: The key is not automatically rotated.
 - **Enable**: The key is automatically rotated. You can select or customize an interval for rotation.

Note

- You can set this parameter only if the Key Type parameter is set to Aliyun_AES_256 or Aliyun_SM4.
- If the Key Material Source parameter is set to External, automatic rotation is not supported.

7. In the **Advanced Settings** section, set the **Key Material Source** parameter.

- **Key Management Service:** Use KMS to generate key material.
- **External:** Import key material from an external source.

Note

You can set this parameter only if the Key Type parameter is set to Aliyun_AES_256 or Aliyun_SM4.

8. Click **Submit**.

Import external keys to a managed HSM

You can import a key from user-managed key infrastructure to a managed HSM. The prerequisite is that you have set **Protection Level** to **HSM** when you create the external key. For more information, see [Import key material](#).

After you trigger the import, KMS performs the following operations:

- Calls the GetParametersForImport operation. During this process, KMS generates a key pair in a managed HSM to import the external key based on the **HSM** protection level and returns the public key of the key pair.
- Calls the ImportKeyMaterial operation. During this process, KMS imports the encrypted external key material to the managed HSM and then obtains the plaintext of the key material by using the key unwrapping mechanism of the managed HSM. The plaintext of the key material can no longer be exported.

Manage and use keys

You can apply all management and cryptographic features supported by KMS to keys created in managed HSMs. You can perform the following operations on these keys:

- Enable or disable keys
- Manage the lifecycle of keys
- Manage key aliases
- Manage key tags
- Call key-related API operations

Integrate managed HSMs with other Apsara Stack services

Keys in managed HSMs can be used to protect native data in other Apsara Stack services, such as ECS, ApsaraDB RDS, and OSS, by using the standard API of KMS. The prerequisite is that the Apsara Stack service supports SSE by using user-managed keys. To use this feature, you need only to configure a CMK in a managed HSM for the Apsara Stack service to implement SSE.

22.8. Key rotation

22.8.1. Overview

Keys are used to protect specific data. Therefore, the security of the data depends on the security of its keys. You can regularly rotate keys between key versions to improve key security and implement security policies and best practices for data protection.

Achieve security goals

You can regularly rotate keys to achieve the following goals:

- Reduce the amount of data encrypted based on each key

The security of a key is inversely proportional to the amount of data encrypted based on the key. This amount is usually defined by the total bytes of data or the total number of messages that are encrypted based on the same key. For example, National Institute of Standards and Technology (NIST) defines the secure lifecycle of a key in Galois/Counter Mode (GCM) as the total number of messages encrypted based on the key. Regular key rotation allows each key to remain secure and minimizes vulnerability to cryptanalytic attacks.

- Respond in advance to security events

In the early days of system design, key rotation was introduced as a routine operations and maintenance (O&M) method. This provides the system with a method to handle security events when they occur, and complies with the fail early, fail often principle of software engineering. If key rotation is not executed in the system until an emergency event has already occurred, the probability of system failure increases exponentially.

- Provide logical isolation of data

Data encrypted before a key rotation is isolated from data encrypted after the key rotation. The impact of key-related security events can be identified and preventive measures can be taken.

- Reduce the window of time to crack keys

Regular rotation of encryption keys ensures that you can control and reduce the window of time during which the key and its encrypted data are vulnerable to being cracked. The interval between rotation tasks during which attackers are able to crack the key is limited. This practice greatly increases the security of your data against cryptanalytic attacks.

Meet the requirements of regulatory compliance

Regular key rotation facilitates compliance with various regulations, which include but are not limited to the following regulations:

- Payment Card Industry Data Security Standard (PCI DSS)
- Cryptography-related industrial standards issued by State Cryptography Administration, such as GM/T 0051-2016
- Cryptography-related standards issued by NIST, such as NIST Publication 800-38D

22.8.2. Automatic key rotation

This topic shows you how to configure automatic rotation of CMKs in KMS.

Key versions

A CMK may have multiple key versions. Each key version represents an independently generated key. Key versions of the same CMK do not have a cryptographic relation to each other. KMS automatically rotates CMKs by generating new key versions.

Key versions are divided into the following types:

- Primary key versions
 - The primary key version of a CMK is an active encryption key. Each CMK has only one primary key version at a point in time.
 - When you call an encryption operation such as GenerateDataKey or Encrypt, KMS uses the primary key version of a specified CMK to encrypt plaintext.
 - You can call the DescribeKey operation to query the PrimaryKeyVersion attribute.
- Non-primary key versions
 - The non-primary key version of a CMK is an inactive encryption key. Each CMK can have zero to multiple non-primary key versions.
 - Each non-primary key version was a primary key version and acted as the active encryption key in the past.

- When a new primary key version is created, KMS does not delete or disable non-primary key versions, which are needed to decrypt data.

Note When you call an encryption operation, the primary key version of a specified CMK is used. When you call a decryption operation, the key version that was used for encryption is used.

You can generate a key version in one of the following ways:

- Create a CMK.

You can call the CreateKey operation to create a CMK. If you set the Origin parameter to *Aliyun_KMS*, KMS generates an initial key version and sets it as the primary key version.

- Execute an automatic rotation policy.

After you configure an automatic rotation policy for a CMK, KMS executes the policy on a regular basis to generate key versions.

Automatic rotation

Configure a key rotation policy.

When you call the CreateKey operation to create a CMK, you can specify an automatic rotation policy for the CMK. You can call the UpdateRotationPolicy operation to update the current automatic rotation policy. When you call the operations, you must configure the following parameters:

- EnableAutomaticRotation: specifies whether to enable automatic rotation.
- RotationInterval: the interval for automatic rotation.

You can call the DescribeKey operation to view the configured automatic rotation policy. The following parameters are returned:

- AutomaticRotation: indicates whether automatic rotation is enabled. The value *Disabled* indicates that automatic rotation is disabled. The value *Enabled* indicates that automatic rotation is enabled. The value *Suspended* indicates that KMS suspends the execution of automatic rotation although automatic rotation is enabled.
- RotationInterval: indicates the interval for automatic rotation.

Execute an automatic rotation policy.

When automatic rotation is enabled, KMS calculates the time of the next rotation by using the following formula:

$$\text{<NextRotationTime> = <LastRotationTime> + <RotationInterval>}$$

where:

- **LastRotationTime** : the time when the last key version was created. You can call the DescribeKey operation and check the **LastRotationDate** parameter to obtain the time.
- **NextRotationTime** : the time when KMS performs the next rotation task to create a key version. You can call the DescribeKey operation and check the **NextRotationDate** parameter to obtain the time.

Notice When you update the RotationInterval parameter of an automatic rotation policy, the value of NextRotationTime may be a point in time in the past. This does not affect the execution of the automatic rotation policy. A new key version is generated on a regular basis. If this situation occurs, KMS immediately executes the automatic rotation policy.

Impacts of CMK status on automatic rotation

A new key version can be created only for an enabled CMK. To enable a CMK, set the KeyState parameter to *Enabled*. Take note of the following items:

- If a CMK is in the *Disabled* or *Pending Deletion* state, do not call the UpdateRotationPolicy operation to update its automatic rotation policy.

- If a CMK enters the *Disabled* or *Pending Deletion* state after you enable automatic rotation for the CMK, KMS suspends the execution of automatic rotation. In this case, if you call the `DescribeKey` operation, the returned value of the `AutomaticRotation` parameter is *Suspended*. When the CMK enters the *Enabled* state again, its automatic rotation policy becomes active.

Limits

The following keys do not support multiple key versions:

- Service-managed keys: the default keys managed by KMS for specific cloud services. These keys belong to the users of cloud services and are used to provide basic encryption protection for user data.
- BYOK-based keys: the keys that you imported to KMS. For more information, see [Import and delete key material](#). The `Origin` attribute of these keys is *EXTERNAL*. KMS does not generate key material or initiate rotation tasks for these keys.

These two types of keys do not support version-based manual or automatic key rotation. A BYOK-based key does not have multiple versions in KMS. Users have strong control over the persistence and lifecycle of BYOK-based keys. This makes management of the keys difficult and error-prone. For example, you must have on-premises key management facilities, data must be synchronized between on-premises facilities and the cloud, and no grace period is provided for key material deletion on the cloud. The complexity of maintaining multiple versions of BYOK-based keys makes key management much more risky. In addition, both primary and non-primary key versions may become unavailable at different points in time. For example, if key versions are deleted by KMS or imported again when they expire, CMKs may not be synchronized and encrypted data may fail to be decrypted due to the use of an invalid CMK. This affects system integrity.

 **Note** For information about alternative solutions in the cases where version-based key rotation is not supported, see [Manual CMK rotation](#).

22.8.3. Manual CMK rotation

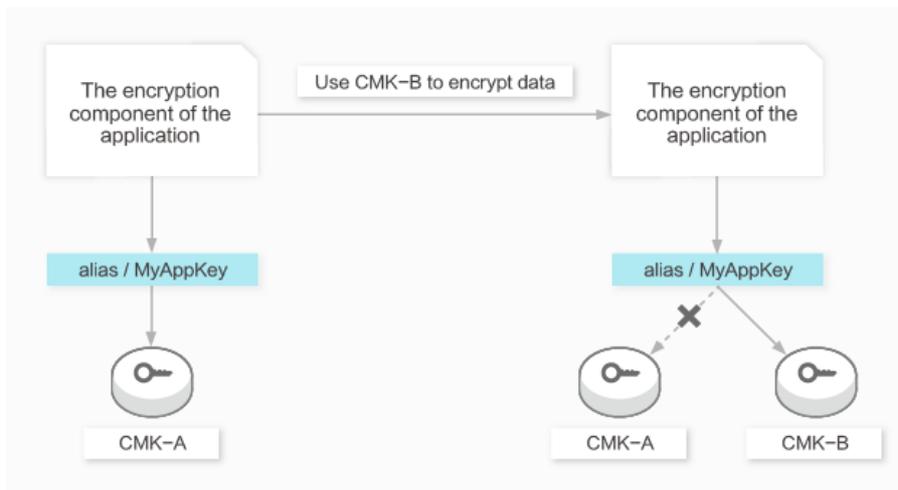
If your CMKs do not support version-based automatic rotation, you can manually rotate the CMKs based on your scenario. This is an alternative solution that can be used regardless of whether automatic CMK rotation is supported.

Scenario of custom data encryption

On-premises or cloud applications can call KMS API operations to implement custom data encryption. Examples:

- Encrypt sensitive data such as ID card numbers, credit card numbers, and home addresses before you write it to databases.
- Encrypt data on a client before you upload the data to Object Storage Service (OSS).
- Encrypt application configuration files that contain sensitive data and service configuration files such as SSL certificates and private keys.

You can use the key alias feature of KMS to rotate encryption keys within applications. The ID and alias of CMKs are not required when you call the Decrypt operation.



In this scenario, the rotation process consists of the following steps:

1. Initial configuration
 - i. The administrator creates a CMK, whose ID is CMK-A.
 - ii. The administrator binds the alias/MyAppKey alias to CMK-A.
 - iii. When the application encryption module calls the Encrypt operation, the application encryption module sets the value of the KeyId parameter to *alias/MyAppKey*. KMS finds that the alias/MyAppKey alias is bound to CMK-A and then uses CMK-A to encrypt data.
 - iv. When the application decryption module calls the Decrypt operation, the application decryption module does not set the KeyId parameter. KMS uses the CMK that is used to encrypt the data to decrypt the data.
2. Manual CMK rotation
 - i. The administrator creates a CMK, whose ID is CMK-B.
 - ii. The administrator calls the UpdateAlias operation to bind the alias/MyAppKey alias to CMK-B.
 - iii. When the application decryption module calls the Encrypt operation, the application decryption module sets the value of the KeyId parameter to *alias/MyAppKey*. KMS finds that the alias/MyAppKey alias is bound to CMK-B and uses CMK-B to encrypt data.
 - iv. When the application decryption module calls the Decrypt operation, the application decryption module does not set the KeyId parameter. KMS uses the CMK that is used to encrypt the data to decrypt the data.

Scenario of server-side encryption by cloud services

Cloud services can implement server-side encryption by calling KMS API operations. In this scenario, the following situations may occur from the perspective of key rotation:

- Automatic rotation policies configured in KMS affect the server-side encryption of cloud services.

Cause: After encryption based on the specified CMK is configured for cloud services, these services call the GenerateDataKey operation of KMS to generate new data keys. When KMS generates a new primary key version, cloud services use the new version to encrypt the newly generated data keys. A typical example of such cloud services is OSS. In this case, if you want to enable automatic key rotation, you must not use service-managed keys or BYOK-based keys imported to KMS because these keys do not support automatic rotation.
- Automatic rotation policies configured in KMS do not affect server-side encryption of cloud services.

Cause: After encryption based on the specified CMK is configured for cloud services, the services call the GenerateDataKey operation of KMS only once to generate keys to encrypt specific resources. When KMS generates a new primary key version, cloud services do not use this version. For example, Elastic Compute Service (ECS) calls the GenerateDataKey operation of KMS once to generate a volume encryption key. This key is not updated after it is created.

If automatic rotation policies configured in KMS do not affect server-side encryption of cloud services or you want to rotate data keys when BYOK-based keys are used, you can modify configurations and copy data to achieve the same effect as key rotation. These methods depend on the features of different cloud services. For more information, see the documents of related cloud services.

23. Log Service

23.1. What is Log Service?

Log Service (SLS) is a one-stop logging service developed by Alibaba Cloud that is widely used by Alibaba Group in big data scenarios. You can use Log Service to collect, query, and consume log data.

Without the need to invest in in-house data collection and processing systems, this enables you to focus on your business, improving business efficiency and helping your business to expand.

Log Service provides the following features:

- **Log collection:** Log Service allows you to collect events, binary logs, and text logs in real time by using multiple methods, such as Logtail and JavaScript.
- **Query and analysis:** Log Service allows you to query and analyze the collected log data and view analysis results on charts and dashboards.
- **Status alert:** Log Service can automatically run query statements at regular intervals after you create an alert task. If the query results meet the conditions of the alert task, Log Service sends an alert to the specified recipients in real time.
- **Real-time consumption:** Log Service provides real-time consumption interfaces through which log consumers can consume log data.

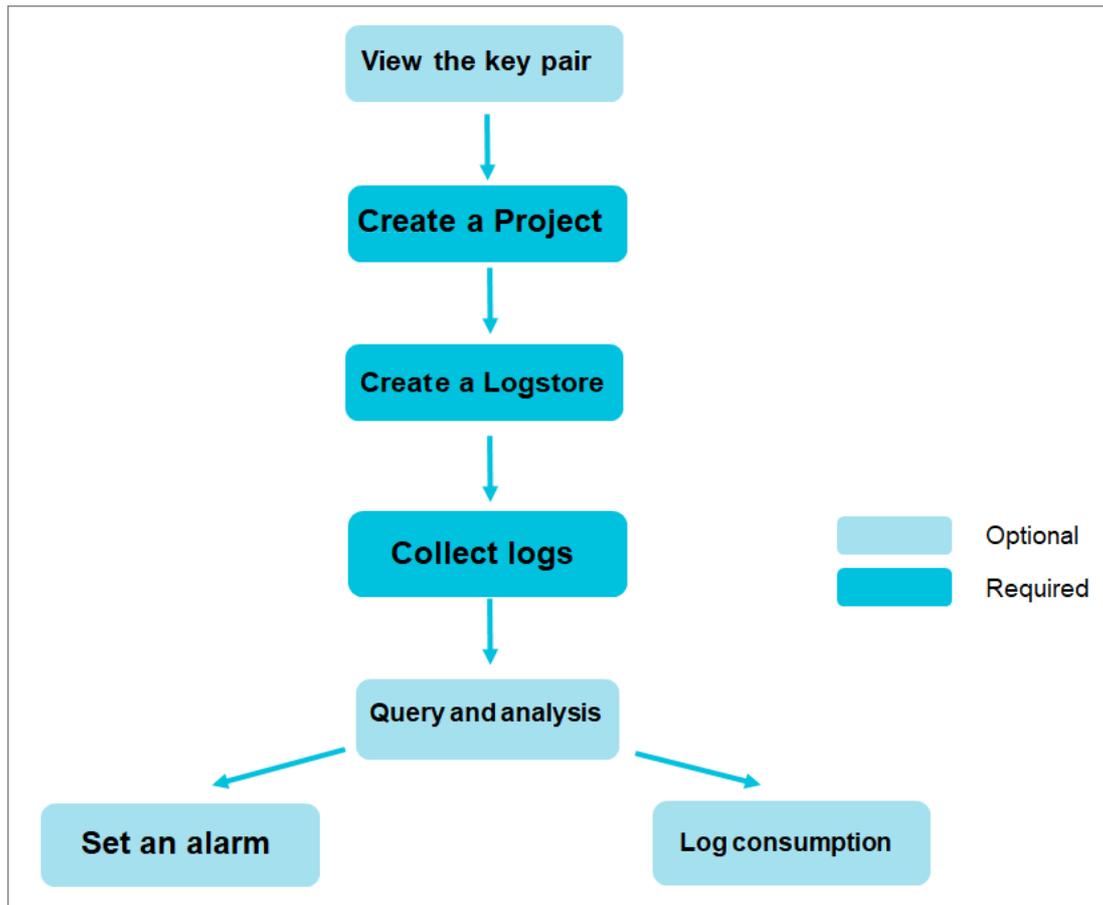
23.2. Quick start

23.2.1. Procedure

This topic provides the basic procedure to use Log Service. You can use this procedure to create projects, create Logstores, and collect log data.

The following figure shows the [Procedure](#).

Procedure



1. Optional. Obtain an AccessKey pair.

Before you can use Log Service through APIs or SDKs, you must have an AccessKey pair.

2. [Create a project](#).

Create a project in a specified region and add a description.

3. [Create a Logstore](#).

Create a Logstore for the project and specify the number of shards.

4. [Collect text logs](#)

Select a method to collect log data based on your business requirements. Text log collection is used as an example.

5. [Enable the index feature and configure indexes for a Logstore](#), and query and analyze logs.

Log Service supports [real-time log query](#) and [analysis](#). After you enable the indexing feature, you can query and analyze logs and configure [Overview](#) and [dashboards](#).

6. [Configure alerts](#).

Log Service allows you to configure alerts based on log query results. Then, Log Service sends alerts by using multiple methods, such as a custom webhook.

7. [Consume logs in real time](#).

Log Service allows you to consume logs by using multiple methods, such as a [Spark Streaming client](#), [Storm spout](#), and [Flink connector](#).

23.2.2. Log on to the Log Service console

This topic describes how to log on to the Log Service console.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

Note When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top navigation bar, choose **Products > Log Service**.
5. On the page that appears, select the organization and region, and then click **SLS**. The home page of the Log Service console is displayed.

23.2.3. Obtain an AccessKey pair

An AccessKey pair consists of an AccessKey ID and an AccessKey secret. The AccessKey pair is used to implement symmetric encryption to verify the identity of the requester. The AccessKey ID is used to identify a user. The AccessKey secret is used to encrypt the signature string. This topic describes how to obtain an AccessKey pair.

Prerequisites

Only the operation administrators or level-1 organization administrators can obtain the AccessKey pair of an organization.

Context

To call Apsara Uni-manager and cloud service APIs, we recommend that you use the AccessKey pair of a personal account. If you use the AccessKey pair of a personal account, you must configure header parameters as described in the following table for access control.

| Parameter | Description |
|-----------------------|--|
| x-acs-regionid | The region ID, such as cn-hangzhou-* |
| x-acs-organizationid | The ID of the organization in the Apsara Uni-manager Management Console. |
| x-acs-resourcegroupid | The ID of the resource set in the Apsara Uni-manager Management Console. |

| Parameter | Description |
|-------------------|---|
| x-acsc-instanceid | The ID of the instance on which you want to perform operations. |

 **Warning** The AccessKey pairs of personal accounts are under control of the Apsara Uni-manager permission system. AccessKey pairs of organization accounts have higher permissions. For security purposes, organization operations must be approved by administrators.

Obtain the AccessKey pair of a personal account

To obtain the AccessKey pair of a personal account, perform the following operations:

1. Log on to the Apsara Uni-manager Management Console.
2. In the upper-right corner of the homepage, move the pointer over the profile picture and click **Personal Information**.
3. In the **Apsara Stack AccessKey Pair** section, view your AccessKey pair.



| Region | AccessKey ID | AccessKey Secret |
|-----------|--------------|--------------------------|
| cn-...d01 | ... | ... Show |

 **Note** The AccessKey pair consists of an AccessKey ID and an AccessKey secret. AccessKey pairs allow you to access Apsara Stack resources with full permissions for your account. You must keep your AccessKey pair confidential.

Obtain the AccessKey pair of an organization

To obtain the AccessKey pair of an organization, perform the following operations:

1. Log on to the Apsara Uni-manager Management Console as an administrator.
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Organizations**.
4. In the organization navigation tree, click an organization name.
5. In the Current Organization section, click **Management Accesskey**.
6. In the Management AccessKey, view the AccessKey pair of the organization.

 **Note** An AccessKey pair is automatically allocated to each level-1 organization. Subordinate organizations use the same AccessKey pair of their level-1 organization.

23.2.4. Manage projects

This topic describes how to create, modify, and delete projects in the Log Service console.

Context

A project in Log Service is a resource management unit. The resources in each project are isolated from resources in other projects. We recommend that you store the log data of different applications in dedicated projects. You can manage Logstores, Logtail configurations, log sources, log data, and machine groups in a project. Each project provides an endpoint for you to access the resources.

A project provides the following features:

- Allows you to store log data from different sources in different Logstores of a project. You can use Log Service to collect log data from multiple sources such as business projects, products, and environments. You can then store the log data of each source in a separate Logstore. This simplifies the downstream processes such as the consuming, exporting, and indexing of log data. In addition, you can manage access permissions at the project level.
- Provides an endpoint for you to access the resources in the project. Log Service allocates an exclusive endpoint to each project. You can use the endpoint to read, write, and manage the log data in the project.

Create a project

Note

- You can create a project only by using the Log Service console.
- You can create up to 50 projects under each Apsara Stack tenant account.

1. [Log on to the Log Service console](#).
2. In the **Projects** section, click **Create Project**.
3. Set the parameters based on your requirements. The following table describes the parameters.

| Parameter | Description |
|--------------|---|
| Project Name | <p>The name of the project. The project name must be unique across all regions. Use the following naming conventions:</p> <ul style="list-style-type: none"> ◦ The name can contain lowercase letters, digits, and hyphens (-). ◦ The name must start and end with a lowercase letter or digit. ◦ The name must be 3 to 63 characters in length. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note After a project is created, its name cannot be modified.</p> </div> |
| Description | <p>The description of the project. After the project is created, the description is displayed in the Projects section. If you need to modify the description after the project is created, find the project in the Projects section, and click Edit in the Actions column. The description must be 0 to 64 characters in length and cannot contain the following characters: <code><>'\"</code>.</p> |
| Region | <p>The region to which the project belongs. We recommend that you select a region that is closer to the log source.</p> <p>After a project is created, its region cannot be modified. This means that projects cannot be migrated across regions.</p> |

4. Click **OK**.

Modify the description of a project

To modify the description of a project, perform the following steps:

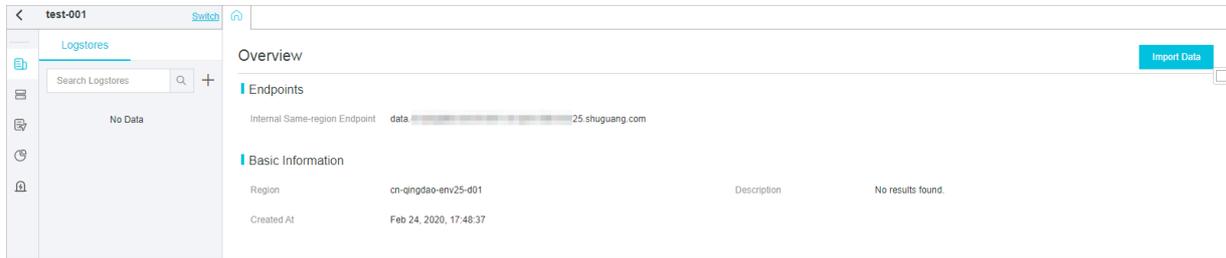
1. In the **Projects** section, find the project.
2. In the **Actions** column, click **Edit**.
3. In the **Modify Project** dialog box, modify the description of the project.

 **Note** You cannot modify the project name or region.

4. Click **OK**.

View the information of a project

To view the information of a project, click the project name in the **Projects** section. On the **Overview** page, you can view the project information such as the endpoint and region.



Delete a project

To delete a project, perform the following steps:

Warning After you delete a project, all logs and configurations in the project are deleted and cannot be restored.

1. In the **Projects** section, find the project.
2. In the **Actions** column, click **Delete**.
3. In the dialog box that appears, select a reason for deletion.
If you select **Other issues**, enter the reason in the text box.

Delete Project ×

 You cannot restore the project data after the project is deleted. Are you sure you want to delete the project?

Project Name: test-001

Reason for Deletion

- The project name is incorrect.
- The region of the project is incorrect.
- Business issue. Log analysis is no longer required.
- The data in the project is for test and must be cleared.
- Cost issue.
- Do not know how to use Log Service.
- Cannot import logs to the project.
- Other issues.

4. Click OK.

23.2.5. Manage Logstores

This topic describes how to create, modify, and delete a Logstore in the Log Service console. A Logstore is a collection of resources inside a project. The log data in a Logstore is collected from the same source.

Context

You can create multiple Logstores in a project. We recommend that you create a Logstore for each type of application log.

Logstores provide the following features:

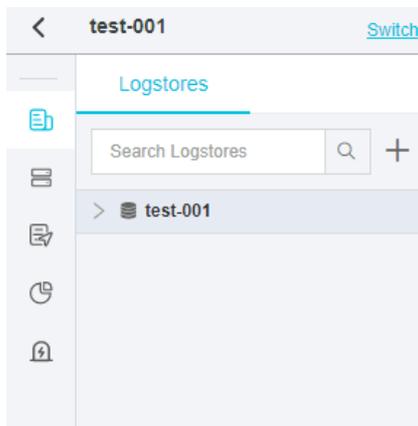
- Real-time log collection
- Log storage and real-time log consumption
- Log indexing and real-time log query

Create a Logstore

 **Note** You can create a maximum of 100 Logstores in each project.

1. [Log on to the Log Service console.](#)

- In the **Projects** section, click the name of a project.
- On the page that appears, click next to the search box.



- In the **Create Logstore** pane, set the parameters and click **OK**.

| Parameter | Description |
|-----------------------|---|
| Logstore Name | <p>The name of the Logstore. The Logstore name must be unique in the project to which the Logstore belongs.</p> <ul style="list-style-type: none"> The name can contain lowercase letters, digits, hyphens (-), and underscores (_). The name must start and end with a lowercase letter or digit. The name must be 3 to 63 characters in length. <p>Note After a Logstore is created, its name cannot be modified.</p> |
| WebTracking | Specifies whether to enable the WebTracking feature for the Logstore. You can use WebTracking to collect the log data of HTML websites, HTML5 websites, iOS apps, or Android apps and forward the data to Log Service. This feature is not enabled by default. |
| Permanent Storage | Specifies whether to permanently store the log data in the Logstore. If you disable this feature, you must specify a retention period for log data. |
| Data Retention Period | The duration for which log data is stored in the Logstore after the log data is collected. Unit: days. Valid values: 1 to 3000. When this period expires, the log data is deleted. |
| Shards | The number of shards in the Logstore. You can divide a Logstore into 1 to 10 shards. |
| Automatic Sharding | Specifies whether to enable the automatic sharding feature. This feature is not enabled by default. If you enable this feature, Log Service automatically splits shards in the Logstore when the data transfer exceeds the capacity of the existing shards. |
| Maximum Shards | The maximum number of shards. This parameter is required if you enable the automatic sharding feature. Maximum value: 64. |

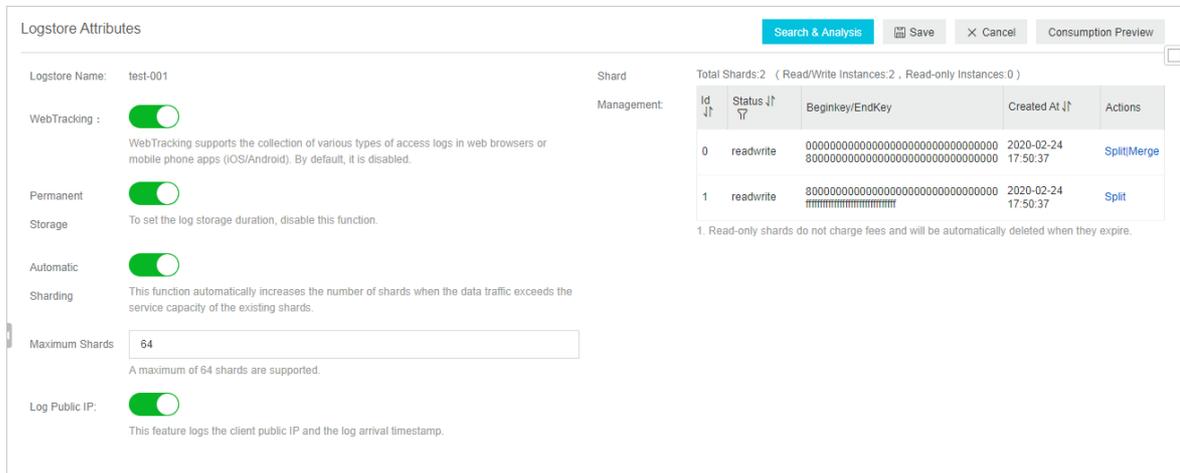
| Parameter | Description |
|---------------|---|
| Log Public IP | <p>Specifies whether to add the following information to the Tag field of each collected log entry:</p> <ul style="list-style-type: none"> __client_ip__ : the public IP address of the log source. __receive_time__ : the time when Log Service receives the log entry. The value is a UNIX timestamp. |

Modify the configurations of a Logstore

To modify the configurations of a Logstore, perform the following steps:

1. In the **Projects** section, click the name of the project to which the Logstore belongs.
2. On the page that appears, click the  icon next to the Logstore, and select **Modify** from the short cut menu.
3. In the upper-right corner of the **Logstore Attributes** page, click **Modify**.

You can modify the **Data Retention Period**, **WebTracking**, **Automatic Sharding**, **Maximum Shards**, and **Log Public IP** parameters. You can also split or merge the existing shards.



4. Click **Save**.

Delete a Logstore

To delete a Logstore, perform the following steps:

Note

- After you delete a Logstore, the log data in the Logstore is deleted and cannot be restored.
- Before you can delete a Logstore, you must delete the Logtail configurations that are associated with the Logstore.

1. In the **Projects** section, click the name of the project to which the Logstore belongs.
2. On the page that appears, click the  icon next to the Logstore, and select **Delete** from the short cut menu.
3. In the message that appears, click **OK**.

Delete : test-001



You cannot restore the data that has been deleted. Are you sure to delete the data?

OK

Cancel

23.2.6. Manage shards

This topic describes how to split, merge, and delete shards in the Log Service console. Logs are stored on shards in a Logstore. Each Logstore can have multiple shards. When you create a Logstore, you must specify the number of shards in the Logstore. After a Logstore is created, you can split or merge the shards.

Hash key

Log Service uses 128-bit MD5 hashes as the hash key of a Logstore. The entire MD5 hash range is [00000000000000000000000000000000,ffffffffffffffffffffffffffffffff]. The hash key range of a Logstore falls within the entire MD5 hash range. When you create a Logstore, you must specify the number (N) of shards in the Logstore. The hash key range of the Logstore is evenly divided into N parts. Each part is assigned to a shard.

The hash key range of a shard is a left-closed and right-open interval that is specified by the following parameters:

- BeginKey: the start of the hash key range. The value of this parameter is included in the range.
- EndKey: the end of the hash key range. The value of this parameter is excluded from the range.

If you split a shard, the hash key range of the shard is evenly split. If you merge two shards, the hash key ranges of the shards are also merged. A hash key range determines the scope of a shard. When you push log data to a Logstore, you can specify a hash key for the log data. Log Service then writes the log data to the shard whose hash key range includes the specified hash key. This is called the hash key mode. If you do not specify a hash key for log data, the load balancing mode is used and Log Service writes the log data to a random available shard. However, when you pull log data from a Logstore, you must specify the shard where the log data is stored.

For example, a Logstore is divided into four shards and the hash key range of the Logstore is [00,FF). [Example shards](#) lists the hash key range of each shard.

Example shards

| Shard | Hash key range |
|--------|----------------|
| Shard0 | [00,40) |
| Shard1 | [40,80) |
| Shard2 | [80,C0) |
| Shard3 | [C0,FF) |

If you set the hash key of log data to 5F, Log Service writes the log data to shard 1 because the hash key range of shard 1 includes 5F. If you set the hash key to 8C, the log data is written to shard 2 because the hash key range of shard 2 includes 8C.

Read/write capacity

Each shard provides an identical read/write capacity. Therefore, the read/write capacity of a Logstore depends on the number of shards in the Logstore. We recommend that you adjust the capacity of a Logstore based on the data traffic. For a Logstore, if the data traffic exceeds the read/write capacity, you can split shards to increase the Logstore capacity. If the data traffic is much less than the read/write capacity, you can merge shards to reduce the Logstore capacity and save costs.

For example, a Logstore consists of two read/write shards and the shards provide a maximum write capacity of 10 MB/s. If log data is written to the Logstore at a rate of 14 MB/s, we recommend that you split one of the shards into two shards. However, if log data is written at a rate of 3 MB/s, you can merge the two shards because the capacity of one shard already meets the read/write requirements.

Note

- If an API operation that writes data to a Logstore constantly returns 403 or 500 errors, you can check the data traffic metrics that are provided by Log Service and determine whether to split shards.
- If the data traffic of a Logstore exceeds the read/write capacity of the Logstore, Log Service provides the best possible service but does not guarantee the service quality.

Shard status

A shard can be in one of the following states:

- Read/write
- Read-only

After a shard is created, the default status of the shard is read/write. If you split or merge shards, the status of the original shards changes to read-only and the new shards are in the read/write state. You can write data to and read data from a read/write shard. However, you can only read data from a read-only shard and cannot write data to the shard.

If you need to split a shard in a Logstore, you must specify the ID of the shard and an MD5 hash. The shard must be in the read/write state. The MD5 hash must be greater than the value of the BeginKey parameter of the shard and less than the value of the EndKey parameter of the shard. After the shard is split, the Logstore has two more shards. The status of the original shard changes from read/write to read-only. You can consume the log data in the original shard but cannot write log data to the shard. The new shards are in the read/write state and are listed below the original shard. The hash key ranges of the new shards cover that of the original shard.

If you need to merge shards in a Logstore, you must specify a read/write shard. The shard cannot be the last read/write shard in the shard list. Log Service finds the shard whose hash key range follows the hash key range of the specified shard, and merges the two shards into a new shard. The status of the original shards changes from read/write to read-only. You can consume the log data in the original shards but cannot write log data to the shards. The new shard is in the read/write state. The hash key range of the new shard covers those of the original shards.

You can perform the following operations on shards in the Log Service console:

- Split a shard.
- Merge shards.

Split a shard

Each shard provides a write capacity of 5 MB/s and a read capacity of 10 MB/s. For a Logstore, if the data traffic exceeds the total read/write capacity of existing shards, we recommend that you split shards to increase the capacity.

1. [Log on to the Log Service console](#).
2. In the **Projects** section, click the name of a project.
3. On the page that appears, click the  icon next to the Logstore, and select **Modify** from the shortcut menu.
4. In the upper-right corner of the **Logstore Attributes** page, click **Modify**.

Logtail retries log collection and locally caches logs to ensure data security.

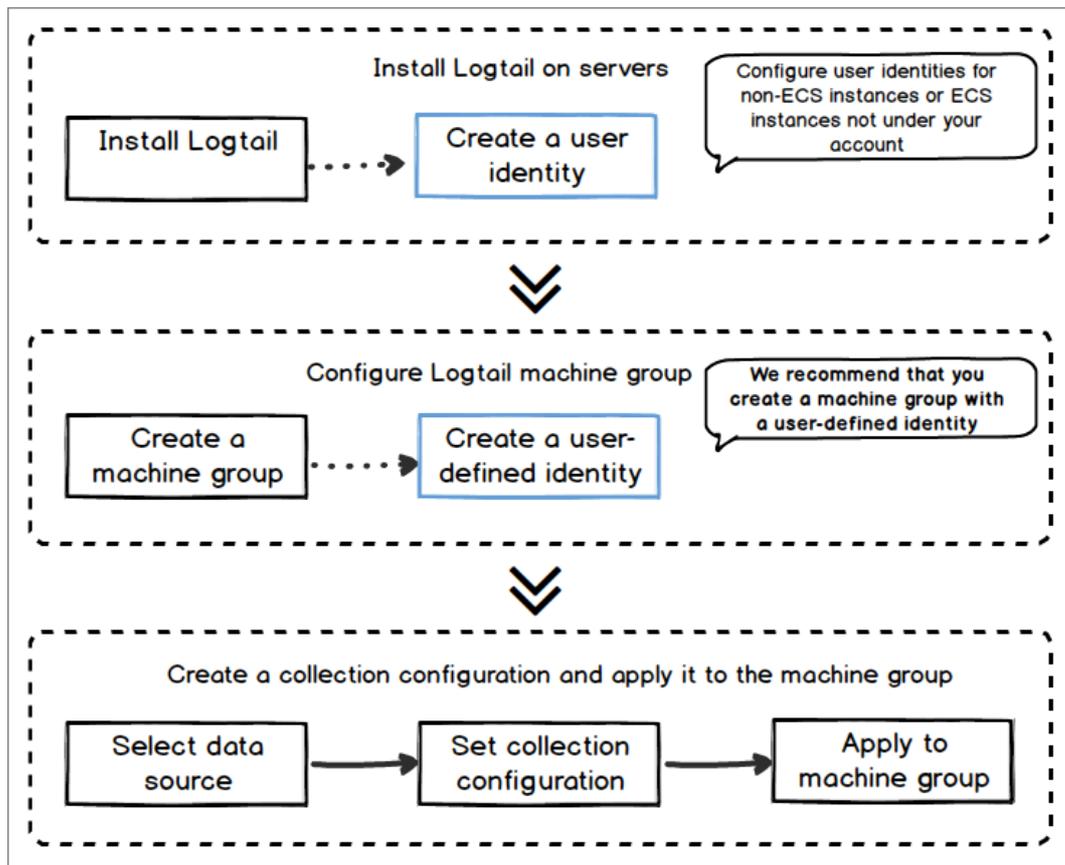
- Provides centralized management based on Log Service. After you install Logtail on the servers from which you want to collect logs, you can configure these servers and the collection method in the Log Service console. You do not need to log on to the servers.
- Provides a comprehensive self-protection mechanism. To minimize the impact of Logtail on the performance of the related servers, Logtail limits the usage of CPU, memory, and network resources.

Processing capabilities and limits

For more information, see [Limits](#).

Configuration process

Configuration process



To collect logs from servers by using Logtail, follow these steps:

1. Install Logtail. For more information about how to install Logtail on a server from which you want to collect logs, see [Install Logtail in Linux](#) and [Install Logtail in Windows](#).
2. Log Service use server groups to manage all servers from which you want to collect logs by using Logtail. Log Service allows you to define server groups by using IP addresses or custom identifiers. You can create a server group as prompted when you apply Logtail configurations to server groups.
3. Create a Logtail configuration and apply it to the server group. For more information about how to create a Logtail configuration, see [Configure text log collection](#).

After the preceding process is complete, logs on the server are automatically collected and sent to the selected Logstore. However, historical logs are not collected. You can use the Log Service console, SDKs, or APIs to query these logs. Log Service allows you to view the status of log collection and check whether errors occur.

For more information, see [Collect logs by using Logtail](#).

Containers

- For information about Alibaba Cloud Container Service for Kubernetes or user-created Kubernetes clusters, see [Collect Kubernetes logs](#).
- For information about other user-created Docker clusters, see [Collect standard Docker logs](#).

Terms

- **Server group:** A server group contains one or more servers from which logs of a specific type are collected. You can apply Logtail configurations to a server group. This enables Log Service to collect logs from all servers in the server group. You can use the Log Service console to manage a server group. For example, you can create, delete, add, or remove a server. Each server group can contain different versions of Windows servers or Linux servers.
- **Logtail:** Logtail is the agent that collects logs from the servers on which Logtail runs. For more information, see [Install Logtail in Linux](#) and [Install Logtail in Windows](#). After you install Logtail on a server, you must create a Logtail configuration and apply it to the server group to which the server belongs.
 - In Linux, Logtail is installed in the `/usr/local/ilogtail` directory. Logtail initiates two separate processes whose names start with `ilogtail`. One is a collection process and the other is a daemon process. The program running log is `/usr/local/ilogtail/ilogtail.LOG`.
 - In Windows, Logtail is installed in the `C:\Program Files\Alibaba\Logtail` (for 32-bit systems) or `C:\Program Files (x86)\Alibaba\Logtail` (for 64-bit systems) directory. You can choose Administrative Tools > Services to view the two Windows services generated from Logtail. One is `LogtailWorker` (log collection process) and the other is `LogtailDaemon`. The program running log is `logtail_*.log` in the installation directory.
- **Logtail configuration:** A Logtail configuration is a set of policies that are used by Logtail to collect logs. You can specify Logtail parameters such as the data source and collection mode. This allows you to customize log collection policies for all servers in a server group. A Logtail configuration determines how to collect a type of logs from a server, parse the logs, and send them to a specified Logstore. You can create a Logtail configuration for a Logstore in the Log Service console. This enables the Logstore to receive logs that are collected by using this Logtail configuration.

Features

Logtail provides the following features:

| Feature | Description |
|------------------------------|--|
| Real-time log collection | <p>Logtail dynamically monitors log files and reads and parses incremental logs in real time. In most cases, logs are sent to Log Service within 3 seconds after they are generated.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note Logtail does not collect historical data. If Logtail reads a log later than 12 hours after the log was generated, Logtail drops the log.</p> </div> |
| Automatic log rotation | <p>Some applications rotate log files based on the file size or date. In the rotation process, the original log files are renamed and empty log files are created. For example, files such as <code>app.LOG.1</code> and <code>app.LOG.2</code> are generated for the <code>app.LOG</code> file after log rotation. You can specify the file (for example, <code>app.LOG</code>) to which collected logs are written. Logtail monitors the log rotation process to ensure that no logs are lost.</p> |
| Multiple data sources | <p>Logtail can collect text logs, syslogs, HTTP logs, and MySQL binlogs.</p> |
| Automatic exception handling | <p>If data transmission fails due to exceptions such as Log Service errors, network errors, or quota exhaustion, Logtail retries log collection based on the specific scenario. If the retry fails, Logtail writes the data to the local cache and resends the data after the exception no longer exists.</p> |

| Feature | Description |
|---|---|
| Flexible collection policy configuration | <p>You can create a Logtail configuration to specify how logs are collected from an ECS instance. You can select log directories and files by using exact match or wildcard match based on actual scenarios. You can also customize the extraction method of collected logs and the names of extracted fields. Log Service allows you to extract logs by using regular expressions.</p> <p>The log data models of Log Service require that each log has a precise timestamp. Logtail provides custom log time formats. This allows you to extract the required timestamps from log data of different formats.</p> |
| Automatic synchronization of Logtail configurations | After you create or update a Logtail configuration in the Log Service console, Logtail receives and validates the configuration within 3 minutes. No data loss occurs during the configuration update process. |
| Automatic upgrade | After you install Logtail on a server, Log Service manages the automatic upgrade of Logtail without manual intervention. No data loss occurs during the Logtail upgrade process. |
| Status monitoring | Logtail monitors its consumption of CPU and memory in real time. This prevents Logtail from excessively consuming your resources. If the resource consumption exceeds the limit, Logtail restarts to avoid affecting other services on the server. Logtail limits network traffic to avoid excessive bandwidth consumption. |
| Data transmission with a signature | <p>Logtail obtains your Alibaba Cloud AccessKey pair and uses it to sign all log data packets before they are sent. This prevents data tampering during transmission.</p> <p> Note Logtail obtains your Alibaba Cloud AccessKey pair over HTTPS to ensure the security of your AccessKey pair.</p> |

Data collection reliability

During data collection, Logtail stores the collected checkpoints to the local directory on a regular basis. If an exception (for example, an unexpected server shut down or a process crash) occurs, Logtail restarts and then collects data from the last recorded checkpoint to prevent data loss. Logtail runs based on the [resource limits](#) specified in the configuration file. If the usage of a resource exceeds the limit for more than 5 minutes, Logtail restarts. After the restart, duplicate data may be generated.

Logtail uses multiple internal mechanisms to improve log collection reliability. However, logs may be lost in the following conditions:

- Logtail is not running but logs are rotated multiple times.
- The log rotation rate is high, for example, one rotation per second.
- The log collection rate is lower than the log generation rate for a long period of time.

23.3.1.1.2. Log collection process of Logtail

This topic describes the process that Logtail uses to collect server logs. The process consists of the following steps: monitor files, read files, process logs, filter logs, aggregate logs, and send logs.

 **Note** After the Logtail configuration is applied to a machine group, unmodified logs on the servers in the machine group are considered as historical logs. Logtail does not collect historical logs in normal running mode. If you want to collect historical logs, see [Import historical logs](#).

Monitor files

After you install Logtail on a server and create a Logtail configuration based on a data source, the Logtail configuration is delivered from Log Service to Logtail in real time. Then, Logtail starts to monitor files based on the Logtail configuration.

1. Logtail scans log directories and files that comply with the specified file naming rules based on the specified log path and maximum monitoring directory depth.

Logtail registers event monitoring and periodical polling for the directory from which Logtail collect logs. In Linux, `Inotify` is used. In Windows, `ReadDirectoryChangesW` is used. This method ensures the timeliness and stability of log collection.

2. If the compliant log files in the specified directory are not modified after the configuration is applied, Logtail does not collect these files. If modification events are generated for log files, Logtail triggers the collection process and reads these files.

Read files

After Logtail detects that a log file has been updated, Logtail reads the log file.

- If Logtail reads the log file for the first time, it checks the file size.
 - If the file size is less than 1 MB, Logtail reads the file from the beginning.
 - If the file size is greater than 1 MB, Logtail reads the file from the last 1 MB of data.
- If Logtail has read the file before, Logtail reads the file from the last checkpoint.
- Logtail can read up to 512 KB of data at a time. Therefore, you must limit the log size to 512 KB.



Notice If you have changed the system time on your server, you must manually restart Logtail. Otherwise, the log time is incorrect and logs are dropped.

Process logs

When Logtail reads a log, it divides the log into multiple lines, parses the log, and sets the time field of the log.

- Divide each log into multiple lines

If you specifies **Regex to Match First Line Only** in the Logtail configuration, the log data read by Logtail at one time is divided into multiple lines based on the specified beginning of the line. If the beginning of the line is not specified, each data block is processed as a log.

- Parse each log

Logtail parses each log based on the Logtail configuration, such as regular expressions, delimiters, and JSON.



Notice A complicated regular expression may lead to high CPU usage. Therefore, we recommend that you use an efficient regular expression.

- Handle parsing failures

Logtail determines how to handle parsing failures based on whether the **Drop Failed to Parse Logs** switch is turned on in the Logtail configuration.

- If the **Drop Failed to Parse Logs** is turned on, Logtail drops the logs that fail to be parsed and reports an error.
- If the **feature** switch is turned off, Logtail uploads the logs that fail to be parsed. In these logs, the key is set to `raw_log` and the value is set to the log content.

- Set the time field of a log

- If the time field of the log is not specified, the log time is the current parsing time.

- If the time field of the log is specified, the following operations are performed:
 - The log time is extracted from the parsed log fields if the difference between the log time and the current time is less than 12 hours.
 - The log is dropped and an error is reported if the difference between the log time and the current time is greater than 12 hours.

Filter logs

After Logtail process logs, it filters the logs based on the filter configuration.

- If the **Filter Configuration** is not specified, Logtail does not filter logs and proceed with the next step.
- If the **Filter Configuration** is specified, Logtail traverses and verifies all the fields of each log.
 - Logtail collects a log if the log matches the filter configuration.
 - Logtail does not collect a log if the log does not match the filter configuration.

Aggregate logs

Logtail sends the logs that match the filter configuration to Log Service. To reduce the number of network requests, Logtail caches the processed and filtered logs for a period of time. Then, Logtail aggregates and packages these logs before sending them to Log Service.

If one of the following conditions is met during caching, logs are immediately packaged and sent to Log Service.

- The log aggregation period exceeds three seconds.
- The number of aggregated logs exceeds 4,096.
- The total size of aggregated logs exceeds 512 KB.

Send logs

Logtail sends the aggregated logs to Log Service. You can set the `max_bytes_per_sec` and `send_request_concurrency` parameters in [Set Logtail startup parameters](#) to adjust the log data sending rate and the maximum concurrent requests. In this case, Logtail ensures that the sending rate and the concurrent requests do not exceed the limits.

If the log data fails to be sent, Logtail retries or stops the operation based on the error message.

| Error | Description | Handling method |
|-----------------|--|--|
| 401 | Logtail is not authorized to collect data. | Logtail drops the log packets. |
| 404 | The specified project or Logstore does not exist in the Logtail configuration. | Logtail drops the log packets. |
| 403 | The shard quota is exhausted. | Logtail waits for three seconds and retries. |
| 500 | A Log Service exception has occurred. | Logtail waits for three seconds and retries. |
| Network timeout | A network connection error has occurred. | Logtail waits for three seconds and retries. |

23.3.1.1.3. Logtail configuration files and record files

This topic describes the basic configuration files and record files of Logtail. When Logtail is active, it uses a series of configuration files and generates record files.

The basic configuration files are as follows:

- [Startup configuration file \(ilogtail_config.json\)](#)
- [Account ID configuration file](#)
- [User-defined identifier file \(user_defined_id\)](#)

- [Logtail configuration file \(user_log_config.json\)](#)

The basic record files are as follows:

- [AppInfo record file \(app_info.json\)](#)
- [Logtail operational log file \(ilogtail.LOG\)](#)
- [Logtail plug-in log file \(logtail_plugin.LOG\)](#)
- [Container path mapping file \(docker_path_config.json\)](#)

Startup configuration file (ilogtail_config.json)

This file is used to query or set Logtail runtime parameters. The file is in the JSON format.

After you install Logtail, you can use the startup configuration file to perform the following operations:

- Change the values of the Logtail runtime parameters.

You can change the CPU usage threshold, usage threshold of terminate and stay resident (TSR) programs, and other settings.

- Check whether the installation commands are correct.

The settings of `config_server_address` and `data_server_list` in this file depend on the parameters and installation commands selected when you installed Logtail. If the region specified in `config_server_address` is unreachable or is different from the region where Log Service resides, the selected parameters or commands are incorrect. In this case, Logtail cannot collect logs and must be reinstalled.

Warning

- The file must be a valid JSON file. Otherwise, Logtail cannot be started.
- If you modify the file, you must restart Logtail to validate your modifications.

The following table describes the default parameters in the startup configuration file. You can also add other parameters. For more information, see [Set Logtail startup parameters](#).

Default parameters

| Parameter | Description |
|------------------------------------|--|
| <code>config_server_address</code> | The address that Logtail uses to receive the configuration file from Log Service. This address depends on the parameters and installation commands that you selected when you installed Logtail. Ensure that the address is reachable and is in the same region as Log Service. |
| <code>data_server_list</code> | The data server address. This address depends on the parameters and installation commands that you selected when you installed Logtail. Ensure that the address is reachable and is in the same region as Log Service. |
| <code>cluster</code> | The name of the region where a server resides. |
| <code>endpoint</code> | The endpoint of Log Service. For more information, see View the information of a project . |
| <code>cpu_usage_limit</code> | The CPU usage threshold, which is calculated by core. |
| <code>mem_usage_limit</code> | The TSR usage threshold. |
| <code>max_bytes_per_sec</code> | The traffic limit on the raw data that is sent by Logtail. If the value of this parameter is greater than 20 Mbit/s, traffic limiting does not take effect. |

| Parameter | Description |
|--------------------------|--|
| process_thread_count | The number of threads that Logtail uses to write data to log files. |
| send_request_concurrency | The number of concurrent requests for sending data packets asynchronously. Logtail sends data packets asynchronously by default. If the write transactions per second (TPS) is high, you can set a greater value for this parameter. |

- File path
 - Linux: The file is stored in `/usr/local/ilogtail/ilogtail_config.json`.
 - Container Service: The file is stored in a Logtail container. The file path is specified in the environment variable `ALIYUN_LOGTAIL_CONFIG` of the Logtail container. You can run the command `docker inspect ${logtail_container_name} | grep ALIYUN_LOGTAIL_CONFIG` to query the file path. For example, the file path is `/etc/ilogtail/conf/cn-hangzhou/ilogtail_config.json`.
 - Windows:
 - 64-bit: The file is stored in `C:\Program Files (x86)\Alibaba\Logtail\ilogtail_config.json`.
 - 32-bit: The file is stored in `C:\Program Files\Alibaba\Logtail\ilogtail_config.json`.

- Sample file

```
$cat /usr/local/ilogtail/ilogtail_config.json
{
  "config_server_address": "http://logtail.cn-hangzhou-intranet.log.aliyuncs.com",
  "data_server_list":
  [
    {
      "cluster": "ap-southeast-2",
      "endpoint": "cn-hangzhou-intranet.log.aliyuncs.com"
    }
  ],
  "cpu_usage_limit": 0.4,
  "mem_usage_limit": 100,
  "max_bytes_per_sec": 2097152,
  "process_thread_count": 1,
  "send_request_concurrency": 4,
  "streamlog_open": false
}
```

Account ID configuration file

This file contains the ID of your Apsara Stack tenant account. The file indicates that the account can collect logs from the server where Logtail is installed. If you want to collect logs from ECS instances that do not belong to your account or from on-premises data centers, you must create an account ID configuration file.

Note

- The file is used only when you collect logs from ECS instances that do not belong to your account and or from on-premises data centers.
- The file can contain only the ID of your Apsara Stack tenant account. It cannot contain the IDs of RAM users under your Apsara Stack tenant account.
- The file name cannot contain a suffix.
- Each Logtail can have multiple account ID configuration files. Each Logtail container can have only one account ID configuration file.

- File path

- Linux: The file is stored in `/etc/ilogtail/users/`.
- Container Service: The file is stored in a Logtail container. The file path is specified in the environment variable `ALIYUN_LOGTAIL_USER_ID` of the Logtail container. You can run the command `docker inspect ${logtail_container_name} | grep ALIYUN_LOGTAIL_USER_ID` to query the file path.
- Windows: The file is stored in `C:\LogtailData\users\`.

- Sample file

```
$ls /etc/ilogtail/users/  
*****
```

User-defined identifier file (`user_defined_id`)

This file is used to configure machine groups with user-defined identifiers. For more information, see [Create a machine group based on a custom ID](#).

 **Note**

- This file is used only when you configure a machine group with user-defined identifiers.
- If you configure multiple user-defined identifiers for a machine group, separate them with line breaks.

- File path

- Linux: The file is stored in `/etc/ilogtail/user_defined_id`.
- Container Service: The file is stored in a Logtail container. The file path is specified in the environment variable `ALIYUN_LOGTAIL_USER_DEFINED_ID` of the Logtail container. You can run the command `docker inspect ${logtail_container_name} | grep ALIYUN_LOGTAIL_USER_DEFINED_ID` to query the file path.
- Windows: The file is stored in `C:\LogtailData\user_defined_id`.

- Sample file

```
$cat /etc/ilogtail/user_defined_id  
aliyun-ecs-rs1e16355
```

Logtail configuration file (`user_log_config.json`)

This file contains the Logtail configuration that Logtail receives from Log Service. The file is in the JSON format and is updated along with configuration updates. You can use this file to check whether the Logtail configuration is delivered to the server where Logtail is installed. If the Logtail configuration file exists and all contents in the file are up to date, the Logtail configuration is delivered.

 **Notice**

- We recommend that you do not modify the Logtail configuration file unless you need to specify sensitive information, such as the AccessKey pair and database password.
- You must upload this file when you submit a ticket.

- File path

- Linux: The file is stored in `/usr/local/ilogtail/user_log_config.json`.
- Container Service: The file is stored in `/usr/local/ilogtail/user_log_config.json`.
- Windows
 - 64-bit: The file is stored in `C:\Program Files (x86)\Alibaba\Logtail\user_log_config.json`.
 - 32-bit: The file is stored in `C:\Program Files\Alibaba\Logtail\user_log_config.json`.

- Sample file

```
$cat /usr/local/ilogtail/user_log_config.json
{
  "metrics": {
    "##1.0##k8s-log-c12ba2028*****939f0b$app-java": {
      "aliuid": "16542189*****50",
      "category": "app-java",
      "create_time": 1534739165,
      "defaultEndpoint": "cn-hangzhou-intranet.log.aliyuncs.com",
      "delay_alarm_bytes": 0,
      "enable": true,
      "enable_tag": true,
      "filter_keys": [],
      "filter_regs": [],
      "group_topic": "",
      "local_storage": true,
      "log_type": "plugin",
      "log_tz": "",
      "max_send_rate": -1,
      "merge_type": "topic",
      "plugin": {
        "inputs": [
          {
            "detail": {
              "IncludeEnv": {
                "aliyun_logs_app-java": "stdout"
              },
              "IncludeLabel": {
                "io.kubernetes.container.name": "java-log-demo-2",
                "io.kubernetes.pod.namespace": "default"
              },
              "Stderr": true,
              "Stdout": true
            },
            "type": "service_docker_stdout"
          }
        ]
      },
      "priority": 0,
      "project_name": "k8s-log-c12ba2028c*****ac1286939f0b",
      "raw_log": false,
      "region": "cn-hangzhou",
      "send_rate_expire": 0,
      "sensitive_keys": [],
      "tz_adjust": false,
      "version": 1
    }
  }
}
```

AppInfo record file (app_info.json)

This file contains the startup time of Logtail. It also contains the IP address and hostname that Logtail obtains. You must check the IP address obtained by Logtail when you configure an [IP address-based machine group](#).

In most cases, Logtail obtains server IP addresses based on the following rules:

- If the IP address of a server is associated with its hostname in the `/etc/hosts` server file, Logtail obtains the IP address.
- If the IP address of a server is not associated with its hostname, Logtail obtains the IP address of the first network interface card (NIC) on the server.

 **Note**

- The AppInfo record file contains only the basic Logtail information, which cannot be manually modified.
- If you modify the hostname or other network settings of the server, you must restart Logtail to obtain a new IP address.

Parameters

| Parameter | Description |
|-----------------|--|
| UUID | The serial number of the server. |
| hostname | The hostname. |
| instance_id | The unique identifier of Logtail. This identifier is randomly generated. |
| ip | <p>The IP address that is obtained by Logtail. If this parameter is not specified, Logtail has not obtained the IP address of a server. In this case, Logtail cannot function properly. You must set an IP address for your server and restart Logtail.</p> <div data-bbox="614 862 1388 1048" style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p> Note If the machine group is an IP address-based machine group, ensure that the IP address specified for the machine group is the same as the value of this parameter. If the two IP address are different, modify the IP address that you specified for the machine group in the Log Service console. Check the IP addresses again after 1 minute.</p> </div> |
| logtail_version | The version of Logtail. |
| os | The version of the operating system. |
| update_time | The last startup time of Logtail. |

- File path
 - Linux: The file is stored in `/usr/local/ilogtail/app_info.json`.
 - Container Service: The file is stored in `/usr/local/ilogtail/app_info.json`.
 - Windows
 - 64-bit: The file is stored in `C:\Program Files (x86)\Alibaba\Logtail\app_info.json`.
 - 32-bit: The file is stored in `C:\Program Files\Alibaba\Logtail\app_info.json`.

• Sample file

```
$cat /usr/local/ilogtail/app_info.json
{
  "UUID": "",
  "hostname": "logtail-ds-slpn8",
  "instance_id": "E5F93BC6-B024-11E8-8831-0A58AC14039E_1**.***.***.***_1536053315",
  "ip": "1**.***.***.***",
  "logtail_version": "0.16.13",
  "os": "Linux; 3.10.0-693.2.2.el7.x86_64; #1 SMP Tue Sep 12 22:26:13 UTC 2017; x86_64",
  "update_time": "2018-09-04 09:28:36"
}
```

Logtail operational log file (ilogtail.LOG)

This file contains operational information about Logtail. Log severity levels are ranked as follows in ascending order: `INFO` , `WARN` , `ERROR` . Logs of the `INFO` level can be ignored.

- File path
 - For Linux: The file is stored in `/usr/local/ilogtail/ilogtail.LOG`.
 - Container Service: The file is stored in `/usr/local/ilogtail/ilogtail.LOG`.
 - Windows
 - 64-bit: The file is stored in `C:\Program Files (x86)\Alibaba\Logtail\logtail_*.log`.
 - 32-bit: The file is stored in `C:\Program Files\Alibaba\Logtail\logtail_*.log`.
- Sample file

```
$tail /usr/local/ilogtail/ilogtail.LOG
[2018-09-13 01:13:59.024679] [INFO] [3155] [build/release64/sls/ilogtail/elogtail.cpp:123] change working dir:/usr/local/ilogtail/
[2018-09-13 01:13:59.025443] [INFO] [3155] [build/release64/sls/ilogtail/AppConfig.cpp:175] load logtail config file , path:/etc/ilogtail/conf/ap-southeast-2/ilogtail_config.json
[2018-09-13 01:13:59.025460] [INFO] [3155] [build/release64/sls/ilogtail/AppConfig.cpp:176] load logtail config file , detail:{
  "config_server_address": "http://logtail.ap-southeast-2-intranet.log.aliyuncs.com",
  "data_server_list": [
    {
      "cluster": "ap-southeast-2",
      "endpoint": "ap-southeast-2-intranet.log.aliyuncs.com"
    }
  ]
}
```

Logtail plug-in log file (logtail_plugin.LOG)

This file contains operational information about plug-ins, such as `stdout` , `binlog` , and `HTTP` plug-ins. Log severity levels are ranked as follows in ascending order: `INFO` , `WARN` , `ERROR` . Logs of the `INFO` level can be ignored.

- File path
 - Linux: The file is stored in `/usr/local/ilogtail/logtail_plugin.LOG`
 - Container Service: The file is stored in `/usr/local/ilogtail/logtail_plugin.LOG`.
 - Windows: The file is not supported.
- Sample file

```
$tail /usr/local/ilogtail/logtail_plugin.LOG
2018-09-13 02:55:30 [INF] [docker_center.go:525] [func1] docker fetch all:start
2018-09-13 02:55:30 [INF] [docker_center.go:529] [func1] docker fetch all:stop
2018-09-13 03:00:30 [INF] [docker_center.go:525] [func1] docker fetch all:start
2018-09-13 03:00:30 [INF] [docker_center.go:529] [func1] docker fetch all:stop
2018-09-13 03:03:26 [INF] [log_file_reader.go:221] [ReadOpen] [##1.0##sls-zc-test-hz-pub$docker-stdout-config,k8s-stdout] open file for read, file:/logtail_host/var/lib/docker/containers/7f46afec6a14de39b59ee9cdfbfa8a70c2fa26f1148b2e2f31bd3410f5b2d624/7f46afec6a14de39b59ee9cdfbfa8a70c2fa26f1148b2e2f31bd3410f5b2d624-json.log offset:40379573 status:794354-64769-40379963
2018-09-13 03:03:26 [INF] [log_file_reader.go:221] [ReadOpen] [##1.0##k8s-log-c12ba2028cfb444238cd9ac1286939f0b$docker-stdout-config,k8s-stdout] open file for read, file:/logtail_host/var/lib/docker/containers/7f46afec6a14de39b59ee9cdfbfa8a70c2fa26f1148b2e2f31bd3410f5b2d624/7f46afec6a14de39b59ee9cdfbfa8a70c2fa26f1148b2e2f31bd3410f5b2d624-json.log offset:40379573 status:794354-64769-40379963
2018-09-13 03:04:26 [INF] [log_file_reader.go:308] [CloseFile] [##1.0##sls-zc-test-hz-pub$docker-stdout-config,k8s-stdout] close file, reason:no read timeout file:/logtail_host/var/lib/docker/containers/7f46afec6a14de39b59ee9cdfbfa8a70c2fa26f1148b2e2f31bd3410f5b2d624/7f46afec6a14de39b59ee9cdfbfa8a70c2fa26f1148b2e2f31bd3410f5b2d624-json.log offset:40379963 status:794354-64769-40379963
2018-09-13 03:04:27 [INF] [log_file_reader.go:308] [CloseFile] [##1.0##k8s-log-c12ba2028cfb444238cd9ac1286939f0b$docker-stdout-config,k8s-stdout] close file, reason:no read timeout file:/logtail_host/var/lib/docker/containers/7f46afec6a14de39b59ee9cdfbfa8a70c2fa26f1148b2e2f31bd3410f5b2d624/7f46afec6a14de39b59ee9cdfbfa8a70c2fa26f1148b2e2f31bd3410f5b2d624-json.log offset:40379963 status:794354-64769-40379963
2018-09-13 03:05:30 [INF] [docker_center.go:525] [func1] docker fetch all:start
2018-09-13 03:05:30 [INF] [docker_center.go:529] [func1] docker fetch all:stop
```

Container path mapping file (docker_path_config.json)

This file is automatically created only when container files are collected. It records path mappings between container files and actual files. The file is in the JSON format.

 **Note** This file is only an information record file. Modifications to this file do not take effect. If you delete this file, another one is automatically created without service interruptions.

- File path

The file is stored in `/usr/local/ilogtail/docker_path_config.json`.

- Sample file

```
$cat /usr/local/ilogtail/docker_path_config.json
{
  "detail": [
    {
      "config_name": "##1.0##k8s-log-c12ba2028cfb444238cd9ac1286939f0b$nginx",
      "container_id": "df19c06e854a0725ea7fca7e0378b0450f7bd3122f94fe3e754d8483fd330d10",
      "params": "{\n  \"ID\": \"df19c06e854a0725ea7fca7e0378b0450f7bd3122f94fe3e754d8483fd330d10\", \n  \"Path\": \n  \"/logtail_host/var/lib/docker/overlay2/947db346695a1f65e63e582ecfd10ae1f57019a1b99260b6c83d00fcd1892874/diff/var/log\", \n  \"Tags\": [\n    \"nginx-type\", \n    \"access-log\", \n    \"_image_name_\", \n    \"registry.cn-hangzhou.aliyuncs.com/log-service/docker-log-test:latest\", \n    \"_container_name_\", \n    \"nginx-log-demo\", \n    \"_pod_name_\", \n    \"nginx-log-demo-h2lzc\", \n    \"_namespace_\", \n    \"default\", \n    \"_pod_uid_\", \n    \"87e56ac3-b65b-11e8-b172-00163f008685\", \n    \"_container_ip_\", \n    \"172.20.4.224\", \n    \"purpose\", \n    \"test\"\n  ]\n}"
    },
    {
      "version": "0.1.0"
    }
  ]
}
```

23.3.1.2. Installation

23.3.1.2.1. Install Logtail in Linux

This topic describes how to install Logtail on a Linux server.

Supported systems

Logtail supports the following x86-64 (64-bit) Linux operating systems:

- Aliyun Linux
- Ubuntu
- Debian
- CentOS
- openSUSE
- Red Hat

Procedure

Note If you have installed Logtail, the installer will uninstall the existing version of Logtail, delete the `/usr/local/ilogtail` directory, and then reinstall Logtail. By default, Logtail runs after the installation and at startup.

1. Run the following command to download the Logtail installer:

```
wget http://${service:sls-backend-server:sls_data.endpoint}/logtail.sh -O logtail.sh; chmod 755 logtail.sh
```

Note You must replace `${service:sls-backend-server:sls_data.endpoint}` in the command with the actual endpoint. For more information about endpoints, see [View the information of a project](#).

2. Run the installation command.

Start Linux PowerShell and run the following command as an administrator to install Logtail:

```
./logtail.sh install
```

3. [Configure an account ID for a server](#).

View the version of Logtail

To view the version of Logtail, open the file in the `/usr/local/ilogtail/app_info.json` directory. The `logtail_version` field shows the version of Logtail.

```
$cat /usr/local/ilogtail/app_info.json
{
  "UUID": "0DF18E97-0F2D-486F-B77F-*****",
  "hostname": "david*****",
  "instance_id": "F4FAFADA-F1D7-11E7-846C-00163E30349E-*****_1515129548",
  "ip": "*****",
  "logtail_version": "0.16.0",
  "os": "Linux; 2.6.32-220.23.2.ali1113.el5.x86_64; #1 SMP Thu Jul 4 20:09:15 CST 2013; x86_64",
  "update_time": "2018-01-05 13:19:08"
}
```

Upgrade Logtail

You can use the Logtail installer (`logtail.sh`) to upgrade Logtail. The installer selects an upgrade method based on the configurations of the existing Logtail.

Note During the upgrade, Logtail is temporarily stopped. Only related files are overwritten. The configuration file, checkpoint file, and logs are retained.

Run the following commands to upgrade Logtail:

```
# Download the Logtail installer.
wget http://${service:sls-backend-server:sls_data.endpoint}/logtail.sh -O logtail.sh; chmod 755 logtail.sh
# Upgrade Logtail.
sudo ./logtail.sh upgrade
```

Response:

```
# The upgrade is successful.
Stop logtail successfully.
ilogtail is running
Upgrade logtail success
{
  "UUID": "****",
  "hostname": "****",
  "instance_id": "****",
  "ip": "****",
  "logtail_version": "0.16.11",
  "os": "Linux; 3.10.0-693.2.2.el7.x86_64; #1 SMP Tue Sep 12 22:26:13 UTC 2017; x86_64",
  "update_time": "2018-08-29 15:01:36"
}
# The upgrade fails because the current version is the latest version.
[Error]: Already up to date.
```

Start and stop Logtail

- Start Logtail

Run the following command as an administrator:

```
/etc/init.d/ilogtailed start
```

- Stop Logtail

Run the following command as an administrator:

```
/etc/init.d/ilogtailed stop
```

Uninstall Logtail

Run Linux PowerShell as an administrator to uninstall Logtail:

```
wget http://${service:sls-backend-server:sls_data.endpoint}/logtail.sh -O logtail.sh
chmod 755 logtail.sh
./logtail.sh uninstall
```

23.3.1.2.2. Install Logtail in Windows

This topic describes how to install Logtail on a Windows server.

Supported systems

Logtail supports the following Windows operating systems:

- Windows 7 (Client) 32-bit
- Windows 7 (Client) 64-bit
- Windows Server 2008 32-bit

- Windows Server 2008 64-bit
- Windows Server 2012 64-bit
- Windows Server 2016 64-bit

Procedure

1. Download the installation package.

Run the following command to download the installation package:

```
wget http://${service:sls-backend-server:sls_data.endpoint}/windows/logtail_installer.zip
```

Note You must replace `${service:sls-backend-server:sls_data.endpoint}` in the command with the actual endpoint. For more information about endpoints, see [View the information of a project](#).

2. Decompress the `logtail_installer.zip` package to the current directory.
3. Run the installation command.

Run Windows PowerShell or Command Prompt as an administrator. Enter the `logtail_installer` directory, and then run the installation command based on the network type.

```
.\logtail_installer.exe install me-east-1
```

Note You must replace `${region}` in the command with the actual endpoint. For more information about endpoints, see [View the information of a project](#).

4. [Configure an account ID for a server](#).

Installation directory

After you run the installation command, Logtail is installed in the specified directory. The directory cannot be changed. In the directory, you can [View the version of Logtail](#) in the `app_info.json` file or [Uninstall Logtail](#).

The installation directory is as follows:

- 32-bit Windows: `C:\Program Files\Alibaba\Logtail`
- 64-bit Windows: `C:\Program Files (x86)\Alibaba\Logtail`

Note You can run 32-bit or 64-bit applications in a 64-bit Windows operating system. However, the operating system stores 32-bit applications in separate x86 folders to ensure compatibility.

Logtail for Windows is a 32-bit application. Therefore, it is installed in the `Program Files (x86)` folder in 64-bit Windows. If Logtail for 64-bit Windows becomes available in the future, it will be installed in the `Program Files` folder.

View the version of Logtail

To view the version of Logtail, go to the [default installation directory](#), and then use the notepad or another text editor to open the `app_info.json` file. The `logtail_version` field shows the version of Logtail.

In the following example, the version of Logtail is 1.0.0.0:

```
{
  "logtail_version": "1.0.0.0"
}
```

Upgrade Logtail

- Automatic upgrade

Logtail later than 1.0.0.0 is automatically upgraded in Windows.

- Manual upgrade

Logtail earlier than 1.0.0.0 must be manually upgraded. The manual upgrade procedure is the same as the installation [procedure](#).

 **Note** During a manual upgrade, the files in the original installation directory are deleted. We recommend that you back up the files before you perform a manual upgrade.

Start and stop Logtail

Open the **Control Panel**, choose System and Security > **Administrative Tools**, and then double-click **Services**.

Find the service based on your Logtail version.

- Logtail 0.x.x.x: LogtailWorker.
- Logtail 1.0.0.0 and later: LogtailDaemon.

Perform the following operations as required:

- Start Logtail

Right-click the service and select **Start** from the shortcut menu.

- Stop Logtail

Right-click the service and select **Stop** from the shortcut menu.

- Restart Logtail

Right-click the service and select **Restart** from the shortcut menu.

Uninstall Logtail

Run Windows PowerShell or Command Prompt as an administrator. Enter the `logtail_installer` directory, and then run the following command:

```
.\logtail_installer.exe uninstall
```

After Logtail is uninstalled, the [installation directory](#) is deleted. However, some residual configuration data is still maintained in the `C:\LogtailData` directory. You can manually delete the data. The residual configuration data includes the following information:

- *checkpoint*: checkpoints of all plug-ins, for example, the Windows event log plug-in.
- *logtail_check_point*: checkpoints of Logtail.
- *users*: IDs of Apsara Stack tenant accounts.

23.3.1.2.3. Set Logtail startup parameters

This topic describes how to set Logtail startup parameters.

Context

You may need to set Logtail startup parameters in the following scenarios:

- You need to collect a large number of log files that consume much memory. You want to maintain the metadata (such as the file signature, collection location, and file name) of each file in the memory.
- The CPU usage is high due to heavy log data traffic.
- The traffic sent to Log Service is heavy due to a large amount of log data.
- You want to collect syslogs or TCP data streams.

Startup configurations

- File path

```
/usr/local/ilogtail/ilogtail_config.json
```

- File format

JSON

- Sample file (only partial configurations are provided)

```
{
  ...
  "cpu_usage_limit": 0.4,
  "mem_usage_limit": 100,
  "max_bytes_per_sec": 2097152,
  "process_thread_count": 1,
  "send_request_concurrency": 4,
  "streamlog_open": false,
  "streamlog_pool_size_in_mb": 50,
  "streamlog_rcv_size_each_call": 1024,
  "streamlog_formats": [],
  "streamlog_tcp_port": 11111,
  "buffer_file_num": 25,
  "buffer_file_size": 20971520,
  "buffer_file_path": "",
  ...
}
```

Startup parameters

| Parameter | Description | Example |
|--------------------------|--|--|
| cpu_usage_limit | The CPU usage threshold for a single core. Data type: double. | For example, the value 0.4 indicates that the CPU usage of Logtail is limited to 40% processing capacity of a single core. In most cases, the processing capacity of a single core is about 24 MB/s in the simple mode and 12 MB/s in the full regex mode. |
| mem_usage_limit | The usage threshold of the resident memory. Data type: integer. Unit: MB. | For example, the value 100 indicates that the memory usage of Logtail is limited to 100 MB. If the threshold is exceeded, Logtail restarts. If you want to collect more than 1,000 log files, you can increase the threshold value. |
| max_bytes_per_sec | The traffic limit on the raw data that is sent by Logtail. Data type: integer. Unit: bytes/s. | For example, the value 2097152 indicates that the data transfer rate of Logtail is limited to 2 MB/s. |
| process_thread_count | The number of threads that Logtail uses to process data. | Default value: 1. Each thread provides a write speed of 24 MB/s in the simple mode and 12 MB/s in the full regex mode. We recommend that you do not modify the default value. |
| send_request_concurrency | Logtail sends data packets asynchronously by default. If the write transactions per second (TPS) is high, you can set this parameter to a greater value. | Twenty asynchronous concurrencies are provided by default. Each concurrency can provide 0.5 MB/s to 1 MB/s network throughput. The number of concurrencies varies with the network delay. |

| Parameter | Description | Example |
|------------------------------|---|---|
| streamlog_open | Specifies whether to receive syslogs. Data type: Boolean. | The value false indicates that syslogs are not received. The value true indicates that syslogs are received. |
| streamlog_pool_size_in_mb | The size of memory pool that the syslog server uses to cache syslogs. Unit: MB. | Logtail requests memory when it starts. Set the memory pool size based on the server memory size and your business requirements. |
| streamlog_rcv_size_each_call | The size of the buffer that Logtail uses when the linux socket rcv API is called. Unit: bytes. Valid values: 1024 to 8192. | You can set a greater value if the syslog traffic is high. |
| streamlog_formats | The method that is used to parse received syslogs. | N/A |
| streamlog_tcp_addr | The associated address that Logtail uses to receive syslogs. Default value: 0.0.0.0. | N/A |
| streamlog_tcp_port | The TCP port that Logtail uses to receive syslogs. | Default value: 11111. |
| buffer_file_num | The maximum number of cached files. If a network exception occurs or the writing quota is exceeded, Logtail writes parsed logs to local files in the installation directory. After the network recovers or a new writing quota is available, Logtail retries to send the logs to Log Service. | Default value: 25. |
| buffer_file_size | The maximum number of bytes that can be contained in each cache file. The maximum disk space available for cache files is the value of buffer_file_num multiplied by the value of buffer_file_size. | Default value: 20971520 bytes (20 MB). |
| buffer_file_path | The directory in which cached files are stored. If you modify this parameter, you must move the files (for example, <i>logtail_buffer_file_*</i>) in the old cache directory to the new directory. Then, Logtail can read, send, and delete the cache files. | The default value is null, which indicates that the cached files are stored in the Logtail installation directory <i>/usr/local/ilogtail</i> . |
| bind_interface | The name of the NIC associated with the local machine, for example, eth1. This parameter is valid only for Logtail that runs in Linux. | By default, the available NICs are automatically associated with the local machine. If you specify this parameter, Logtail will use the specified NIC to upload logs. |

| Parameter | Description | Example |
|----------------------|---|---|
| check_point_filename | The full path in which the checkpoint file is stored. This parameter is used to customize the path to store the checkpoint file of Logtail. | Default value: <code>/tmp/logtail_checkpoint</code> . We recommend that Docker users modify this path and mount the directory where the checkpoint file resides to the host. Otherwise, duplicate collection occurs due to checkpoint data loss when the container is released. For example, you can set <code>check_point_filename</code> to <code>/data/logtail/checkpoint.dat</code> in Docker and add <code>-v /data/docker1/logtail:/data/logtail</code> to the Docker startup command. Then, the <code>/data/docker1/logtail</code> directory of the host is mounted to the <code>/data/logtail</code> directory of Docker. |

Note

- The preceding table lists only the common startup parameters. If the `ilogtail_config.json` file contains parameters that are not listed in the table, the default settings are used for these parameters.
- We recommend that you do not add unnecessary parameters to the `ilogtail_config.json` file.

Modify configurations

1. Configure the `ilogtail_config.json` file as needed.

Ensure that the modified configurations are in the valid JSON format.

2. Restart Logtail to apply the modified configurations.

```
/etc/init.d/ilogtailed stop
/etc/init.d/ilogtailed start
/etc/init.d/ilogtailed status
```

23.3.1.3. Logtail machine group

23.3.1.3.1. Overview

Log Service uses machine groups to manage the servers from which you want to collect logs by using Logtail.

A machine group is a virtual group that contains multiple servers. If you want to use a Logtail configuration file to collect logs from multiple servers, you can add the servers to a machine group. Then, you can apply the Logtail configuration file to the machine group.

To define a machine group, you can use one of the following methods:

- IP address: Add the IP addresses of all servers to a machine group. Each server can be identified by using its unique IP address.
- Custom ID: Use a custom ID to identify the machine group and use the same ID for servers in the machine group.

 Note Windows and Linux servers cannot be added to the same machine group.

IP address-based machine groups

You can add multiple servers to a machine group by adding their IP addresses to the machine group. Then, you can create a Logtail configuration file for all the servers at the same time.

- If you use ECS instances and have not associated them with hostnames or changed their network types, you can add their private IP addresses to the machine group.

- In other cases, you must add the server IP addresses obtained by Logtail to a machine group. The IP address of each server is recorded in the IP address field of the `app_info.json` file on the server.

 **Note** The `app_info.json` file records the internal information of Logtail. This file includes the server IP addresses obtained by Logtail. If you modify the IP address field of the file, the IP addresses obtained by Logtail remain unchanged.

Logtail obtains a server IP address by using the following methods:

- If the IP address of a server is associated with the hostname in the `/etc/hosts` file of the server, Logtail obtains this IP address.
- If the IP address of a server is not associated with the hostname, Logtail obtains the IP address of the first network interface controller (NIC) on the server.

For more information, see [Create an IP address-based server group](#).

Custom ID-based machine groups

You can use custom IDs to dynamically define machine groups.

An application system consists of multiple modules. You can scale out each module by adding multiple servers to the module. If you want to collect logs by module, you can create a machine group for each module. Therefore, you must specify a custom ID for each server in each module. For example, a website consists of an HTTP request processing module, a caching module, a logic processing module, and a storage module. The custom IDs of these modules can be `http_module`, `cache_module`, `logic_module`, and `store_module`.

For more information, see [Create a machine group based on a custom ID](#).

23.3.1.3.2. Create a machine group based on a server IP address

This topic describes how to create a machine group based on a server IP address. You can create a machine group based on a server IP address that is obtained by using a Logtail configuration file. You can then use the same Logtail configuration file to collect logs from the machine group.

Prerequisites

- A project is created. A Logstore is created in the project.
- One or more servers are available. The IP addresses of the servers are obtained.
- Logtail is installed on the servers. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).
- The ID of your Apsara Stack tenant account is configured on the server. For more information, see [Configure an account ID for a server](#).

Procedure

1. [Log on to the Log Service console](#).
2. In the **Projects** section, click the name of a project.
3. In the left-side navigation pane, click the **Machine Groups** icon.
4. In the pane that appears, click the  icon next to **Machine Group** and select **Create Machine Group** from the shortcut menu.

You can also create a machine group in the Logtail configuration wizard.

5. Create a machine group.

- i. Enter a machine group name in the **name** field.

The machine group name must be 3 to 128 characters in length and can contain lowercase letters, digits, hyphens (-), and underscores (_). It must start and end with a lowercase letter or digit.

 **Notice** After the machine group is created, its name cannot be modified.

- ii. Select **IP Addresses** from the **Identifier** drop-down list.
- iii. Enter a topic name in the **Topic** field.
For more information about topics, see [Generate a topic](#).
- iv. Enter the IP addresses of the servers in the **IP Addresses** text box.

 **Notice**

- Separate each IP address with a line break.
- Do not add Windows servers and Linux servers to the same machine group.

6. Click **OK**.

Result

You can view the created machine group in the **Machine Groups** pane.



23.3.1.3.3. Create a machine group based on a custom ID

This topic describes how to create a machine group based on a custom ID.

Context

You can use a custom ID to identify a machine group in the following scenarios:

- Servers reside in multiple custom network environments such as virtual private clouds (VPCs). IP addresses of different servers may be the same. In this scenario, Log Service cannot distinguish between servers based on IP addresses.
- You want to implement automatic server discovery. To do this, you only need to set a custom ID of a new server to the custom ID of an existing machine group. Log Service automatically identifies the server and adds it to the machine group.

Procedure

1. Set a custom ID on a server.

- **Linux Logtail**

Set a custom ID in the `/etc/ilogtail/user_defined_id` file.

For example, if you need to set a custom ID of a server to `userdefined`, run the following command to open the file:

```
# vim /etc/ilogtail/user_defined_id
```

In the file, enter `userdefined`.

- **Windows Logtail**

Set a custom ID in the `C:\LogtailData\user_defined_id` file.

For example, if you need to set a custom ID of a server, run the following command:

```
C:\LogtailData>more user_defined_id
userdefined_windows
```

Notice

- o A machine group cannot include both Linux and Windows servers. Therefore, do not set a custom ID of a Linux server and a custom ID of a Windows server to the same value.
- o You can set multiple custom IDs for a single server. Separate each custom ID with a line break.
- o If the `/etc/ilogtail/` or `C:\LogtailData` directory or the `/etc/ilogtail/user_defined_id` or `C:\LogtailData\user_defined_id` file does not exist, create the directory or the file.

2. Create a machine group.

- Log on to the Log Service console.
- In the **Projects** section, click a project.
- In the left-side navigation pane, click the **Machine Groups** icon.
- Click the  icon next to **Machine Groups**, and select **Create Machine Group** from the shortcut menu.
- Set the parameters of the machine group.

- **name:** Enter a machine group name.

The machine group name must be 3 to 128 characters in length and can contain lowercase letters, digits, hyphens (-), and underscores (_). It must start and end with a lowercase letter or digit.

Note After the machine group is created, its name cannot be modified.

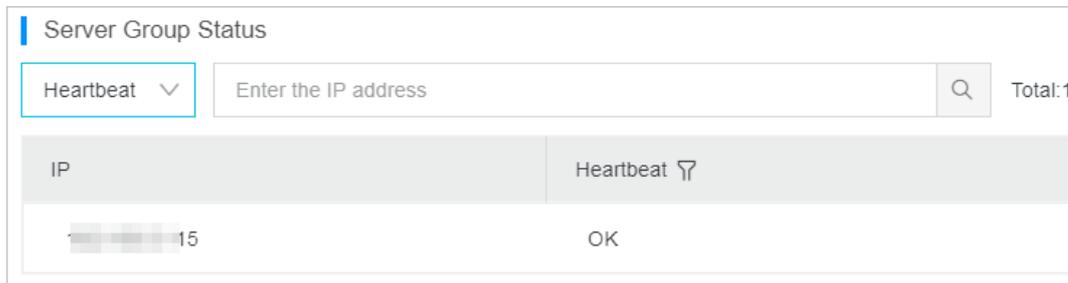
- **Identifier:** Select **Custom ID**.
- **Topic:** Enter a topic name for the machine group. For more information, see [Generate a topic](#).
- **Custom Identifier:** Enter the custom ID that you set in Step 1.

- Click **OK**.

Note If you need to add a server to the machine group, set a custom ID of the server to the custom ID of the machine group. The server is then listed in the **Machine Group Status** section.

3. View the status of the machine group.

In the **Machine Groups** pane, click the name of the machine group. On the **Machine Group Settings** page, view the status of the machine group. You can view the IP address list of the servers in the machine group and their heartbeat status.



Note

- The **Machine Group Status** section lists the IP addresses of the servers whose custom ID is the same as the custom ID that you set for the machine group.

For example, the custom ID of a machine group is userdefined and the IP addresses in the **Machine Group Status** section are 10.10.10.10, 10.10.10.11, and 10.10.10.12. This means that the custom ID userdefined is set for the three servers. If you need to add the server whose IP address is 10.10.10.13 to the machine group, set the custom ID userdefined for the server. Then, the IP address of the server is displayed in the **Machine Group Status** section. Log Service collects logs from the server based on the Logtail configuration file of the machine group.
- The heartbeat status indicates whether the connection between a server and Log Service is normal. For information about how to troubleshoot heartbeat errors, see [What can I do if no heartbeat packet is received from a Logtail client?](#)

Delete the custom IDs of a server

If you need to change the ID of a server from a custom ID to the server IP address, delete the `user_defined_id` file. The change takes effect within 1 minute.

- In Linux, run the following command to delete the file:

```
rm -f /etc/ilogtail/user_defined_id
```

- In Windows, run the following command to delete the file:

```
del C:\LogtailData\user_defined_id
```

Time required for a change to take effect

After you create, edit, or delete a `user_defined_id` file, the change takes effect within 1 minute.

If you need the change to take immediate effect, restart Logtail.

- In Linux, run the following commands to restart Logtail:

```
/etc/init.d/ilogtaild stop
/etc/init.d/ilogtaild start
```

- In Windows, perform the following steps to restart Logtail:

Open the Control Panel. In the window that appears, choose **Control Panel > Administrative Tools > Services**. Right-click **LogtailWorker**, and select **Restart** from the shortcut menu.

Example

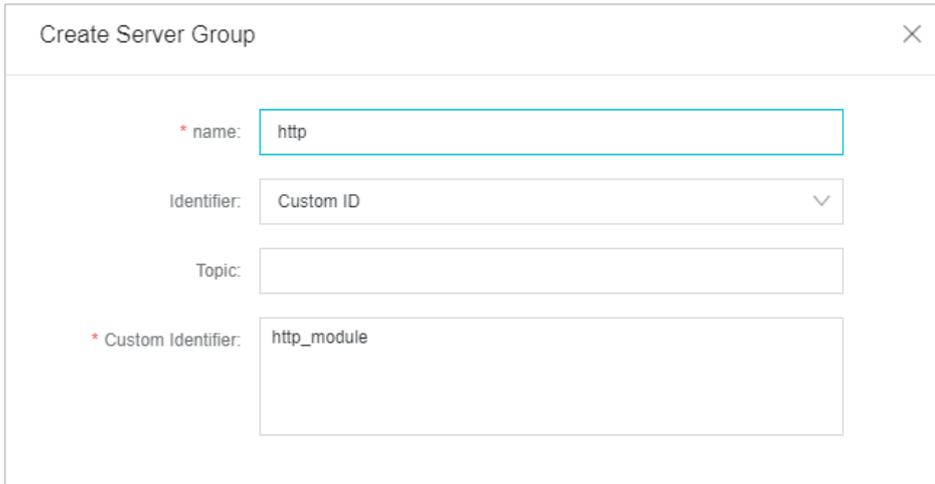
An application consists of multiple modules. Each module runs on multiple servers. For example, a website consists of an HTTP request processor, a cache, a logic processor, and a storage. You may scale out each module by adding multiple servers. You need to collect logs from both the existing and new servers.

1. Set a custom ID for each server.

Install Logtail on the servers and set a custom ID for each server. In this example, you can use four custom IDs: `http_module`, `cache_module`, `logic_module`, and `store_module`. Each custom ID corresponds to a module.

2. Create a machine group for each module.

When you create a machine group for a module, enter the custom ID of the module in the **Custom Identifier** field.



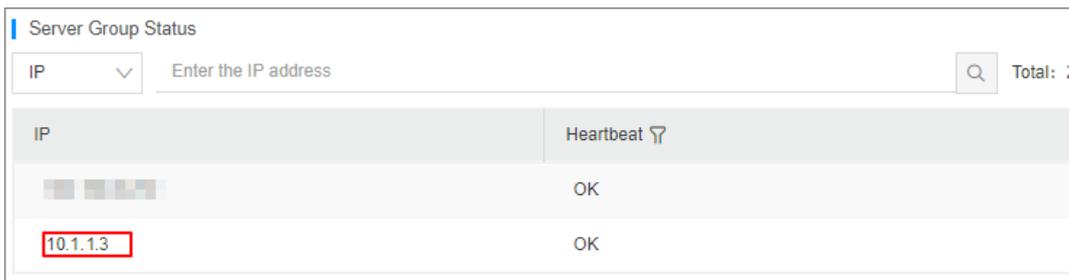
The screenshot shows a 'Create Server Group' dialog box with the following fields:

- * name:**
- Identifier:**
- Topic:**
- * Custom Identifier:**

3. View the status of the machine group.

On the **Machine Group Settings** page of the machine group, you can view the status of the machine group in the **Machine Group Status** section. You can view the list of servers in the machine group and their heartbeat status.

4. If you need to add a server whose IP address is 10.1.1.3 to the machine group whose custom ID is `http_module`, set the custom ID `http_module` for the server. Then, you can view the server in the **Machine Group Status** section.



| IP | Heartbeat |
|-----------|-----------|
| [blurred] | OK |
| 10.1.1.3 | OK |

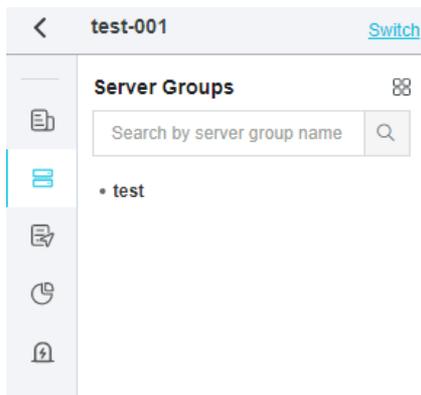
23.3.1.3.4. View server groups

This topic describes how to view the server groups of a project on the **Server Groups** page in the Log Service console.

Procedure

1. [Log on to the Log Service console](#).
2. Find the target project in the project list and click the project name.
3. In the left-side navigation pane of the page that appears, click the **Server Groups** icon to display the list of server groups.

You can view all server groups of the project.



23.3.1.3.5. Modify a server group

This topic describes how to modify a server group in the Log Service console. After you create a server group, you can modify the parameters of the server group.

Procedure

1. [Log on to the Log Service console.](#)
2. Find the target project in the project list and click the project name.
3. In the left-side navigation pane of the page that appears, click the **Server Groups** icon to display the list of server groups.
4. Click the name of the server group to be modified. On the **Server Group Settings** page, click **Modify**.

Note The name of the server group cannot be modified.

5. Modify the parameters of the server group, and then click **Save**.

23.3.1.3.6. View the status of a server group

This topic describes how to view the status of a server group in the Log Service console. You can view the heartbeat information of Logtail to check whether Logtail is installed on the servers in a server group.

Procedure

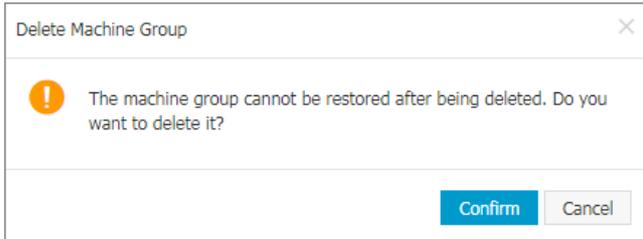
1. [Log on to the Log Service console.](#)
2. Find the target project in the project list and click the project name.
3. In the left-side navigation pane of the page that appears, click the **Server Groups** icon to display the list of server groups.
4. Click the name of the server group. On the **Server Group Settings** page, check the server group status.
 - If the heartbeat is OK, Logtail is installed on the servers in the server group and Logtail is connected to Log Service.
 - If the heartbeat status is FAIL, Logtail fails to connect to Log Service. If the FAIL state persists, perform troubleshooting based on the instructions provided in [What can I do if no heartbeat packet is received from a Logtail client?](#)

23.3.1.3.7. Delete a server group

This topic describes how to delete a server group in the Log Service console. You can delete a server group if you no longer need to collect logs from the server group.

Procedure

1. [Log on to the Log Service console.](#)
2. Find the target project in the project list and click the project name.
3. In the left-side navigation pane of the page that appears, click the **Server Groups** icon to display the list of server groups.
4. Find the server group that you want to delete, click the  icon next to the server group, and then select **Delete**.
5. In the dialog box that appears, click **OK**.



23.3.1.3.8. Manage server group configurations

This topic describes how to manage server group configurations in the Log Service console. Log Service uses server groups to manage the servers from which you collect logs by using Logtail. In the Log Service console, you can create, view, modify, and delete server groups. You can also view the status of server groups, manage server group configurations, and apply server group identifiers.

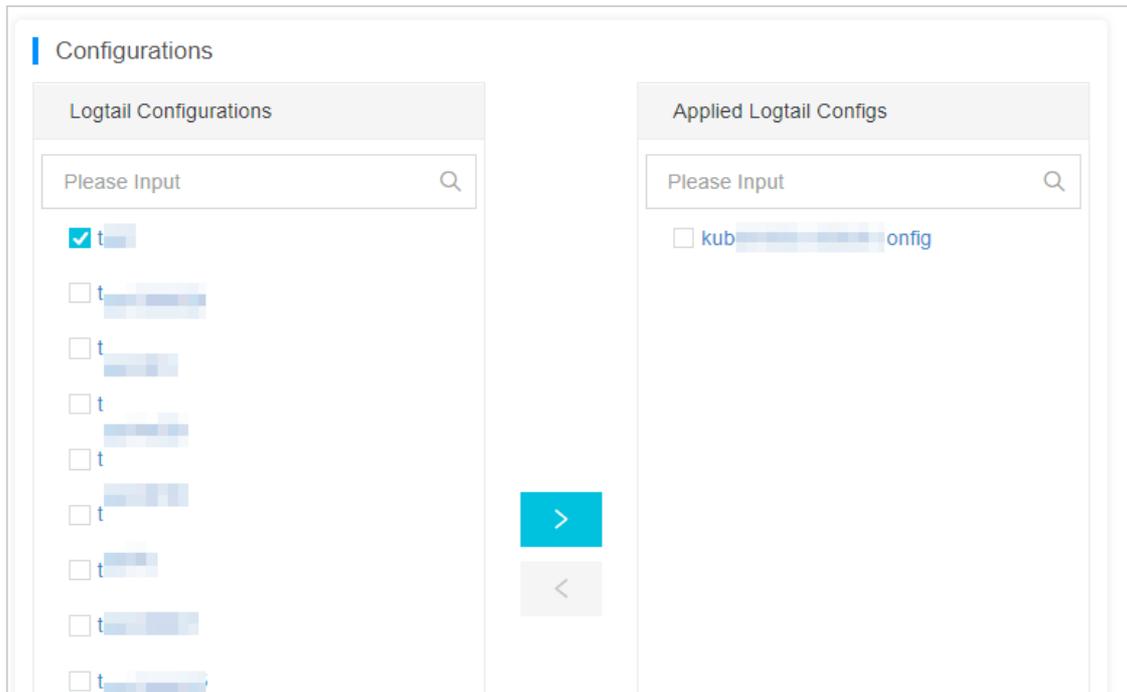
Context

Log Service allows you to manage the Logtail configurations that you create for Logtail installed on the servers in a server group. You can apply Logtail configurations to a server group. The Logtail configurations determine what logs are collected on each server, how the logs are parsed, and which Logstore the logs are written to.

Procedure

1. [Log on to the Log Service console.](#)
2. Find the target project in the project list and click the project name.
3. In the left-side navigation pane of the page that appears, click the **Server Groups** icon to display the list of server groups.
4. Click the name of the server group whose configurations you want to modify. On the **Server Group Settings** page, click **Modify**.
5. In the **Configurations** section, modify the Logtail configuration that you want to apply to the server group and click **Save**.

After a Logtail configuration is added, it is delivered to Logtail on each server in the server group. After a Logtail configuration is removed, it is removed from Logtail.



23.3.1.3.9. Manage a Logtail configuration

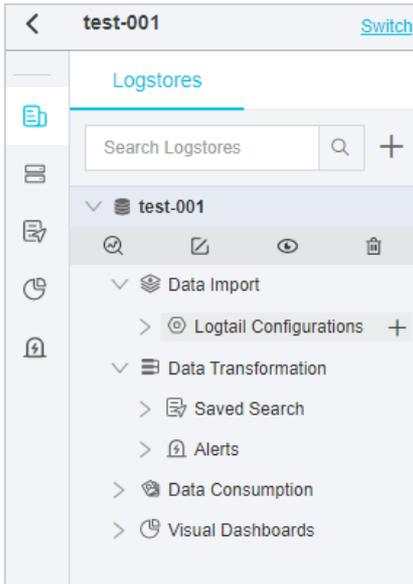
This topic describes how to manage a Logtail configuration in the Log Service console. Before you can collect logs from a server, you must install Logtail on the server. After you install Logtail, you must create a Logtail configuration in the Log Service console and apply the Logtail configuration to the server. You can create and modify Logtail configurations in Logstores.

Create a Logtail configuration

For information about how to create a Logtail configuration in the Log Service console, see [Configure text log collection](#).

View Logtail configurations

1. [Log on to the Log Service console](#).
2. Find the target project in the project list and click the project name.
3. In the left-side navigation pane, click the closing angle bracket (>) next to the target Logstore and choose **Data Import > Logtail Configurations**. Each item under **Logtail Configurations** indicates a Logtail configuration.



Modify a Logtail configuration

Under **Logtail Configurations**, click the name of the Logtail configuration. On the **Logtail Config** page, click **Modify**.

You can also change the log collection mode of the Logtail configuration, and then apply the Logtail configuration to the server group again. The process of modifying a Logtail configuration is the same as the process of creating a Logtail configuration.

Delete a Logtail configuration

Click the  icon next to the Logtail configuration, and then select **Delete**.

After the Logtail configuration is deleted, it is disassociated from the server group. Logtail no longer collects logs specified by the Logtail configuration.

23.3.1.3.10. Configure an account ID on a server

This topic describes how to configure the ID of an Apsara Stack tenant account on a server.

Prerequisites

Logtail is installed on the server. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).

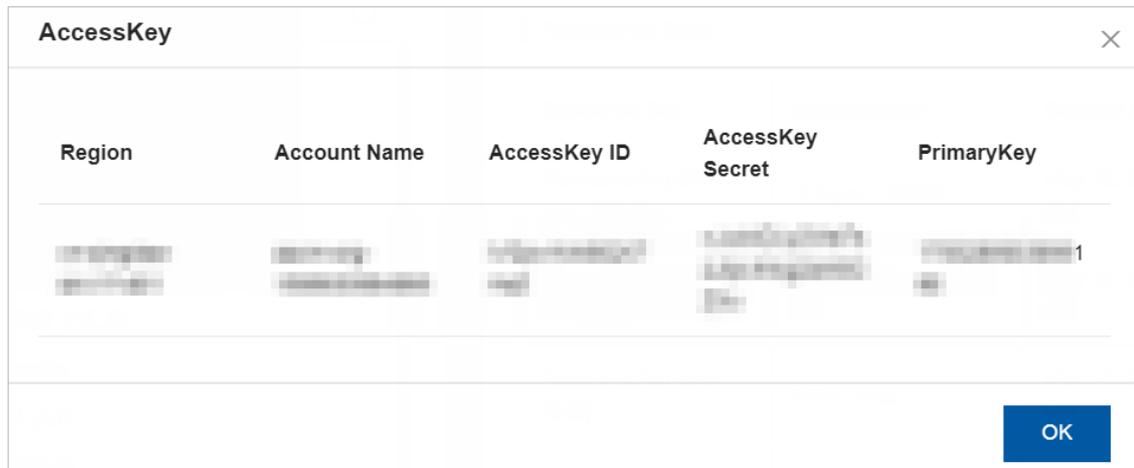
Context

You must configure the ID if the server is an on-premises server or an ECS instance that belong to another Apsara Stack tenant account. You must also configure the ID if the server is a cloud server that is provided by a third-party vendor. Then, the Apsara Stack tenant account can use Logtail to collect logs from the server. If you do not configure the account ID on the server, Log Service cannot receive the heartbeat of the server and Logtail cannot collect logs from the server.

Procedure

1. View the ID of your Apsara Stack tenant account.
 - i. Log on to the Apsara Uni-manager Management Console.
For more information, see [Log on to the Log Service console](#).
 - ii. In the top navigation bar, click **Enterprise**.
 - iii. In the left-side navigation pane, click **Organizations**.

- iv. Select the destination account and click **Obtain an accesskey**
- v. In the **AccessKey** dialog box, view the account ID.



2. Log on to the server and configure the account ID on the server.

- o Linux server:

In the `/etc/ilogtail/users` directory, create a file. Set the name of the file to the account ID. If the directory does not exist, create the directory first. You can configure multiple account IDs for a server. Examples:

```
touch /etc/ilogtail/users/1*****
touch /etc/ilogtail/users/1*****
```

If you no longer need to collect logs from the server to a Log Service project, run the following command to delete the account ID:

```
rm /etc/ilogtail/users/1*****
```

- o Windows server:

In the `C:\LogtailData\users` directory, create a file. Set the name of the file to the account ID. To delete the account ID, delete the file.

For example, you can delete the file of an account ID from the `C:\LogtailData\users\1*****` directory.

Note

- After you configure the ID of an Apsara Stack tenant account on a server, the account is authorized to use Logtail to collect logs from the server. If an account is no longer used to collect logs from the server, delete the account ID file from the server at the earliest opportunity.
- After you configure or delete an account ID, the change takes effect within 1 minute.

23.3.1.4. Text logs

23.3.1.4.1. Configure text log collection

This topic describes how to configure Logtail in the Log Service console to collect text logs from specified servers.

Prerequisites

Logtail is installed. Logtail can be installed on a Windows or Linux operating system. For more information, see [Install Logtail in Linux](#) and [Install Logtail in Windows](#).

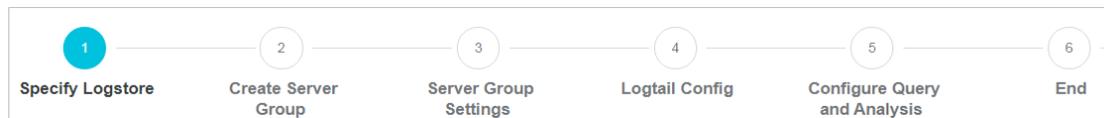
Limits

- Each log file can be collected by using only one Logtail configuration. If you want to collect a log file by using more than one Logtail configuration, we recommend that you use symbolic links. For example, to collect a log file by using two Logtail configurations in the `/home/log/nginx/log` directory, you can use the original log path for one Logtail configuration. Then, run the `ln -s /home/log/nginx/log /home/log/nginx/link_log` command to create a symbolic link for this directory and use the symbolic link as the log path for the other Logtail configuration.
- Logtail supports only Windows or Linux operating systems. For more information, see [Logtail overview](#).

Logtail configuration procedure

You can specify Logtail configurations in the Log Service console. Logtail supports various collection modes, such as simple mode, NGINX configuration mode, Apache configuration mode, IIS configuration mode, delimiter mode, JSON mode, and full regex mode.

Configuration procedure



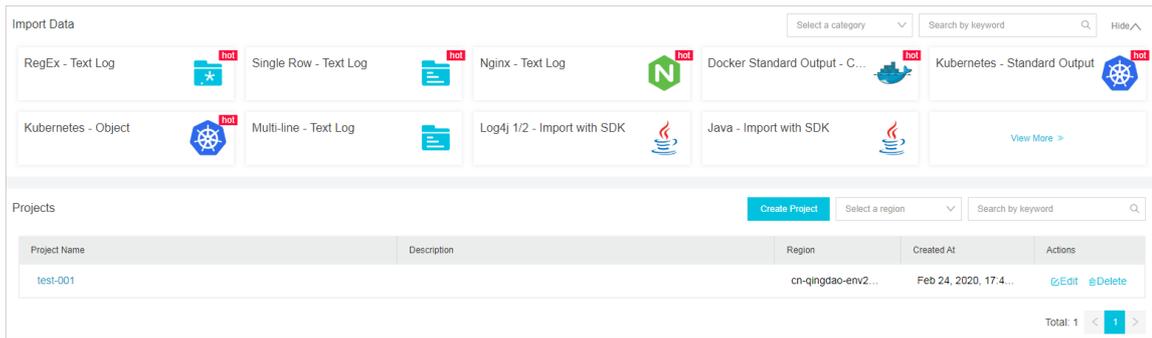
Collection modes

Logtail supports various collection modes, such as simple mode, NGINX configuration mode, Apache configuration mode, IIS configuration mode, delimiter mode, JSON mode, and full regex mode.

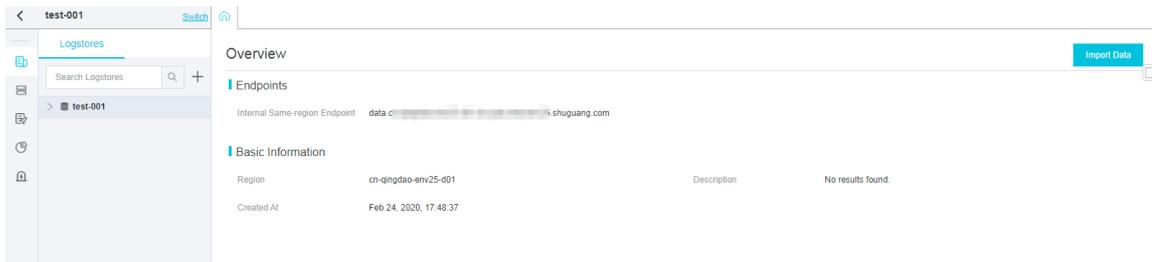
- Simple mode
Logtail can be used to collect logs in the simple mode. For more information, see [Collect logs by line](#).
- Full regex mode
Logtail can be used to collect logs in the full regex mode. For more information, see [Use regular expressions to collect logs](#).
- Delimiter mode
Logtail can be used to collect logs in the delimiter mode. For more information, see [Collect DSV formatted logs](#).
- JSON mode
Logtail can be used to collect logs in the JSON mode. For more information, see [Collect JSON logs](#).
- NGINX configuration mode
Logtail can be used to collect logs in the NGINX configuration mode. For more information, see [Collect NGINX logs](#).
- IIS configuration mode
Logtail can be used to collect logs in the IIS configuration mode. For more information, see [Collect IIS logs](#).
- Apache configuration mode
Logtail can be used to collect logs in the Apache configuration mode. For more information, see [Collect Apache logs](#).

Procedure

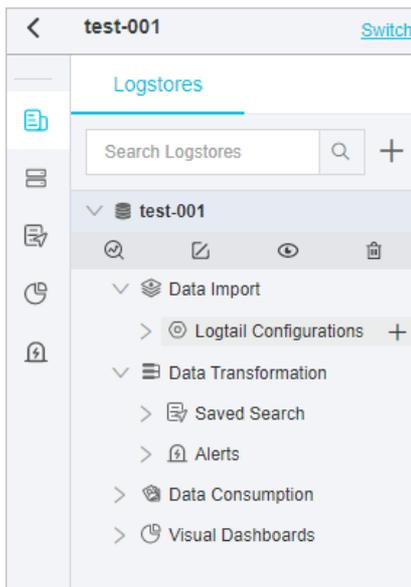
1. [Log on to the Log Service console](#).
2. Select a data source.
You can use one of the following three methods to select a data source:
 - On the homepage of the Log Service console, select a data source in the **Import Data** section.



- o In the **Projects** section, click a project name. On the **Overview** page, click **Import Data**.



- o On the **Logstores** tab in the left-side navigation pane, find a Logstore and click the closing angle bracket (>) in front of the Logstore name. Then, click the plus sign (+) next to **Data Import**.



Select a data source based on your business requirements. Log Service supports the following log sources of text logs: **RegEx-Text Log**, **Single Row-Text Log**, **Multi-Row-Text Log**, **Delimiter Mode-Text Log**, **JSON-Text Log**, **Nginx-Text Log**, **IIS-Text Log**, and **Apache-Text Log**.

3. Select a Logstore, and then click **Next**.

Select an existing project and Logstore. You can also click **Create Now** to create a project and Logstore. For more information, see [Manage a Logstore](#).

If you enter the configuration procedure by clicking the plus sign (+) next to **Data Import** under a Logstore on the Logstores tab, the system skips this step.

4. Create a server group.

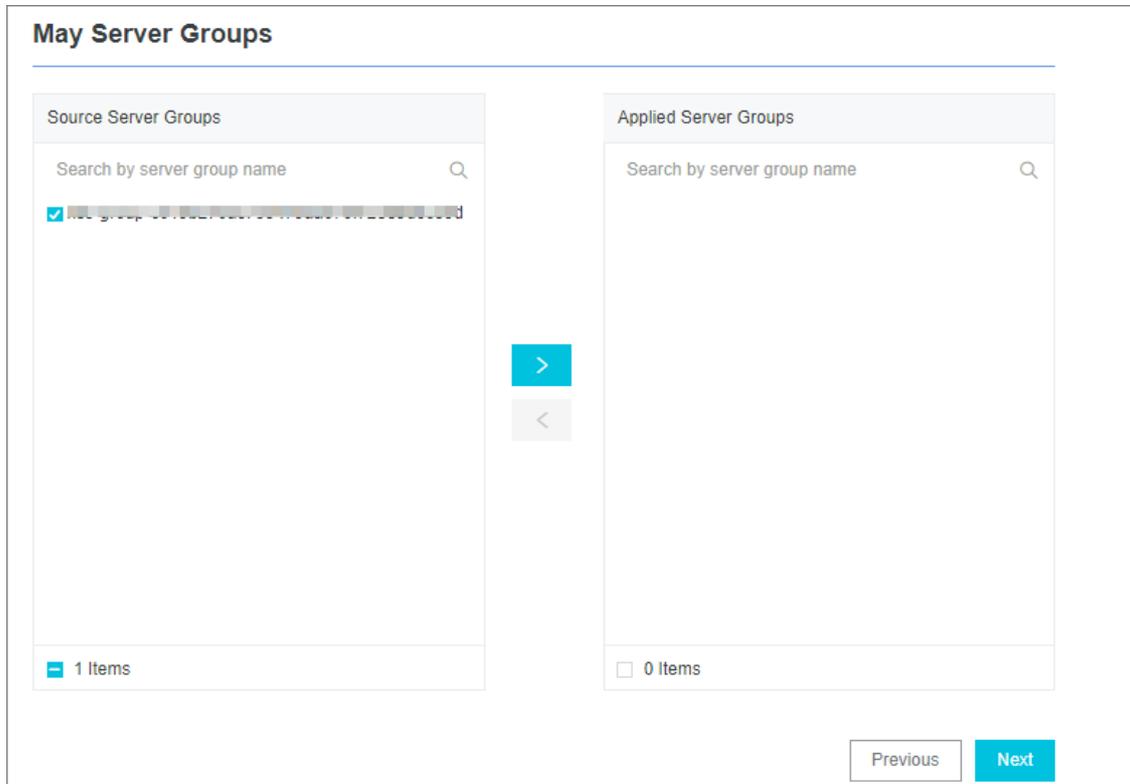
Before you create a server group, ensure that Logtail is installed.

Install Logtail as prompted. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).

After you install Logtail, click **Complete Installation** to create a server group. For more information, see [Overview](#). If you have created a server group, click **Use Existing Server Groups** to select the server group.

5. Configure the server group, and then click Next.

Select a server group and move the group from **Source Server Groups** to **Applied Server Groups**.



6. Specify Logtail parameters.

Logtail parameters vary based on collection modes. For more information, see the relevant parameters for specific collection modes.

7. (Optional)Specify **Advanced Options** and click **Next**.

Specify **Advanced Options** based on your business requirements. We recommend that you do not modify the default settings unless otherwise required.

| Parameter | Description |
|---------------------------|---|
| Enable Plug-in Processing | Specifies whether to enable plug-in processing. If you turn on this switch, you can use the plug-in of Logtail to process text logs. |
| Upload Raw Log | Specifies whether to upload raw logs. If you turn on this switch, raw logs are written to the <code>__raw__</code> field and uploaded with the parsed logs. |
| Topic Generation Mode | <ul style="list-style-type: none"> ◦ Null - Do not generate topic: This mode is selected by default. In this mode, the topic is set to an empty string and you can query logs without the need to enter a topic. ◦ Server Group Topic Attributes: This mode is used to differentiate log data that is generated by different frontend servers. ◦ File Path RegEx: If you select this mode, you must enter a value in the Custom RegEx field to extract part of the path as the topic. This mode is used to differentiate log data that is generated by users or instances. |

| Parameter | Description |
|----------------------|--|
| Custom RegEx | Specifies a custom regular expression. If you select File Path RegEx for Topic Generation Mode , you must enter a custom regular expression. |
| Log File Encoding | <ul style="list-style-type: none"> utf8: indicates UTF-8 encoding. gbk: indicates GBK encoding. |
| Timezone | <p>Specifies the time zone where logs are collected.</p> <ul style="list-style-type: none"> System Timezone: This option is selected by default. It indicates that the time zone where logs are collected is the same as the time zone to which the server belongs. Custom: Select a time zone. |
| Timeout | <p>If a log file is not updated within a specific period of time, Logtail considers the file to be timed out.</p> <ul style="list-style-type: none"> Never: All log files are continuously monitored and never time out. 30 Minute Timeout: If a log file is not updated within 30 minutes, Logtail considers the log file to be timed out and no longer monitors the file. |
| Filter Configuration | <p>Only logs that meet all filter conditions are collected.</p> <p>Examples:</p> <ul style="list-style-type: none"> Collect logs that meet a condition: Set the condition to <code>Key:level Regex:WARNING ERROR</code>. It indicates that only logs with the severity level of WARNING or ERROR are collected. Filter logs that do not meet a condition: <ul style="list-style-type: none"> Set the condition to <code>Key:level Regex:^(?!.*(INFO DEBUG)).*</code>. It indicates that logs with the severity level of INFO or DEBUG are not collected. Set the condition to <code>Key:url Regex:.*(?!.*(healthcheck)).*</code>. It indicates that logs whose URL contains the keyword healthcheck are not collected. For example, logs in which the key is url and the value is <code>/inner/healthcheck/jiankong.html</code> are not collected. |

8. Configure an index.

Configure an index based on your business requirements. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

After you complete the settings, you can start to collect logs.

23.3.1.4.2. Collect logs by line

This topic describes how to collect logs by line and configure indexes. You can specify the required settings in the Log Service console.

Context

To collect logs by line, you must select the simple mode. The simple mode can be divided into two types:

- Singleline mode

In this mode, each line of log data is considered as a log. Two logs in a log file are separated by a line break.

Logtail does not extract log fields in this mode. The default regular expression is `(.*)`. Logtail records the system time of the current server as the timestamp of a log. You can modify or manage advanced Logtail settings after you have completed the configuration procedure. For more information, see [Manage a Logtail configuration](#).

- Multi-line mode

In the multi-line mode, a regular expression is used to match the first line of a log. The settings in the multi-line mode are similar to the settings in the full regex mode. For more information, see [Use regular expressions to collect logs](#).

Procedure

1. [Log on to the Log Service console](#).

2. Select a data source.

Select **Single Row-Text Log**.

3. Select a Logstore, and then click Next.

Select an existing project and Logstore. You can also click **Create Now** to create a project and Logstore. For more information, see [Manage a Logstore](#).

If you enter the configuration procedure by clicking the plus sign (+) next to **Data Import** under a Logstore on the Logstores tab, the system skips this step.

4. Create a server group.

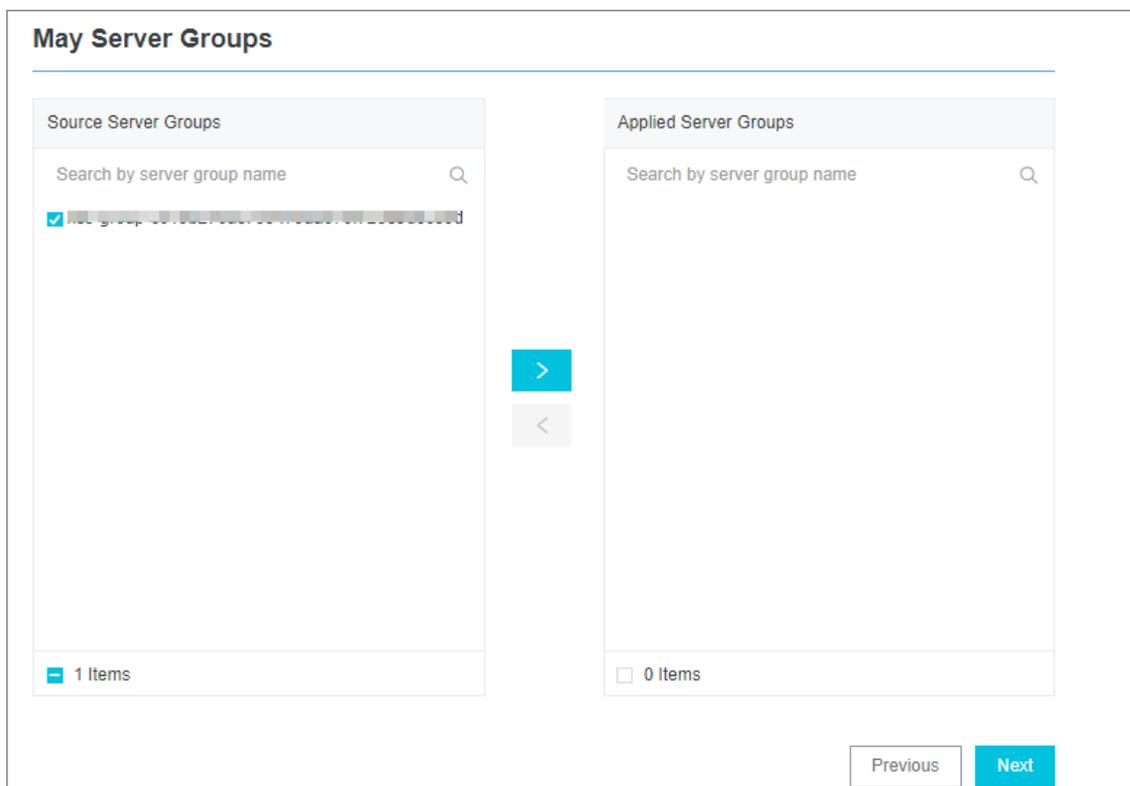
Before you create a server group, ensure that Logtail is installed.

Install Logtail as prompted. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).

After you install Logtail, click **Complete Installation** to create a server group. For more information, see [Overview](#). If you have created a server group, click **Use Existing Server Groups** to select the server group.

5. Configure the server group, and then click Next.

Select a server group and move the group from **Source Server Groups** to **Applied Server Groups**.



6. Create a Logtail configuration.

The following table lists the Logtail parameters.

| Parameter | Description |
|------------------------------------|---|
| Config Name | <p>The configuration name must be 3 to 128 characters in length, and can contain lowercase letters, digits, hyphens (-), and underscores (_). It must start and end with a lowercase letter or digit.</p> <p> Note The configuration name cannot be modified after it is created.</p> |
| Log Path | <p>The directory and name of the log file.</p> <ul style="list-style-type: none"> The specified log file name can be a complete file name or a file name that contains wildcards. Recursive directory matching is used in the log file search. If this matching method is applied, all files that match the specified file name in the specified directory and its sub-directories are monitored. <ul style="list-style-type: none"> Example 1: <code>/apsara/nuwa/.../*.log</code> indicates the files whose extension is <code>.log</code> in the <code>/apsara/nuwa</code> directory and its sub-directories are monitored. Example 2: <code>/var/logs/app_*/.../*.log*</code> indicates the files whose file name contains <code>.log</code> in the following directories are monitored: the sub-directories of the <code>/var/logs</code> directory that match the <code>app_*</code> format and the sub-directories of these matching sub-directories. <p> Note</p> <ul style="list-style-type: none"> Each log file can be collected by using only one Logtail configuration. Only the asterisk (<code>*</code>) and question mark (<code>?</code>) can be used as wildcards in the log path. |
| Docker File | <p>If the log file to be collected is in a Docker container, you can configure the internal path and container tag. Logtail monitors the creation and destruction of the container, filters logs of the container based on the tag, and collects the filtered logs.</p> |
| Mode | <p>If you have specified Single Row-Text Log for the data source, the default mode is Simple Mode. You can change the mode.</p> |
| Maximum Directory Monitoring Depth | <p>The maximum number of directory layers that can be recursively monitored when logs are collected from the data source. Valid values: 0 to 1000. The value 0 indicates that only the directory that is specified in the log path is monitored.</p> |

7. (Optional)Specify **Advanced Options** and click **Next**.

Specify **Advanced Options** based on your business requirements. We recommend that you do not modify the default settings unless otherwise required.

| Parameter | Description |
|---------------------------|---|
| Enable Plug-in Processing | <p>Specifies whether to enable plug-in processing. If you turn on this switch, you can use the plug-in of Logtail to process text logs.</p> |

| Parameter | Description |
|-----------------------|--|
| Upload Raw Log | Specifies whether to upload raw logs. If you turn on this switch, raw logs are written to the <code>__raw__</code> field and uploaded with the parsed logs. |
| Topic Generation Mode | <ul style="list-style-type: none"> ◦ Null - Do not generate topic: This mode is selected by default. In this mode, the topic is set to an empty string and you can query logs without the need to enter a topic. ◦ Server Group Topic Attributes: This mode is used to differentiate log data that is generated by different frontend servers. ◦ File Path RegEx: If you select this mode, you must enter a value in the Custom RegEx field to extract part of the path as the topic. This mode is used to differentiate log data that is generated by users or instances. |
| Custom RegEx | Specifies a custom regular expression. If you select File Path RegEx for Topic Generation Mode , you must enter a custom regular expression. |
| Log File Encoding | <ul style="list-style-type: none"> ◦ <code>utf8</code>: indicates UTF-8 encoding. ◦ <code>gbk</code>: indicates GBK encoding. |
| Timezone | <p>Specifies the time zone where logs are collected.</p> <ul style="list-style-type: none"> ◦ System Timezone: This option is selected by default. It indicates that the time zone where logs are collected is the same as the time zone to which the server belongs. ◦ Custom: Select a time zone. |
| Timeout | <p>If a log file is not updated within a specific period of time, Logtail considers the file to be timed out.</p> <ul style="list-style-type: none"> ◦ Never: All log files are continuously monitored and never time out. ◦ 30 Minute Timeout: If a log file is not updated within 30 minutes, Logtail considers the log file to be timed out and no longer monitors the file. |
| Filter Configuration | <p>Only logs that meet all filter conditions are collected.</p> <p>Examples:</p> <ul style="list-style-type: none"> ◦ Collect logs that meet a condition: Set the condition to <code>Key:level Regex:WARNING ERROR</code>. It indicates that only logs with the severity level of WARNING or ERROR are collected. ◦ Filter logs that do not meet a condition: <ul style="list-style-type: none"> ▪ Set the condition to <code>Key:level Regex:^(?!.*(INFO DEBUG)).*</code>. It indicates that logs with the severity level of INFO or DEBUG are not collected. ▪ Set the condition to <code>Key:url Regex:.*(?!.*(healthcheck)).*</code>. It indicates that logs whose URL contains the keyword healthcheck are not collected. For example, logs in which the key is url and the value is <code>/inner/healthcheck/jiankong.html</code> are not collected. |

8. Configure an index.

Configure an index based on your business requirements. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

After you complete the settings, you can start to collect logs by line.

23.3.1.4.3. Use regular expressions to collect logs

This topic describes how to collect logs by using regular expressions and configure indexes. You can specify the required settings in the Log Service console.

Context

If you need to collect multi-line logs and extract fields from logs, we recommend that you use regular expressions. Log Service can generate a regular expression based on a sample log that you enter in the **Import Data** wizard. However, you must modify the expression to match fields in the sample log as expected. For more information, see [How do I test a regular expression?](#).

Procedure

1. [Log on to the Log Service console.](#)

2. Select a data source.

Select **RegEx-Text Log**.

3. Select a Logstore, and then click Next.

Select an existing project and Logstore. You can also click **Create Now** to create a project and Logstore. For more information, see [Manage a Logstore](#).

If you enter the configuration procedure by clicking the plus sign (+) next to **Data Import** under a Logstore on the Logstores tab, the system skips this step.

4. Create a server group.

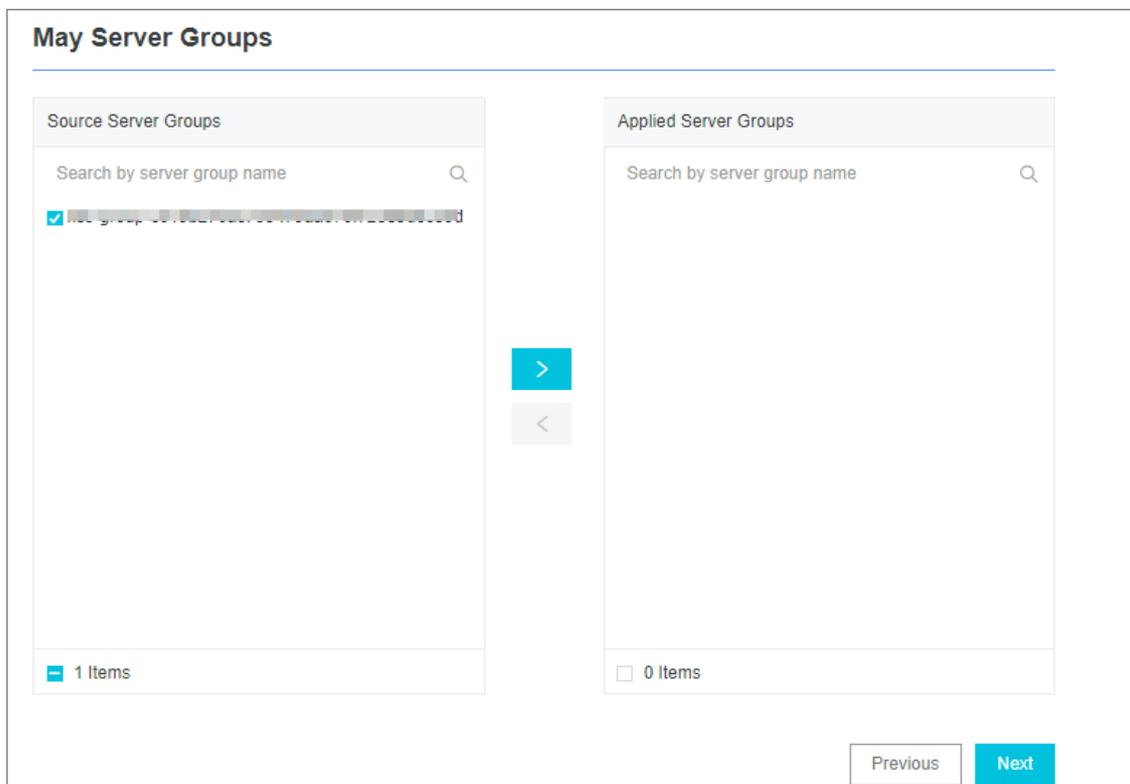
Before you create a server group, ensure that Logtail is installed.

Install Logtail as prompted. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).

After you install Logtail, click **Complete Installation** to create a server group. For more information, see [Overview](#). If you have created a server group, click **Use Existing Server Groups** to select the server group.

5. Configure the server group, and then click Next.

Select a server group and move the group from **Source Server Groups** to **Applied Server Groups**.



6. Create a Logtail configuration.

The following table lists the Logtail parameters.

| Parameter | Description |
|---------------------------|---|
| Config Name | <p>The configuration name must be 3 to 128 characters in length, and can contain lowercase letters, digits, hyphens (-), and underscores (_). It must start and end with a lowercase letter or digit.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note The configuration name cannot be modified after it is created.</p> </div> |
| Log Path | <p>The directory and name of the log file.</p> <ul style="list-style-type: none"> ○ The specified log file name can be a complete file name or a file name that contains wildcards. ○ Recursive directory matching is used in the log file search. If this matching method is applied, all files that match the specified file name in the specified directory and its sub-directories are monitored. <ul style="list-style-type: none"> ■ Example 1: <code>/apsara/nuwa/.../*.log</code> indicates the files whose extension is <code>.log</code> in the <code>/apsara/nuwa</code> directory and its sub-directories are monitored. ■ Example 2: <code>/var/logs/app_.../*.log*</code> indicates the files whose file name contains <code>.log</code> in the following directories are monitored: the sub-directories of the <code>/var/logs</code> directory that match the <code>app_*</code> format and the sub-directories of these matching sub-directories. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note</p> <ul style="list-style-type: none"> ■ Each log file can be collected by using only one Logtail configuration. ■ Only the asterisk (<code>*</code>) and question mark (<code>?</code>) can be used as wildcards in the log path. </div> |
| Docker File | <p>If the log file to be collected is in a Docker container, you can configure the internal path and container tag. Logtail monitors the creation and destruction of the container, filters logs of the container based on the tag, and collects the filtered logs.</p> |
| Mode | <p>If you have specified RegEx-Text Log for the data source, the default mode is Full Regex Mode. You can change the mode.</p> |
| Singleline | <p>The singleline mode is enabled by default. In this mode, logs are separated by line. To collect multi-line logs, such as Java program logs, you must disable the Singleline mode and configure Regex to Match First Line.</p> |
| Log Sample | <p>Enter a sample log that is retrieved from the data source. Then, Log Service generates a regular expression.</p> |
| Regex to Match First Line | <p>You can click Auto Generate or Manual. After you enter a sample log entry and click Auto Generate, the system generates a regular expression. If no regular expression is generated, you can switch to the manual mode and enter a regular expression for verification.</p> |

| Parameter | Description |
|------------------------------------|--|
| Extract Field | To analyze and process specific fields in logs, you can turn on the Extract Field switch. Then, the specified fields are converted to key-value pairs and sent to Log Service. You must specify a regular expression to parse the log content. |
| RegEx | <p>If you turn on the Extract Field switch, you must specify this setting.</p> <ul style="list-style-type: none"> Automatically generate a regular expression You can select the fields to be extracted from the sample log and then click Generate Regular Expression. The system generates a regular expression. Enter a regular expression You can also enter a regular expression. Click Manually to switch to the manual mode. After you enter a regular expression, click Validate to check whether the regular expression can parse the log content. For more information, see How do I test a regular expression?. |
| Extracted Content | <p>If you turn on the Extract Field switch, you must specify this setting.</p> <p>After a regular expression is automatically generated or manually specified, you must specify the key name for each extracted field.</p> |
| Use System Time | <p>If you turn on the Extract Field switch, you must specify this setting.</p> <p>If you turn off the Use System Time switch, you must specify a field as the time field and name this field <code>time</code>. After you specify the <code>time</code> field, click Auto Generate in the Time Conversion Format field to automatically parse the time. For more information, see Configure the time format.</p> |
| Drop Failed to Parse Logs | <p>Specifies whether to upload logs to Log Service if the logs fail to be parsed.</p> <ul style="list-style-type: none"> If you turn on this switch, logs that fail to be parsed are not uploaded to Log Service. If you turn off this switch, raw logs are uploaded to Log Service when logs fail to be parsed. |
| Maximum Directory Monitoring Depth | The maximum number of directory layers that can be recursively monitored when logs are collected from the data source. Valid values: 0 to 1000. The value 0 indicates that only the directory specified in the log path is monitored. |

7. (Optional)Specify **Advanced Options** and click **Next**.

Specify **Advanced Options** based on your business requirements. We recommend that you do not modify the default settings unless otherwise required.

| Parameter | Description |
|---------------------------|---|
| Enable Plug-in Processing | Specifies whether to enable plug-in processing. If you turn on this switch, you can use the plug-in of Logtail to process text logs. |
| Upload Raw Log | Specifies whether to upload raw logs. If you turn on this switch, raw logs are written to the <code>__raw__</code> field and uploaded with the parsed logs. |

| Parameter | Description |
|-----------------------|---|
| Topic Generation Mode | <ul style="list-style-type: none"> ◦ Null - Do not generate topic: This mode is selected by default. In this mode, the topic is set to an empty string and you can query logs without the need to enter a topic. ◦ Server Group Topic Attributes: This mode is used to differentiate log data that is generated by different frontend servers. ◦ File Path RegEx: If you select this mode, you must enter a value in the Custom RegEx field to extract part of the path as the topic. This mode is used to differentiate log data that is generated by users or instances. |
| Custom RegEx | Specifies a custom regular expression. If you select File Path RegEx for Topic Generation Mode , you must enter a custom regular expression. |
| Log File Encoding | <ul style="list-style-type: none"> ◦ utf8: indicates UTF-8 encoding. ◦ gbk: indicates GBK encoding. |
| Timezone | <p>Specifies the time zone where logs are collected.</p> <ul style="list-style-type: none"> ◦ System Timezone: This option is selected by default. It indicates that the time zone where logs are collected is the same as the time zone to which the server belongs. ◦ Custom: Select a time zone. |
| Timeout | <p>If a log file is not updated within a specific period of time, Logtail considers the file to be timed out.</p> <ul style="list-style-type: none"> ◦ Never: All log files are continuously monitored and never time out. ◦ 30 Minute Timeout: If a log file is not updated within 30 minutes, Logtail considers the log file to be timed out and no longer monitors the file. |
| Filter Configuration | <p>Only logs that meet all filter conditions are collected.</p> <p>Examples:</p> <ul style="list-style-type: none"> ◦ Collect logs that meet a condition: Set the condition to <code>Key:level Regex:WARNING ERROR</code>. It indicates that only logs with the severity level of WARNING or ERROR are collected. ◦ Filter logs that do not meet a condition: <ul style="list-style-type: none"> ▪ Set the condition to <code>Key:level Regex:^(?!.*(INFO DEBUG)).*</code>. It indicates that logs with the severity level of INFO or DEBUG are not collected. ▪ Set the condition to <code>Key:url Regex:^(?!.*(healthcheck)).*</code>. It indicates that logs whose URL contains the keyword healthcheck are not collected. For example, logs in which the key is url and the value is <code>/inner/healthcheck/jiankong.html</code> are not collected. |

8. Configure an index.

Configure an index based on your business requirements. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

After you complete the settings, you can start to use regular expressions to collect logs.

23.3.1.4.4. Collect DSV formatted logs

This topic describes how to collect delimiter-separated values (DSV) formatted logs and configure indexes. You can specify the required settings in the Log Service console.

Context

DSV formatted logs use line feeds as boundaries. Each line is a log entry. The fields of each log entry are separated with a fixed delimiter. The characters that can be used as delimiters include the tab (`\t`), space, vertical bar (`|`), comma (`,`), and semicolon (`;`). A field that contains a delimiter must be enclosed in double quotation marks (`"`), which are used as quotes.

Log formats

Common DSV formatted logs include comma-separated values (CSV) and tab-separated values (TSV) formatted logs.

A delimiter can contain a **single character** or **multiple characters**.

Single-character delimiter

You can specify a single-character delimiter and a quote in the console.

- **Delimiter:** The fields of each log entry are separated with a single-character delimiter, such as the tab (`\t`), vertical bar (`|`), space, comma (`,`), and semicolon (`;`). You can also specify a non-printable character as the delimiter.

Note A double quotation mark (`"`) cannot be used as a delimiter.

If a double quotation mark (`"`) is included in a log entry but not used as a quote, it must be escaped and processed as double quotation marks (`""`). When Log Service parse logs, it restores double quotation marks (`""`) into a double quotation mark (`"`). You can use a double quotation mark (`"`) on each boundary of a field as a quote. You can also use a double quotation mark (escaped as `""`) in the content of a field. If the use of a double quotation mark (`"`) does not comply with the defined format, you can use the simple mode or full regex mode to parse fields.

For example, assume that you use commas (`,`) as delimiters and include double quotation marks (`""`) and commas (`,`) in a field. Enclose the field with quotes and escape the double quotation marks into `""`. For example, a processed log is `1999,Chevy,"Venture ""Extended Edition, Very Large""",5000.00`. The log can be parsed into five fields: `1999`, `Chevy`, `Venture "Extended Edition, Very Large"`, empty field, and `5000.00`.

- **Quote:** If a log field contains delimiters, you must specify a quote to enclose the field. Otherwise, the field cannot be parsed as expected. Log Service parses the content enclosed in quotes as one field. Only delimiters can exist between fields.

You can use one of the following characters as the quote: tab (`\t`), vertical bar (`|`), space, comma (`,`), semicolon (`;`), and non-printable characters.

For example, a log is `1997,Ford,E350,"ac,abs,moon",3000.00`. In this example, the comma (`,`) is used as the delimiter and the double quotation mark (`"`) is used as the quote. The log entry can be parsed into five fields: `1997`, `Ford`, `E350`, `ac,abs,moon`, and `3000.00`. Among the five fields, `ac,abs,moon` enclosed in quotes is regarded as one field.

Note Log Service allows you to use a non-printable character as a delimiter or quote. Non-printable characters are characters whose decimal ASCII codes are within the range of 1 to 31 and 127. If you use a non-printable character as a delimiter or quote, you must find the hexadecimal ASCII code of this character and enter the character in the following format: `0xthe hexadecimal ASCII code of the non-printable character`. For example, to use the non-printable character whose decimal ASCII code is 1 and hexadecimal ASCII code is 01, you must enter `0x01`.

Multi-character delimiter

Each multi-character delimiter can contain two or three characters, such as `||`, `&&&`, and `^_^`. If you specify a multi-character delimiter, Log Service parses logs only based on the delimiter. You do not need to use quotes to enclose log fields.

 **Note** You must ensure that log fields do not contain the delimiter. Otherwise, Log Service cannot parse these fields as expected.

For example, if the delimiter is set to `&&`, the log `1997&&Ford&&E350&&ac&abs&moon&&3000.00` is parsed into five fields: `1997`, `Ford`, `E350`, `ac&abs&moon`, and `3000.00`.

Sample logs

• Single-character delimiter

```
05/May/2016:13:30:28,10.10.*.*,"POST /PutData? Category=YunOsAccountOpLog&AccessKeyId=*****&Date=Fri%2C%2028%20Jun%202013%2006%3A53%3A30%20GMT&Topic=raw&Signature=*****  
** HTTP/1.1",200,18204,aliyun-sdk-java  
05/May/2016:13:31:23,10.10.*.*,"POST /PutData? Category=YunOsAccountOpLog&AccessKeyId=*****&Date=Fri%2C%2028%20Jun%202013%2006%3A53%3A30%20GMT&Topic=raw&Signature=*****  
** HTTP/1.1",401,23472,aliyun-sdk-java
```

• Multi-character delimiter

```
05/May/2016:13:30:28&&10.200.*.*&&POST /PutData? Category=YunOsAccountOpLog&AccessKeyId=*****  
**&&Date=Fri%2C%2028%20Jun%202013%2006%3A53%3A30%20GMT&Topic=raw&Signature=pD12XYLmGxKQ%2Bmk  
d6x7hAgQ7b1c%3D HTTP/1.1&&200&&18204&&aliyun-sdk-java  
05/May/2016:13:31:23&&10.200.*.*&&POST /PutData? Category=YunOsAccountOpLog&AccessKeyId=*****  
**&&Date=Fri%2C%2028%20Jun%202013%2006%3A53%3A30%20GMT&Topic=raw&Signature=*****  
***** HTTP/1.1&&401&&23472&&aliyun-sdk-java
```

Procedure

1. [Log on to the Log Service console](#).

2. Select a data source.

Select **Delimiter-Text Log**.

3. Select a Logstore, and then click Next.

Select an existing project and Logstore. You can also click **Create Now** to create a project and Logstore. For more information, see [Manage a Logstore](#).

If you enter the configuration procedure by clicking the plus sign (+) next to **Data Import** under a Logstore on the Logstores tab, the system skips this step.

4. Create a server group.

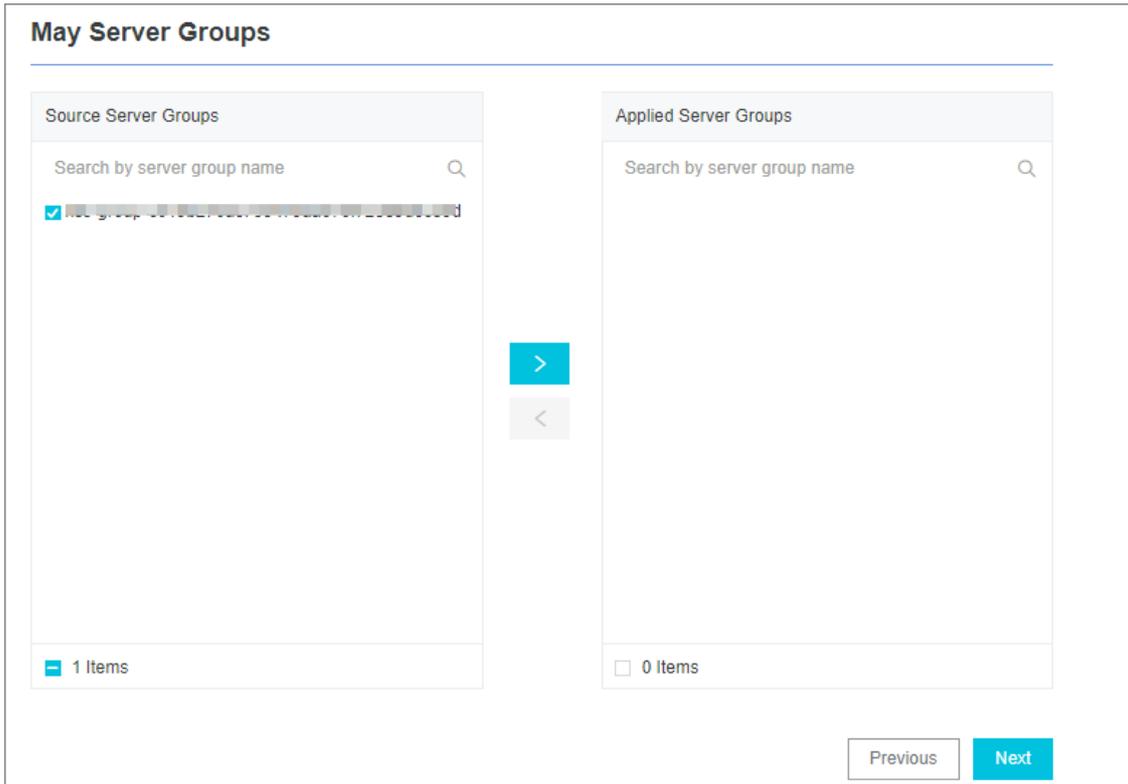
Before you create a server group, ensure that Logtail is installed.

Install Logtail as prompted. For more information, see [Inst all Logtail in Linux](#) or [Inst all Logtail in Windows](#).

After you install Logtail, click **Complete Installation** to create a server group. For more information, see [Overview](#). If you have created a server group, click **Use Existing Server Groups** to select the server group.

5. Configure the server group, and then click Next.

Select a server group and move the group from **Source Server Groups** to **Applied Server Groups**.



6. Create a Logtail configuration file.

The following table lists the Logtail parameters.

| Parameter | Description |
|-------------|--|
| Config Name | <p>The name must be 3 to 128 characters in length, and can contain only lowercase letters, digits, hyphens (-), and underscores (_). The name must start and end with a lowercase letter or digit.</p> <p>Note The configuration name cannot be modified after it is created.</p> |

| Parameter | Description |
|-------------|--|
| Log Path | <p>The directory and name of the log file.</p> <ul style="list-style-type: none"> ◦ The specified log file name can be a complete file name or a file name that contains wildcards. ◦ Recursive directory matching is used in the log file search. If this matching method is applied, all files that match the specified file name in the specified directory and its sub-directories are monitored. <ul style="list-style-type: none"> ▪ Example 1: <code>/apsara/nuwa/.../*.log</code> indicates the files whose extension is <code>.log</code> in the <code>/apsara/nuwa</code> directory and its sub-directories are monitored. ▪ Example 2: <code>/var/logs/app_*.../*.log*</code> indicates the files whose file name contains <code>.log</code> in the following directories are monitored: the sub-directories of the <code>/var/logs</code> directory that match the <code>app_*</code> format and the sub-directories of these matching sub-directories. <div style="background-color: #e1f5fe; padding: 10px; margin-top: 10px;"> <p> Note</p> <ul style="list-style-type: none"> ▪ Each log file can be collected by using only one Logtail configuration file. ▪ Only the asterisk (<code>*</code>) and question mark (<code>?</code>) can be used as wildcards in the log path. </div> |
| Docker File | <p>If you collect logs from Docker containers, you can configure the paths and tags of the containers. Logtail monitors the containers when they are created and destroyed, filters the logs of the containers by tag, and collects the filtered logs.</p> |
| Mode | <p>If you have specified Delimiter-Text Log for the data source, the default mode is Delimiter Mode. You can change the mode.</p> |
| Log Sample | <p>Enter a sample log entry that is retrieved from a log source in an actual scenario. Then, Log Service extracts a regular expression from the log entry.</p> |
| Delimiter | <p>Select a delimiter.</p> <p>Select a delimiter based on the log format. Otherwise, logs may fail to be parsed.</p> <div style="background-color: #e1f5fe; padding: 10px; margin-top: 10px;"> <p> Note If you use a non-printable character as a delimiter, you must find the hexadecimal ASCII code of this character and enter the character in the following format: <code>0xthe hexadecimal ASCII code of the non-printable character</code> . For example, to use the non-printable character whose decimal ASCII code is 1 and hexadecimal ASCII code is 01, you must enter <code>0x01</code> .</p> </div> |

| Parameter | Description |
|------------------------------------|--|
| Quote | <p>Select a quote.</p> <p>Select a quote based on the log format. Otherwise, logs may fail to be parsed.</p> <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e6f2ff;"> <p> Note If you use a non-printable character as a quote, you must find the hexadecimal ASCII code of this character and enter the character in the following format: <code>0xthe hexadecimal ASCII code of the non-printable character</code>. For example, to use the non-printable character whose decimal ASCII code is 1 and hexadecimal ASCII code is 01, you must enter <code>0x01</code>.</p> </div> |
| Extracted Content | <p>After you enter a sample log and select a delimiter, Log Service extracts log fields based on the delimiter and defines the fields as values. You must specify a key for each value.</p> |
| Incomplete Entry Upload | <p>This feature specifies whether to upload a log entry whose number of parsed fields is less than the number of the specified keys. If you turn on this switch, the log entry is uploaded. Otherwise, the log entry is dropped.</p> <p>For example, if you set the delimiter to the vertical bar (), the log entry <code>11 22 33 44 55</code> can be parsed into the following fields: <code>11</code>, <code>22</code>, <code>33</code>, <code>44</code>, and <code>55</code>. You can set the keys to <code>A</code>, <code>B</code>, <code>C</code>, <code>D</code>, and <code>E</code>. If you turn on the Incomplete Entry Upload switch, the <code>55</code> field is uploaded as the value of the <code>D</code> key when Log Service collects the log entry <code>11 22 33 55</code>. If you turn off the Incomplete Entry Upload switch, Log Service drops the log entry because the fields and keys do not match.</p> |
| Use System Time | <p>If you turn on the Extract Field switch, you must specify this setting.</p> <p>If you turn off the Use System Time switch, you must specify a field as the time field and name this field <code>time</code>. After you specify the <code>time</code> field, click Auto Generate in the Time Conversion Format field to automatically parse the time. For more information, see Configure the time format.</p> |
| Drop Failed to Parse Logs | <p>Specifies whether to upload logs to Log Service if the logs fail to be parsed.</p> <ul style="list-style-type: none"> ◦ If you turn on this switch, logs that fail to be parsed are not uploaded to Log Service. ◦ If you turn off this switch, raw logs are uploaded to Log Service when logs fail to be parsed. |
| Maximum Directory Monitoring Depth | <p>The maximum number of directory layers that can be recursively monitored when logs are collected from the data source. Valid values: 0 to 1000. The value 0 indicates that only the directory that is specified in the log path is monitored.</p> |

7. (Optional)Specify **Advanced Options** and click **Next**.

Specify **Advanced Options** based on your business requirements. We recommend that you do not modify the default settings unless otherwise required.

| Parameter | Description |
|---------------------------|---|
| Enable Plug-in Processing | <p>Specifies whether to enable plug-in processing. If you turn on this switch, you can use the plug-in of Logtail to process text logs.</p> |

| Parameter | Description |
|-----------------------|---|
| Upload Raw Log | Specifies whether to upload raw logs. If you turn on this switch, raw logs are written to the <code>__raw__</code> field and uploaded with the parsed logs. |
| Topic Generation Mode | <ul style="list-style-type: none"> ◦ Null - Do not generate topic: This mode is selected by default. In this mode, the topic is set to an empty string and you can query logs without the need to enter a topic. ◦ Server Group Topic Attributes: This mode is used to differentiate log data that is generated by different frontend servers. ◦ File Path RegEx: If you select this mode, you must enter a value in the Custom RegEx field to extract part of the path as the topic. This mode is used to differentiate log data that is generated by users or instances. |
| Custom RegEx | Specifies a custom regular expression. If you select File Path RegEx for Topic Generation Mode , you must enter a custom regular expression. |
| Log File Encoding | <ul style="list-style-type: none"> ◦ <code>utf8</code>: indicates UTF-8 encoding. ◦ <code>gbk</code>: indicates GBK encoding. |
| Timezone | <p>Specifies the time zone where logs are collected.</p> <ul style="list-style-type: none"> ◦ System Timezone: This option is selected by default. It indicates that the time zone where logs are collected is the same as the time zone to which the server belongs. ◦ Custom: Select a time zone. |
| Timeout | <p>If a log file is not updated within a specific period of time, Logtail considers the file to be timed out.</p> <ul style="list-style-type: none"> ◦ Never: All log files are continuously monitored and never time out. ◦ 30 Minute Timeout: If a log file is not updated within 30 minutes, Logtail considers the log file to be timed out and no longer monitors the file. |
| Filter Configuration | <p>Only logs that meet all filter conditions are collected.</p> <p>Examples:</p> <ul style="list-style-type: none"> ◦ Collect logs that meet a condition: Set the condition to <code>Key:level Regex:WARNING ERROR</code>. It indicates that only logs with the severity level of WARNING or ERROR are collected. ◦ Filter logs that do not meet a condition: <ul style="list-style-type: none"> ▪ Set the condition to <code>Key:level Regex:^(?!.*(INFO DEBUG)).*</code>. It indicates that logs with the severity level of INFO or DEBUG are not collected. ▪ Set the condition to <code>Key:url Regex:^(?!.*(healthcheck)).*</code>. It indicates that logs whose URL contains the keyword healthcheck are not collected. For example, logs in which the key is url and the value is <code>/inner/healthcheck/jiankong.html</code> are not collected. |

8. Configure an index.

Configure an index based on your business requirements. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

After you complete the settings, you can start to collect DSV formatted logs.

23.3.1.4.5. Collect JSON logs

This topic describes how to use Logtail to collect JSON logs and configure indexes. You can specify the required settings in the Log Service console.

Context

Logtail can parse JSON objects from logs. It extracts the keys and values from the first layer of an object as the names and values of log fields. The valid data types of field values include object, array, and primitive data types such as string or number.

JSON logs can be written in the following two types of structures:

- Object: a collection of key-value pairs.
- Array: an ordered list of values.

Lines of JSON logs are separated with `\n`. Each line is extracted as a single log.

Logtail can parse only JSON logs of the object type. If you want to parse JSON logs of other types, such as JSON arrays, you must use regular expressions to extract the fields or specify the simple mode to collect logs by line.

Sample log

A sample JSON log is as follows:

```
{"url": "POST /PutData? Category=YunOsAccountOpLog&AccessKeyId=U0Ujpek*****&Date=Fri%2C%2028%20Jun%202013%2006%3A53%3A30%20GMT&Topic=raw&Signature=pD12XYLmGxKQ%2Bmkd6x7hAgQ7b1c%3D HTTP/1.1", "ip": "10.200.98.220", "user-agent": "aliyun-sdk-java", "request": {"status": "200", "latency": "18204"}, "time": "05/May/2016:13:30:28"}
{"url": "POST /PutData? Category=YunOsAccountOpLog&AccessKeyId=U0Ujpek*****&Date=Fri%2C%2028%20Jun%202013%2006%3A53%3A30%20GMT&Topic=raw&Signature=pD12XYLmGxKQ%2Bmkd6x7hAgQ7b1c%3D HTTP/1.1", "ip": "10.200.98.210", "user-agent": "aliyun-sdk-java", "request": {"status": "200", "latency": "10204"}, "time": "05/May/2016:13:30:29"}
```

Procedure

1. [Log on to the Log Service console](#).

2. Select a data source.

Select **JSON-Text Log**.

3. Select a Logstore, and then click Next.

Select an existing project and Logstore. You can also click **Create Now** to create a project and Logstore. For more information, see [Manage a Logstore](#).

If you enter the configuration procedure by clicking the plus sign (+) next to **Data Import** under a Logstore on the Logstores tab, the system skips this step.

4. Create a server group.

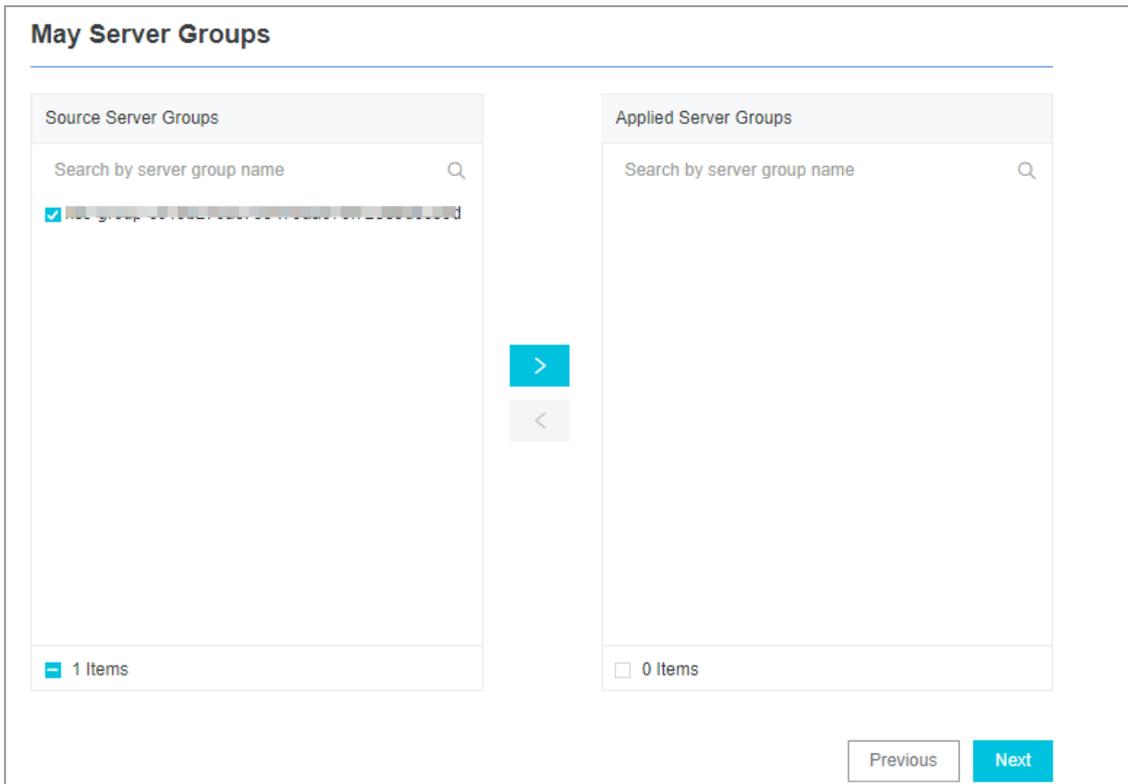
Before you create a server group, ensure that Logtail is installed.

Install Logtail as prompted. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).

After you install Logtail, click **Complete Installation** to create a server group. For more information, see [Overview](#). If you have created a server group, click **Use Existing Server Groups** to select the server group.

5. Configure the server group, and then click Next.

Select a server group and move the group from **Source Server Groups** to **Applied Server Groups**.



6. Create a Logtail configuration.

The following table lists the Logtail parameters.

| Parameter | Description |
|-------------|---|
| Config Name | <p>The configuration name must be 3 to 128 characters in length, and can contain lowercase letters, digits, hyphens (-), and underscores (_). It must start and end with a lowercase letter or digit.</p> <p>Note The configuration name cannot be modified after it is created.</p> |

| Parameter | Description |
|------------------------------------|---|
| Log Path | <p>The directory and name of the log file.</p> <ul style="list-style-type: none"> The specified log file name can be a complete file name or a file name that contains wildcards. Recursive directory matching is used in the log file search. If this matching method is applied, all files that match the specified file name in the specified directory and its sub-directories are monitored. <ul style="list-style-type: none"> Example 1: <code>/apsara/nuwa/.../*.log</code> indicates the files whose extension is <code>.log</code> in the <code>/apsara/nuwa</code> directory and its sub-directories are monitored. Example 2: <code>/var/logs/app_*/.../*.log*</code> indicates the files whose file name contains <code>.log</code> in the following directories are monitored: the sub-directories of the <code>/var/logs</code> directory that match the <code>app_*</code> format and the sub-directories of these matching sub-directories. <div style="background-color: #e1f5fe; padding: 10px; border: 1px solid #cfcfcf;"> <p> Note</p> <ul style="list-style-type: none"> Each log file can be collected by using only one Logtail configuration. Only the asterisk (<code>*</code>) and question mark (<code>?</code>) can be used as wildcards in the log path. </div> |
| Docker File | <p>If the log file to be collected is in a Docker container, you can configure the internal path and container tag. Logtail monitors the creation and destruction of the container, filters logs of the container based on the tag, and collects the filtered logs.</p> |
| Mode | <p>If you have specified JSON-Text Log for the data source, the default mode is JSON Mode. You can change the mode.</p> |
| Use System Time | <p>If you turn on the Extract Field switch, you must specify this parameter.</p> <p>If you turn off the Use System Time switch, you must specify a field as the time field and name this field <code>time</code>. After you specify the <code>time</code> field, click Auto Generate in the Time Conversion Format field to automatically parse the time. For more information, see Configure the time format.</p> |
| Drop Failed to Parse Logs | <p>Specifies whether to upload logs to Log Service if the logs fail to be parsed.</p> <ul style="list-style-type: none"> If you turn on this switch, logs that fail to be parsed are not uploaded to Log Service. If you turn off this switch, raw logs are uploaded to Log Service when logs fail to be parsed. |
| Maximum Directory Monitoring Depth | <p>The maximum number of directory layers that can be recursively monitored when logs are collected from the data source. Valid values: 0 to 1000. The value 0 indicates that only the directory specified in the log path is monitored.</p> |

7. (Optional)Specify **Advanced Options** and click **Next**.

Specify **Advanced Options** based on your business requirements. We recommend that you do not modify the default settings unless otherwise required.

| Parameter | Description |
|---------------------------|--|
| Enable Plug-in Processing | Specifies whether to enable plug-in processing. If you turn on this switch, you can use the plug-in of Logtail to process text logs. |
| Upload Raw Log | Specifies whether to upload raw logs. If you turn on this switch, raw logs are written to the <code>__raw__</code> field and uploaded with the parsed logs. |
| Topic Generation Mode | <ul style="list-style-type: none"> ◦ Null - Do not generate topic: This mode is selected by default. In this mode, the topic is set to an empty string and you can query logs without the need to enter a topic. ◦ Server Group Topic Attributes: This mode is used to differentiate log data that is generated by different frontend servers. ◦ File Path RegEx: If you select this mode, you must enter a value in the Custom RegEx field to extract part of the path as the topic. This mode is used to differentiate log data that is generated by users or instances. |
| Custom RegEx | Specifies a custom regular expression. If you select File Path RegEx for Topic Generation Mode , you must enter a custom regular expression. |
| Log File Encoding | <ul style="list-style-type: none"> ◦ <code>utf8</code>: indicates UTF-8 encoding. ◦ <code>gbk</code>: indicates GBK encoding. |
| Timezone | <p>Specifies the time zone where logs are collected.</p> <ul style="list-style-type: none"> ◦ System Timezone: This option is selected by default. It indicates that the time zone where logs are collected is the same as the time zone to which the server belongs. ◦ Custom: Select a time zone. |
| Timeout | <p>If a log file is not updated within a specific period of time, Logtail considers the file to be timed out.</p> <ul style="list-style-type: none"> ◦ Never: All log files are continuously monitored and never time out. ◦ 30 Minute Timeout: If a log file is not updated within 30 minutes, Logtail considers the log file to be timed out and no longer monitors the file. |
| Filter Configuration | <p>Only logs that meet all filter conditions are collected.</p> <p>Examples:</p> <ul style="list-style-type: none"> ◦ Collect logs that meet a condition: Set the condition to <code>Key:level Regex:WARNING ERROR</code>. It indicates that only logs with the severity level of WARNING or ERROR are collected. ◦ Filter logs that do not meet a condition: <ul style="list-style-type: none"> ▪ Set the condition to <code>Key:level Regex:^(?!.*(INFO DEBUG)).*</code>. It indicates that logs with the severity level of INFO or DEBUG are not collected. ▪ Set the condition to <code>Key:url Regex:.*(?!.*(healthcheck)).*</code>. It indicates that logs whose URL contains the keyword healthcheck are not collected. For example, logs in which the key is url and the value is <code>/inner/healthcheck/jiankong.html</code> are not collected. |

8. Configure an index.

Configure an index based on your business requirements. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

After you complete the settings, you can start to collect JSON logs.

23.3.1.4.6. Collect NGINX logs

This topic describes how to collect NGINX logs and configure indexes. You can connect Log Service to NGINX and specify the required settings in the Log Service console.

Context

The NGINX log format and path are specified in the `/etc/nginx/nginx.conf` configuration file.

NGINX log format

In the configuration file, the format of NGINX logs is defined as follows:

```
log_format main '$remote_addr - $remote_user [$time_local] "$request" '
                '$request_time $request_length '
                '$status $body_bytes_sent "$http_referer" '
                '"$http_user_agent";
```

The path of the log file is declared as follows. The "main" portion that follows the path indicates that logs are written in the preceding format.

```
access_log /var/logs/nginx/access.log main
```

Sample log

A sample NGINX log is as follows:

```
192.168.1.2 - - [10/Jul/2015:15:51:09 +0800] "GET /ubuntu.iso HTTP/1.0" 0.000 129 404 168 "-" "Wget/1.11.4 Red Hat modified"
```

NGINX log fields

| Field | Description |
|---------------------|---|
| remote_addr | The IP address of the client. |
| remote_user | The username of the client. |
| request | The URL and HTTP protocol of the request. |
| status | The status of the request. |
| body_bytes_sent | The number of bytes in the response that is returned to the client, excluding the size of the response header. |
| connection | The serial number of a connection. |
| connection_requests | The number of requests that are received from a connection. |
| msec | The time when the log is written. The time is measured in seconds, accurate to milliseconds. |
| pipe | Indicates whether the request is pipelined. If the request is pipelined, the field value is <code>p</code> . Otherwise, the field value is <code>.</code> . |

| Field | Description |
|-------------------|--|
| http_referer | The URL of the web page linked to the resource that is being requested. |
| "http_user_agent" | The browser information of the client. The information must be enclosed by double quotation marks (""). |
| request_length | The length of the request. The length includes the request line, request header, and request body. |
| request_time | The time period for which the request is processed. The time period is measured in seconds, accurate to milliseconds. The time period starts when the first byte is read from the client and ends when the log is written after the last byte is sent to the client. |
| [%time_local] | The local time in the Common Log Format. The time must be enclosed by brackets []. |

Procedure

1. [Log on to the Log Service console.](#)

2. Select a data source.

Select **Nginx-Text Log**.

3. Select a Logstore, and then click Next.

Select an existing project and Logstore. You can also click **Create Now** to create a project and Logstore. For more information, see [Manage a Logstore](#).

If you enter the configuration procedure by clicking the plus sign (+) next to **Data Import** under a Logstore on the Logstores tab, the system skips this step.

4. Create a server group.

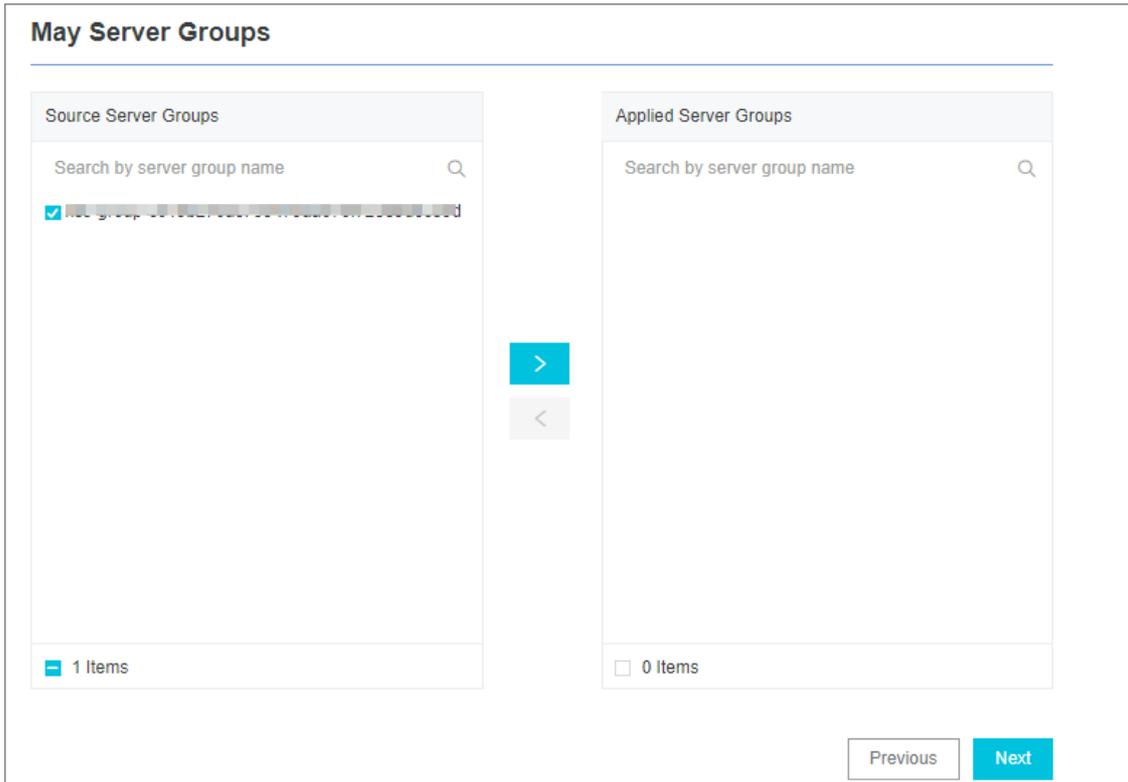
Before you create a server group, ensure that Logtail is installed.

Install Logtail as prompted. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).

After you install Logtail, click **Complete Installation** to create a server group. For more information, see [Overview](#). If you have created a server group, click **Use Existing Server Groups** to select the server group.

5. Configure the server group, and then click Next.

Select a server group and move the group from **Source Server Groups** to **Applied Server Groups**.



6. Create a Logtail configuration.

The following table lists the Logtail parameters.

| Parameter | Description |
|-------------|---|
| Config Name | <p>The configuration name must be 3 to 128 characters in length, and can contain lowercase letters, digits, hyphens (-), and underscores (_). It must start and end with a lowercase letter or digit.</p> <p>Note The configuration name cannot be modified after it is created.</p> |

| Parameter | Description |
|------------------------------------|---|
| Log Path | <p>The directory and name of the log file.</p> <ul style="list-style-type: none"> The specified log file name can be a complete file name or a file name that contains wildcards. Recursive directory matching is adopted in the log file search. If this matching method is applied, all files that match the specified file name in the specified directory and its sub-directories are monitored. <ul style="list-style-type: none"> Example 1: <code>/apsara/nuwa/.../*.log</code> indicates the files whose extension is <code>.log</code> in the <code>/apsara/nuwa</code> directory and its sub-directories are monitored. Example 2: <code>/var/logs/app_*/.../*.log*</code> indicates the files whose file name contains <code>.log</code> in the following directories are monitored: the sub-directories of the <code>/var/logs</code> directory that match the <code>app_*</code> format and the sub-directories of these matching sub-directories. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 10px; margin-top: 10px;"> <p> Note</p> <ul style="list-style-type: none"> Each log file can be collected by using only one Logtail configuration. Only the asterisk (<code>*</code>) and question mark (<code>?</code>) can be used as wildcards in the log path. </div> |
| Docker File | If the log file to be collected is in a Docker container, you can configure the internal path and container tag. Logtail monitors the creation and destruction of the container, filters logs of the container based on the tag, and collects the filtered logs. |
| Mode | If you have specified Nginx-Text Log for the data source, the default mode is NGINX Configuration Mode . You can change the mode. |
| NGINX Log Configuration | Enter the log configuration section that is specified in a standard NGINX configuration file. The section starts with <code>log_format</code> . |
| NGINX Key | Log Service reads the keys of NGINX logs. |
| Drop Failed to Parse Logs | <p>Specifies whether to upload logs to Log Service if the logs fail to be parsed.</p> <ul style="list-style-type: none"> If you turn on this switch, logs that fail to be parsed are not uploaded to Log Service. If you turn off this switch, raw logs are uploaded to Log Service when logs fail to be parsed. |
| Maximum Directory Monitoring Depth | The maximum number of directory layers that can be recursively monitored when logs are collected from the data source. Valid values: 0 to 1000. The value 0 indicates that only the directory specified in the log path is monitored. |

7. (Optional)Specify **Advanced Options** and click **Next**.

Specify **Advanced Options** based on your business requirements. We recommend that you do not modify the default settings unless otherwise required.

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|---------------------------|--|
| Enable Plug-in Processing | Specifies whether to enable plug-in processing. If you turn on this switch, you can use the plug-in of Logtail to process text logs. |
| Upload Raw Log | Specifies whether to upload raw logs. If you turn on this switch, raw logs are written to the <code>__raw__</code> field and uploaded with the parsed logs. |
| Topic Generation Mode | <ul style="list-style-type: none"> ◦ Null - Do not generate topic: This mode is selected by default. In this mode, the topic is set to an empty string and you can query logs without the need to enter a topic. ◦ Server Group Topic Attributes: This mode is used to differentiate log data that is generated by different frontend servers. ◦ File Path RegEx: If you select this mode, you must enter a value in the Custom RegEx field to extract part of the path as the topic. This mode is used to differentiate log data that is generated by users or instances. |
| Custom RegEx | Specifies a custom regular expression. If you select File Path RegEx for Topic Generation Mode , you must enter a custom regular expression. |
| Log File Encoding | <ul style="list-style-type: none"> ◦ <code>utf8</code>: indicates UTF-8 encoding. ◦ <code>gbk</code>: indicates GBK encoding. |
| Timezone | <p>Specifies the time zone where logs are collected.</p> <ul style="list-style-type: none"> ◦ System Timezone: This option is selected by default. It indicates that the time zone where logs are collected is the same as the time zone to which the server belongs. ◦ Custom: Select a time zone. |
| Timeout | <p>If a log file is not updated within a specific period of time, Logtail considers the file to be timed out.</p> <ul style="list-style-type: none"> ◦ Never: All log files are continuously monitored and never time out. ◦ 30 Minute Timeout: If a log file is not updated within 30 minutes, Logtail considers the log file to be timed out and no longer monitors the file. |
| Filter Configuration | <p>Only logs that meet all filter conditions are collected.</p> <p>Examples:</p> <ul style="list-style-type: none"> ◦ Collect logs that meet a condition: Set the condition to <code>Key:level Regex:WARNING ERROR</code>. It indicates that only logs with the severity level of WARNING or ERROR are collected. ◦ Filter logs that do not meet a condition: <ul style="list-style-type: none"> ▪ Set the condition to <code>Key:level Regex:^(?!.*(INFO DEBUG)).*</code>. It indicates that logs with the severity level of INFO or DEBUG are not collected. ▪ Set the condition to <code>Key:url Regex:.*(?!.*(healthcheck)).*</code>. It indicates that logs whose URL contains the keyword healthcheck are not collected. For example, logs in which the key is url and the value is <code>/inner/healthcheck/jiankong.html</code> are not collected. |

8. Configure an index.

Configure an index based on your business requirements. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

After you complete the settings, you can start to collect NGINX logs.

23.3.1.4.7. Collect IIS logs

This topic describes how to collect Internet Information Services (IIS) logs and configure indexes. You can specify the required settings in the Log Service console.

Context

To meet log analysis requirements, we recommend that you use the W3C Extended Log File Format. To use this format, click **Select Fields** in the IIS Manager, and then select `sc-bytes` and `cs-bytes` in the Standard Fields list.

Log format

The W3C Extended Log File Format is as follows:

```
logExtFileFlags="Date, Time, ClientIP, UserName, SiteName, ComputerName, ServerIP, Method, UriStem, UriQuery, Http
Status, Win32Status, BytesSent, BytesRecv, TimeTaken, ServerPort, UserAgent, Cookie, Referer, ProtocolVersion, Host,
HttpSubStatus"
```

- Field prefixes

| Prefix | Description |
|--------|------------------------------|
| s- | The server action. |
| c- | The client action. |
| cs- | The client-to-server action. |
| sc- | The server-to-client action. |

- Fields

| Field | Description |
|----------------|---|
| date | The date on which the client sends the request. |
| time | The time when the client sends the request. |
| s-sitename | The Internet service name and instance number of the site visited by the client. |
| s-computername | The name of the server on which the log is generated. |
| s-ip | The IP address of the server on which the log is generated. |
| cs-method | The HTTP request method that is used by the client, for example, GET or POST. |
| cs-uri-stem | The URI resource requested by the client. |
| cs-uri-query | The query string that follows the question mark (?) in the HTTP request. |
| s-port | The port number of the server to which the client is connected. |
| cs-username | The username used by the client to access the server. Authenticated users are referenced as <code>domain\username</code> . Anonymous users are indicated by a hyphen (-). |
| c-ip | The IP address of the client that sends the request. |

| Field | Description |
|-----------------|--|
| cs-version | The protocol version that is used by the client, for example, HTTP 1.0 or HTTP 1.1. |
| user-agent | The browser that is used by the client. |
| Cookie | The content of the sent or received cookie. A hyphen (-) is used if no cookie is sent or received. |
| referer | The site that the client last visited. This site provides a link to the current site. |
| cs-host | The header name of the host. |
| sc-status | The HTTP or FTP status code that is returned by the server. |
| sc-substatus | The HTTP substatus code that is returned by the server. |
| sc-win32-status | The Windows status code that is returned by the server. |
| sc-bytes | The number of bytes that are sent by the server. |
| cs-bytes | The number of bytes that are received by the server. |
| time-taken | The processing time of the request. Unit: milliseconds. |

Procedure

1. [Log on to the Log Service console.](#)

2. Select a data source.

Select **IIS-Text Log**.

3. Select a Logstore, and then click Next.

Select an existing project and Logstore. You can also click **Create Now** to create a project and Logstore. For more information, see [Manage a Logstore](#).

If you enter the configuration procedure by clicking the plus sign (+) next to **Data Import** under a Logstore on the Logstores tab, the system skips this step.

4. Create a server group.

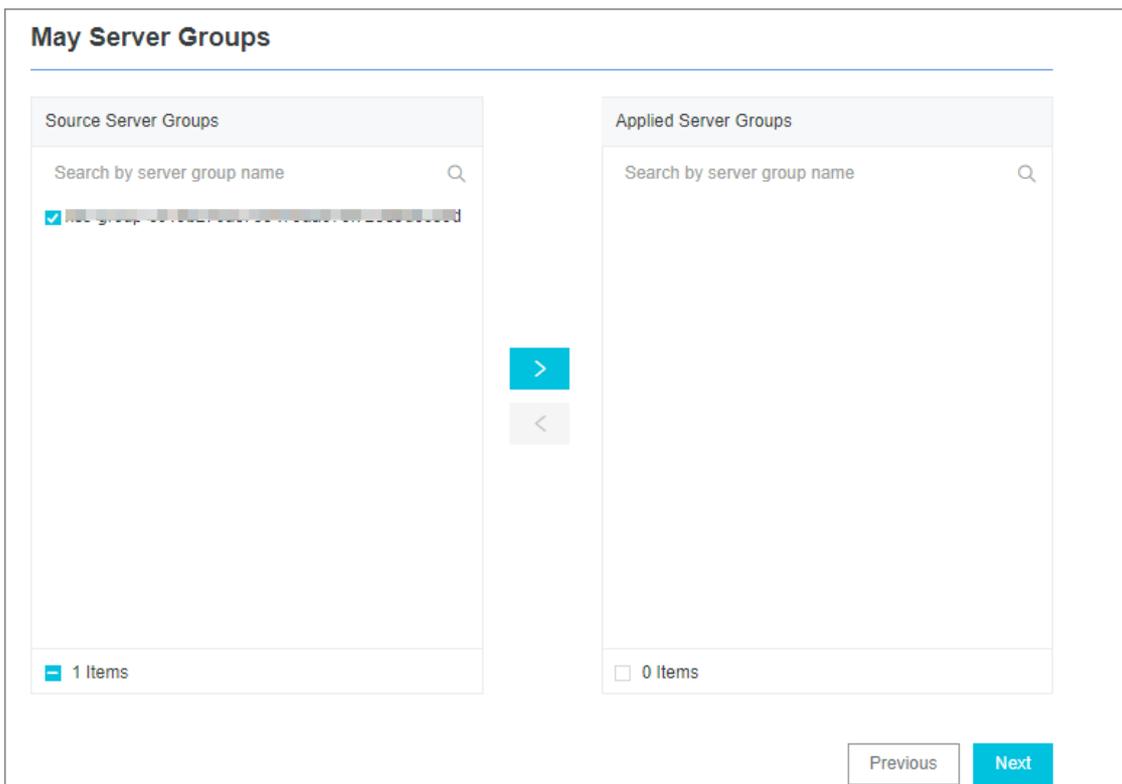
Before you create a server group, ensure that Logtail is installed.

Install Logtail as prompted. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).

After you install Logtail, click **Complete Installation** to create a server group. For more information, see [Overview](#). If you have created a server group, click **Use Existing Server Groups** to select the server group.

5. Configure the server group, and then click Next.

Select a server group and move the group from **Source Server Groups** to **Applied Server Groups**.



6. Create a Logtail configuration.

The following table lists the Logtail parameters.

| Parameter | Description |
|-------------|---|
| Config Name | <p>The configuration name must be 3 to 128 characters in length, and can contain lowercase letters, digits, hyphens (-), and underscores (_). It must start and end with a lowercase letter or digit.</p> <p>Note The configuration name cannot be modified after it is created.</p> |

| Parameter | Description |
|---------------------------|---|
| Log Path | <p>The directory and name of the log file.</p> <ul style="list-style-type: none"> ◦ The specified log file name can be a complete file name or a file name that contains wildcards. ◦ Recursive directory matching is used in the log file search. If this matching method is applied, all files that match the specified file name in the specified directory and its sub-directories are monitored. <ul style="list-style-type: none"> ▪ Example 1: <code>/apsara/nuwa/.../*.log</code> indicates the files whose extension is <code>.log</code> in the <code>/apsara/nuwa</code> directory and its sub-directories are monitored. ▪ Example 2: <code>/var/logs/app_*.../*.log*</code> indicates the files whose file name contains <code>.log</code> in the following directories are monitored: the sub-directories of the <code>/var/logs</code> directory that match the <code>app_*</code> format and the sub-directories of these matching sub-directories. <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 10px; margin-top: 10px;"> <p> Note</p> <ul style="list-style-type: none"> ▪ Each log file can be collected by using only one Logtail configuration. ▪ Only the asterisk (<code>*</code>) and question mark (<code>?</code>) can be used as wildcards in the log path. </div> |
| Docker File | <p>If the log file to be collected is in a Docker container, you can configure the internal path and container tag. Logtail monitors the creation and destruction of the container, filters logs of the container based on the tag, and collects the filtered logs. For more information, see Collect container text logs.</p> |
| Mode | <p>If you have specified IIS-Text Log for the data source, the default mode is IIS Configuration Mode. You can change the mode.</p> |
| Log Format | <p>Select the log format of your IIS server logs. Valid values:</p> <ul style="list-style-type: none"> ◦ IIS: Microsoft IIS log file format ◦ NCSA: NCSA Common log file format ◦ W3C: W3C Extended Log File Format |
| IIS Configuration | <p>Enter the log configuration section that is specified in an IIS configuration file.</p> <ul style="list-style-type: none"> ◦ If you select IIS or NCSA, the fields of the IIS log format are preconfigured. ◦ If you select W3C, enter the content that is specified for the <code>logFile logExtFileFlags</code> in the configuration file. For more information, see Specify the IIS Configuration field. |
| IIS Key Name | <p>Log Service reads the keys of IIS logs.</p> |
| Drop Failed to Parse Logs | <p>Specifies whether to upload logs to Log Service if the logs fail to be parsed.</p> <ul style="list-style-type: none"> ◦ If you turn on this switch, logs that fail to be parsed are not uploaded to Log Service. ◦ If you turn off this switch, raw logs are uploaded to Log Service when logs fail to be parsed. |

| Parameter | Description |
|------------------------------------|---|
| Maximum Directory Monitoring Depth | The maximum number of directory layers that can be recursively monitored when logs are collected from the data source. Valid values: 0 to 1000. The value 0 indicates that only the directory that is specified in the log path is monitored. |

7. Specify the IIS Configuration field.

i. Open the IIS configuration file.

- Default path of the IIS5 configuration file: *C:\WINNT\system32\inetrv\MetaBase.bin*
- Default path of the IIS6 configuration file: *C:\WINDOWS\system32\inetrv\MetaBase.xml*
- Default path of the IIS7 configuration file: *C:\Windows\System32\inetrv\config\applicationHost.config*

ii. Find the `logFile logExtFileFlags` field and copy the text in the quotation marks that follow the field name.

iii. Paste the text into the quotation marks (") in the IIS Configuration field.

8. (Optional)Specify **Advanced Options** and click **Next**.

Specify **Advanced Options** based on your business requirements. We recommend that you do not modify the default settings unless otherwise required.

| Parameter | Description |
|---------------------------|---|
| Enable Plug-in Processing | Specifies whether to enable plug-in processing. If you turn on this switch, you can use the plug-in of Logtail to process text logs. |
| Upload Raw Log | Specifies whether to upload raw logs. If you turn on this switch, raw logs are written to the <code>__raw__</code> field and uploaded with the parsed logs. |
| Topic Generation Mode | <ul style="list-style-type: none"> ◦ Null - Do not generate topic: This mode is selected by default. In this mode, the topic is set to an empty string and you can query logs without the need to enter a topic. ◦ Server Group Topic Attributes: This mode is used to differentiate log data that is generated by different frontend servers. ◦ File Path RegEx: If you select this mode, you must enter a value in the Custom RegEx field to extract part of the path as the topic. This mode is used to differentiate log data that is generated by users or instances. |
| Custom RegEx | Specifies a custom regular expression. If you select File Path RegEx for Topic Generation Mode , you must enter a custom regular expression. |
| Log File Encoding | <ul style="list-style-type: none"> ◦ <code>utf8</code>: indicates UTF-8 encoding. ◦ <code>gbk</code>: indicates GBK encoding. |
| Timezone | <p>Specifies the time zone where logs are collected.</p> <ul style="list-style-type: none"> ◦ System Timezone: This option is selected by default. It indicates that the time zone where logs are collected is the same as the time zone to which the server belongs. ◦ Custom: Select a time zone. |

| Parameter | Description |
|----------------------|--|
| Timeout | <p>If a log file is not updated within a specific period of time, Logtail considers the file to be timed out.</p> <ul style="list-style-type: none"> Never: All log files are continuously monitored and never time out. 30 Minute Timeout: If a log file is not updated within 30 minutes, Logtail considers the log file to be timed out and no longer monitors the file. |
| Filter Configuration | <p>Only logs that meet all filter conditions are collected.</p> <p>Examples:</p> <ul style="list-style-type: none"> Collect logs that meet a condition: Set the condition to <code>Key:level Regex:WARNING ERROR</code>. It indicates that only logs with the severity level of WARNING or ERROR are collected. Filter logs that do not meet a condition: <ul style="list-style-type: none"> Set the condition to <code>Key:level Regex:^(?!.*(INFO DEBUG)).*</code>. It indicates that logs with the severity level of INFO or DEBUG are not collected. Set the condition to <code>Key:url Regex:.*(?!.*(healthcheck)).*</code>. It indicates that logs whose URL contains the keyword healthcheck are not collected. For example, logs in which the key is url and the value is <code>/inner/healthcheck/jiankong.html</code> are not collected. |

9. Configure an index.

Configure an index based on your business requirements. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

After you complete the settings, you can start to collect IIS logs.

23.3.1.4.8. Collect Apache logs

This topic describes how to collect Apache logs and configure indexes. You can specify the required settings in the Log Service console.

Log formats

The Apache configuration file defines two log formats: combined log format and common log format. You can also customize a log format.

- Syntax of the combined log format:

```
LogFormat "%h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-Agent}i\"" combined
```

- Syntax of the common log format:

```
LogFormat "%h %l %u %t \"%r\" %>s %b"
```

- Syntax of a custom log format:

```
LogFormat "%h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-Agent}i\" %D %f %k %p %q %R %T %l %O" customized
```

You must specify the log format, log file directory, and log file name in the Apache configuration file. For example, the following declaration in the configuration file indicates that the combined log format is used. The log file directory is `/var/log/apache2/access_log` and the log file name is `access_log`.

```
CustomLog "/var/log/apache2/access_log" combined
```

Apache log fields

| Format string | Key name | Description |
|---------------|-------------------------|---|
| %a | client_addr | The IP address of the client in the request. |
| %A | local_addr | The local private IP address. |
| %b | response_size_bytes | The size of the response. Unit: bytes. If no bytes are sent, the value is "-" . |
| %B | response_bytes | The size of the response. Unit: bytes. If no bytes are sent, the value is 0. |
| %D | request_time_msec | The time period for which the request is processed. Unit: milliseconds. |
| %h | remote_addr | The name of the remote host. |
| %H | request_protocol_supple | The request protocol. |
| %l | remote_ident | The identity information that is provided by a remote computer. |
| %m | request_method_supple | The request method. |
| %p | remote_port | The port number of the server. |
| %P | child_process | The ID of the child process. |
| %q | request_query | The query string. If it does not exist, the value is an empty string. |
| "%r" | request | The request, which includes the method name, address, and HTTP protocol. |
| %s | status | The HTTP status code for the response. |
| %>s | status | The HTTP status code for the final response. |
| %f | filename | The name of the requested file. |
| %k | keep_alive | The number of keep-alive requests. |
| %R | response_handler | The type of the handler that generates the response on the server. |
| %t | time_local | The local time when the server receives the request. |
| %T | request_time_sec | The time period for which the request is processed. Unit: seconds. |
| %u | remote_user | The username that you used to log on to the client. |
| %U | request_uri_supple | The requested URL, excluding query strings. |
| %v | server_name | The name of the server. |

| Format string | Key name | Description |
|------------------|-----------------------|---|
| %V | server_name_canonical | The server name based on the UseCanonicalName setting. |
| %I | bytes_received | The number of bytes that are received by the server. To use this field, you must enable the mod_logio module. |
| %O | bytes_sent | The number of bytes that are sent by the server. To use this field, you must enable the mod_logio module. |
| "%{User-Agent}i" | http_user_agent | The information about the client. |
| "%{Referer}i" | http_referer | The URL of the web page linked to the resource that is being requested. |

Sample log

```
192.168.1.2 - - [02/Feb/2016:17:44:13 +0800] "GET /favicon.ico HTTP/1.1" 404 209 "http://localhost/x1.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11_3) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/48.0.2564.97 Safari/537.36"
```

Procedure

1. [Log on to the Log Service console.](#)

2. Select a data source.

Select **Apache-Text Log**.

3. Select a Logstore, and then click Next.

Select an existing project and Logstore. You can also click **Create Now** to create a project and Logstore. For more information, see [Manage a Logstore](#).

If you enter the configuration procedure by clicking the plus sign (+) next to **Data Import** under a Logstore on the Logstores tab, the system skips this step.

4. Create a server group.

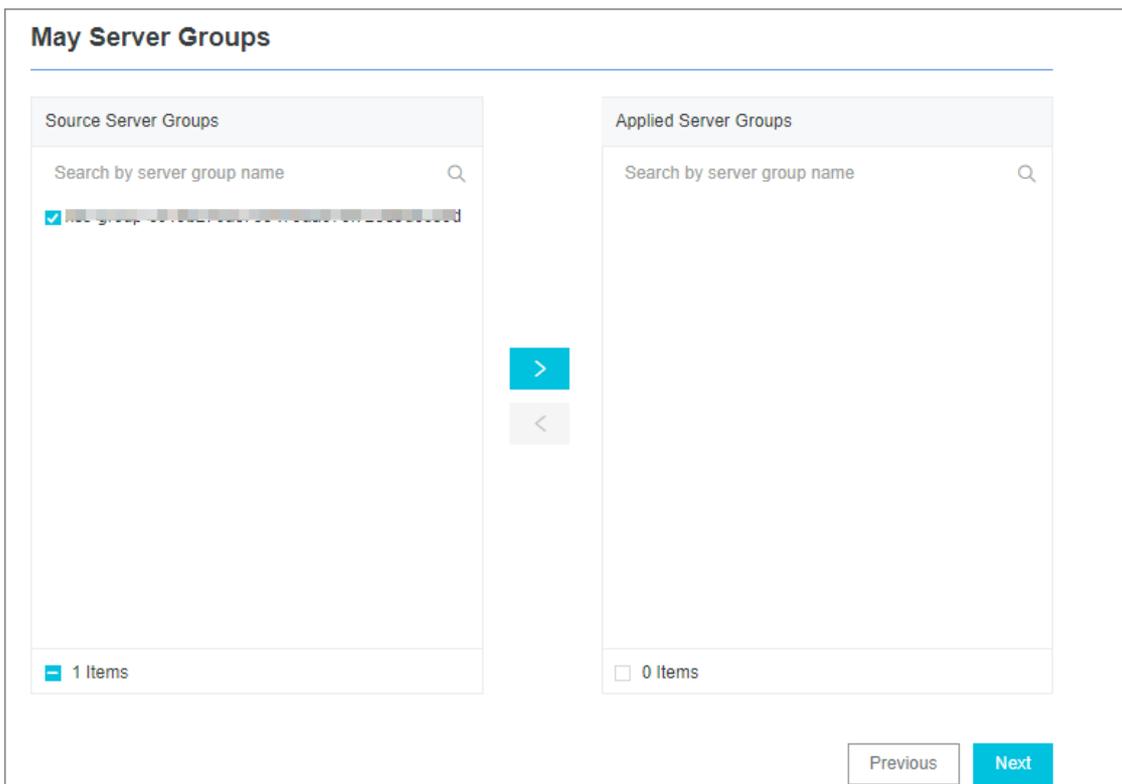
Before you create a server group, ensure that Logtail is installed.

Install Logtail as prompted. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).

After you install Logtail, click **Complete Installation** to create a server group. For more information, see [Overview](#). If you have created a server group, click **Use Existing Server Groups** to select the server group.

5. Configure the server group, and then click Next.

Select a server group and move the group from **Source Server Groups** to **Applied Server Groups**.



6. Create a Logtail configuration.

The following table lists the Logtail parameters.

| Parameter | Description |
|-------------|---|
| Config Name | <p>The configuration name must be 3 to 128 characters in length, and can contain lowercase letters, digits, hyphens (-), and underscores (_). It must start and end with a lowercase letter or digit.</p> <p>Note The configuration name cannot be modified after it is created.</p> |

| Parameter | Description |
|--------------------------------|---|
| Log Path | <p>The directory and name of the log file.</p> <ul style="list-style-type: none"> ◦ The specified log file name can be a complete file name or a file name that contains wildcards. ◦ Recursive directory matching is used in the log file search. If this matching method is applied, all files that match the specified file name in the specified directory and its sub-directories are monitored. <ul style="list-style-type: none"> ▪ Example 1: <code>/apsara/nuwa/.../*.log</code> indicates the files whose extension is <code>.log</code> in the <code>/apsara/nuwa</code> directory and its sub-directories are monitored. ▪ Example 2: <code>/var/logs/app_*.../*.log*</code> indicates the files whose file name contains <code>.log</code> in the following directories are monitored: the sub-directories of the <code>/var/logs</code> directory that match the <code>app_*</code> format and the sub-directories of these matching sub-directories. <div style="background-color: #e6f2ff; padding: 10px; margin-top: 10px;"> <p> Note</p> <ul style="list-style-type: none"> ▪ Each log file can be collected by using only one Logtail configuration. ▪ Only the asterisk (<code>*</code>) and question mark (<code>?</code>) can be used as wildcards in the log path. </div> |
| Docker File | <p>If the log file to be collected is in a Docker container, you can configure the internal path and container tag. Logtail monitors the creation and destruction of the container, filters logs of the container based on the tag, and collects the filtered logs. For more information, see Collect container text logs.</p> |
| Mode | <p>If you have specified Apache-Text Log for the data source, the default mode is Apache Configuration Mode. You can change the mode.</p> |
| Log Format | <p>Select a log format based on the format declared in your Apache log configuration file. To facilitate the query and analysis of log data, we recommend that you use a custom Apache log format.</p> |
| APACHE Logformat Configuration | <p>Enter the log configuration section that is specified in the Apache configuration file. The section starts with LogFormat.</p> <div style="background-color: #e6f2ff; padding: 10px; margin-top: 10px;"> <p> Note If the specified Log Format is Common or Combined, the system enters a commonly used syntax of the log format. Check whether the log format is the same as that defined in the Apache configuration file.</p> </div> |
| APACHE Key Name | <p>Log Service reads the keys of Apache logs. Confirm the key names on the Logtail configuration page.</p> |
| Drop Failed to Parse Logs | <p>Specifies whether to upload logs to Log Service if the logs fail to be parsed.</p> <ul style="list-style-type: none"> ◦ If you turn on this switch, logs that fail to be parsed are not uploaded to Log Service. ◦ If you turn off this switch, raw logs are uploaded to Log Service when logs fail to be parsed. |

| Parameter | Description |
|------------------------------------|---|
| Maximum Directory Monitoring Depth | The maximum number of directory layers that can be recursively monitored when logs are collected from the data source. Valid values: 0 to 1000. The value 0 indicates that only the directory specified in the log path is monitored. |

7. (Optional)Specify **Advanced Options** and click **Next**.

Specify **Advanced Options** based on your business requirements. We recommend that you do not modify the default settings unless otherwise required.

| Parameter | Description |
|---------------------------|---|
| Enable Plug-in Processing | Specifies whether to enable plug-in processing. If you turn on this switch, you can use the plug-in of Logtail to process text logs. |
| Upload Raw Log | Specifies whether to upload raw logs. If you turn on this switch, raw logs are written to the <code>__raw__</code> field and uploaded with the parsed logs. |
| Topic Generation Mode | <ul style="list-style-type: none"> ◦ Null - Do not generate topic: This mode is selected by default. In this mode, the topic is set to an empty string and you can query logs without the need to enter a topic. ◦ Server Group Topic Attributes: This mode is used to differentiate log data that is generated by different frontend servers. ◦ File Path RegEx: If you select this mode, you must enter a value in the Custom RegEx field to extract part of the path as the topic. This mode is used to differentiate log data that is generated by users or instances. |
| Custom RegEx | Specifies a custom regular expression. If you select File Path RegEx for Topic Generation Mode , you must enter a custom regular expression. |
| Log File Encoding | <ul style="list-style-type: none"> ◦ <code>utf8</code>: indicates UTF-8 encoding. ◦ <code>gbk</code>: indicates GBK encoding. |
| Timezone | <p>Specifies the time zone where logs are collected.</p> <ul style="list-style-type: none"> ◦ System Timezone: This option is selected by default. It indicates that the time zone where logs are collected is the same as the time zone to which the server belongs. ◦ Custom: Select a time zone. |
| Timeout | <p>If a log file is not updated within a specific period of time, Logtail considers the file to be timed out.</p> <ul style="list-style-type: none"> ◦ Never: All log files are continuously monitored and never time out. ◦ 30 Minute Timeout: If a log file is not updated within 30 minutes, Logtail considers the log file to be timed out and no longer monitors the file. |

| Parameter | Description |
|----------------------|---|
| Filter Configuration | <p>Only logs that meet all filter conditions are collected.</p> <p>Examples:</p> <ul style="list-style-type: none"> ○ Collect logs that meet a condition: Set the condition to <code>Key:level Regex:WARNING ERROR</code>. It indicates that only logs with the severity level of WARNING or ERROR are collected. ○ Filter logs that do not meet a condition: <ul style="list-style-type: none"> ■ Set the condition to <code>Key:level Regex:^(?!.*(INFO DEBUG)).*</code>. It indicates that logs with the severity level of INFO or DEBUG are not collected. ■ Set the condition to <code>Key:url Regex:^(?!.*(healthcheck)).*</code>. It indicates that logs whose URL contains the keyword healthcheck are not collected. For example, logs in which the key is url and the value is <code>/inner/healthcheck/jiankong.html</code> are not collected. |

8. Configure an index.

Configure an index based on your business requirements. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

After you complete the settings, you can start to collect Apache logs.

23.3.1.4.9. Configure parsing scripts

This topic describes how to configure log contents for log collection.

Specify a method to separate log lines

A complete access log such as an NGINX access log occupies a line. Separate multiple log entries with line breaks. For example, the following shows two access logs:

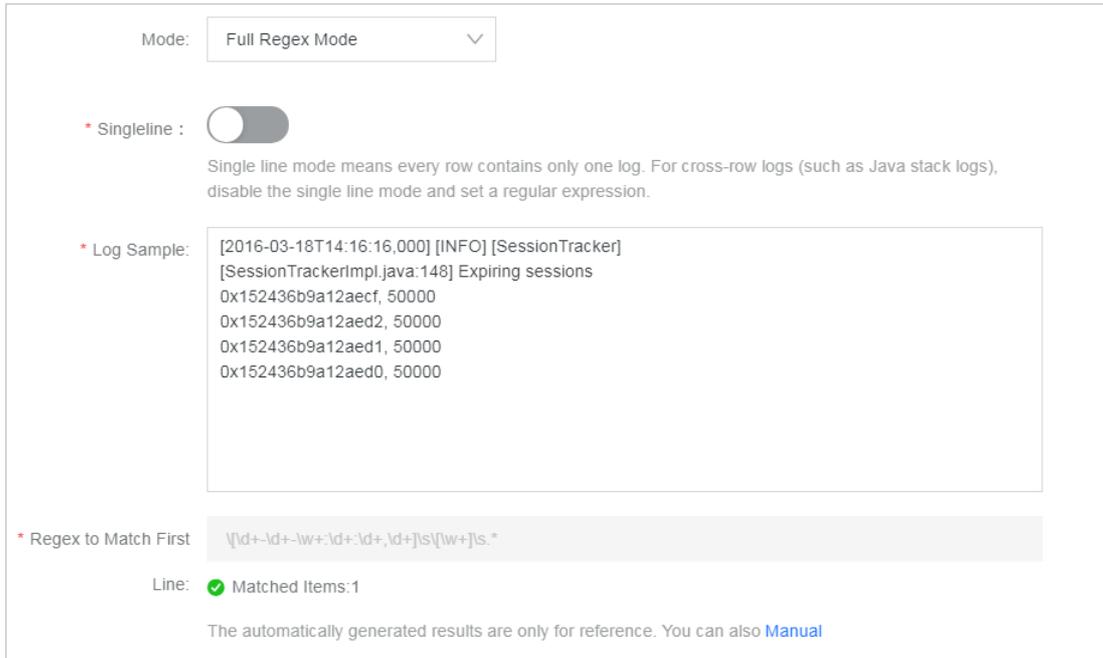
```
10.1.1.1 - - [13/Mar/2016:10:00:10 +0800] "GET / HTTP/1.1" 0.011 180 404 570 "-" "Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; 360se)"
10.1.1.1 - - [13/Mar/2016:10:00:11 +0800] "GET / HTTP/1.1" 0.011 180 404 570 "-" "Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; 360se)"
```

For Java applications, a log entry usually spans several lines. Therefore, log entries are separated based on the identifier at the beginning of each log entry. The following example shows a Java application log.

```
[2016-03-18T14:16:16,000] [INFO] [SessionTracker] [SessionTrackerImpl.java:148] Expiring sessions
0x152436b9a12aecf, 50000
0x152436b9a12aed2, 50000
0x152436b9a12aed1, 50000
0x152436b9a12aed0, 50000
```

The preceding Java application log entries each start with a time field. The regular expression that matches these time fields is `[\d+-\d+-\d+:\d+:\d+]\s.*`. You can enter information in the Log Service console as shown in the following figure.

Full regular expression mode



Extract log fields

To conform to the data models of Log Service, a log contains one or more key-value pairs. If you want to extract specific fields for analysis, you must set a regular expression. If you do not want to process the contents of a log, you can treat the log as a key-value pair.

You can determine whether to extract fields from the preceding NGINX access log.

- Extract fields

The regular expression is `(\S+)\s-\s-\s{[(\S+)\s{^}]+\s"}(\w+).*`. The extracted fields are `10.1.1.1`, `13/Mar/2016:10:00`, and `GET`.

- Extract all

The regular expression is `(.*)`. The extracted field is `10.1.1.1 -- [13/Mar/2016:10:00:10+0800] "GET / HTTP/1.1" 0.011 180 404 570 "-" "Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; 360se)"`.

Specify a log time

To conform to the data models of Log Service, a log must have a time field in the Unix timestamp format. You can use the system time when Logtail collects a log or the time in the log contents as the log time.

For the preceding NGINX access log:

- If you extract the time field from the log contents as the log time, the time is `13/Mar/2016:10:00:10` and the time expression is `%d/%b/%Y:%H:%M:%S`.
- If you use the system time when the log was collected as the log time, the log time is converted into a timestamp.

23.3.1.4.10. Configure the time format

Each log in Log Service must have a timestamp that records the log generation time. When you collect logs from log files, Logtail must extract the timestamp string of each log and parse it into a timestamp. Therefore, you need to specify a timestamp format to facilitate parsing.

In Linux, Logtail supports all time formats provided by the `strptime()` function. If a timestamp string can match one of the time formats that are provided by the `strptime()` function, Logtail can parse and use the timestamp string.

Note

- The log timestamp is accurate to seconds. Therefore, you need to specify seconds in a time format, without the need for other information such as milliseconds or microseconds.
- In addition, you need to configure the time field rather than other information.

Common log time formats supported by Logtail

The timestamp strings of logs have diverse formats. To make configuration easier, the following table lists the common log time formats supported by Logtail:

| Format | Description | Example |
|--------|---|---------------|
| %a | The abbreviation of a day in a week. | Fri |
| %A | The day in a week. | Friday |
| %b | The abbreviation of a month. | Jan |
| %B | The month name. | January |
| %d | The numerical day in a month. Valid values: 01 to 31. | 07 and 31 |
| %h | The abbreviation of a month. The format is equivalent to %b . | Jan |
| %H | The hour in the 24-hour format. | 22 |
| %I | The hour in the 12-hour format. | 11 |
| %m | The numerical month. | 08 |
| %M | The numerical minute. Valid values: 00 to 59. | 59 |
| %n | A line break. | Line break |
| %p | The local time in the a.m. or p.m. format. | AM and PM |
| %r | The time in the 12-hour format. The format is equivalent to %I:%M:%S %p . | 11:59:59 AM |
| %R | The time includes hours and minutes. The format %R is equivalent to %H:%M . | 23:59 |
| %S | The numerical second. Valid values: 00 to 59. | 59 |
| %t | A tab. | Tab |
| %y | The two-digit numerical year. Valid values: 00 to 99. | 04 and 98 |
| %Y | The four-digit numerical year. | 2004 and 1998 |

| Format | Description | Example |
|--------|---|--|
| %C | The numerical century. Valid values: 00 to 99. | 16 |
| %e | The numerical day in a month. Valid values: 1 to 31. A single digit is preceded by a space. | 7 and 31 |
| %j | The numerical day in a year. Valid values: 001 to 366. | 365 |
| %u | The numerical day in a week. Valid values: 1 to 7, in which 1 represents Monday. | 2 |
| %U | The numerical week in a year. Sunday is the first day of a week. Valid values: 00 to 53. | 23 |
| %V | The numerical week in a year. Monday is the first day of a week. If a week has four or more days that start from January 1, the week is treated as the first week. Otherwise, the next week is treated as the first week. Valid values: 01 to 53. | 24 |
| %w | The numerical day in a week. Valid values: 0 to 6, in which 0 represents Sunday. | 5 |
| %W | The numerical week in a year. Monday is the first day of a week. Valid values: 00 to 53. | 23 |
| %c | The standard date and time. | To specify more information such as the long date and short date, you can use the preceding formats to provide exact expression. |
| %x | The standard date. | To specify more information such as the long date and short date, you can use the preceding formats to provide exact time expressions. |
| %X | The standard time. | To specify more information such as the long date and short date, you can use the preceding formats to provide exact expression. |
| %s | The Unix timestamp. | 1476187251 |

Example

The following table lists the common log time formats, examples, and corresponding time expressions.

| Log time format | Example | Time expression |
|-----------------|---------------------|-------------------|
| Custom | 2017-12-11 15:05:07 | %Y-%m-%d %H:%M:%S |

| Log time format | Example | Time expression |
|-----------------|-------------------------------------|-----------------------|
| Custom | [2017-12-11 15:05:07.012] | [%Y-%m-%d %H:%M:%S] |
| RFC822 | 02 Jan 06 15:04 MST | %d %b %y %H:%M |
| RFC822Z | 02 Jan 06 15:04 -0700 | %d %b %y %H:%M |
| RFC850 | Monday, 02-Jan-06 15:04:05 MST | %A, %d-%b-%y %H:%M:%S |
| RFC1123 | Mon, 02 Jan 2006 15:04:05 MST | %A, %d-%b-%y %H:%M:%S |
| RFC3339 | 2006-01-02T15:04:05Z07:00 | %Y-%m-%dT%H:%M:%S |
| RFC3339Nano | 2006-01-02T15:04:05.999999999Z07:00 | %Y-%m-%dT%H:%M:%S |

23.3.1.4.11. Import historical logs

Logtail collects incremental logs by default. If you want to import historical logs, use the historical log importing feature of Logtail.

Prerequisites

To collect logs from Linux servers, use Logtail 0.16.15 or later. To collect logs from Windows servers, use Logtail 1.0.0.1 or later. To ensure successful log collection, update Logtail to the latest version.

Context

Logtail collects log files based on events. The system captures events by detecting or polling files for changes at intervals. Additionally, Logtail can load events from local files to trigger log collection. Logtail implements historical log collection based on these local events.

You can import historical log files from the Logtail installation directory. The location of the directory varies based on the operating system.

- Linux: `/usr/local/ilogtail`
- Windows:
 - 32-bit: `C:\Program Files\Alibaba\Logtail`
 - 64-bit: `C:\Program Files (x86)\Alibaba\Logtail`

Note

- The maximum interval between the time a local event is generated and the time the local event is imported is one minute.
- Loading local configurations is a special action. Therefore, Logtail sends the `LOAD_LOCAL_EVENT_ALARM` alert to your server to notify you of this action.
- If you want to import a large number of log files, we recommend that you modify the Logtail startup configuration to increase the upper limit of CPU to 2.0 GHz or more and the upper limit of the memory size to 512 MB or more. For more information, see [Set Logtail startup parameters](#).

Procedure

1. Configure log collection.

If a collection configuration is only used to import historical log files, you can specify a collection directory that does not exist. For more information, see [Configure text log collection](#).

2. Obtain a unique identifier for a collection configuration.

Obtain the unique identifier in the *user_log_config.json* file stored in the installation directory of Logtail. In Linux, use the **grep** command in the directory to query the unique identifier. In Windows, use tools such as Notepad to query the unique identifier.

To query a unique identifier in a Linux operating system, run the following command:

```
grep "##" /usr/local/ilogtail/user_log_config.json | awk '{print $1}'
##1.0##log-config-test$multi"
##1.0##log-config-test$secs-test"
##1.0##log-config-test$metric_system_test"
##1.0##log-config-test$redis-status"
```

3. Add local events.

Local events are stored in the *local_event.json* file that resides in the installation directory of Logtail. The file is in the JSON format. The syntax is:

```
[
  {
    "config": "${your_config_unique_id}",
    "dir": "${your_log_dir}",
    "name": "${your_log_file_name}"
  },
  {
    ...
  }
  ...
]
```

o Parameters

| Parameter | Description | Example |
|-----------|--|--|
| config | The unique identifier that is obtained in Step 2. | ##1.0##log-config-test\$secs-test |
| dir | The directory where logs are stored. Note The directory cannot end with a slash (/). | /data/logs |
| name | The name of a log file. Wildcards are supported. | For example, access.log.2018-08-08 and access.log* |

Note To prevent Logtail from loading invalid JSON files, save local event configurations to a temporary file, edit the configurations in the temporary file, and copy the contents to the *local_event.json* file.

o Configuration examples

In a Windows system, you can use tools such as Notepad to add local events to the *local_event.json* file. In a Linux system, add local events as follows:

```
$ cat /usr/local/ilogtail/local_event.json
[
  {
    "config": "##1.0##log-config-test$secs-test",
    "dir": "/data/log/",
    "name": "access.log",
  },
  {
    "config": "##1.0##log-config-test$secs-test",
    "dir": "/tmp",
    "name": "access.log.2017-08-09"
  }
]
```

What's next

- Check whether Logtail has loaded configurations

After you save the `local_event.json` file, Logtail loads the configuration file to the memory within one minute and clears the contents of the `local_event.json` file.

To check whether Logtail has read local events, use the following methods:

- If the contents of the `local_event.json` file are cleared, it indicates that Logtail has read the local events.
 - Check whether the `ilogtail.LOG` file in the Logtail installation directory contains the `process local event` keywords. If the contents of the `local_event.json` are cleared and these keywords cannot be found, the local configuration file may be screened due to invalid contents.
- Check whether the configuration is loaded but no data is collected

Possible causes are as follows:

- The configuration is invalid.
- The local `config` file does not exist.
- The log file does not exist in the path specified in the collection configuration.
- The log file has been collected by Logtail.

23.3.1.4.12. Generate a topic

This topic describes how to generate a topic in the Log Service console. After you generate a topic, you can use the topic to group logs. You can specify topics for logs when these logs are written. You can use a topic as a filter when you query logs.

Topic generation modes

You can set a topic when you use Logtail to collect logs or when you use API operations or SDKs to upload logs. The following topic generation modes are available in the Log Service console: **Null - Do not generate topic**, **Server Group Topic Attributes**, and **File Path RegEx**.

- **Null - Do not generate topic**

When you configure Logtail in the Log Service console to collect text logs, the default topic generation mode is **Null - Do not generate topic**. In this mode, no topic is generated and query logs without specifying a topic.

- **Server Group Topic Attributes**

You can use this mode to identify logs that are generated from multiple servers. Logs from multiple servers can be stored in the same file or directory. To identify these logs based on topics during log collection, you can create server groups and add the servers into different groups. When you create server groups, you must specify a unique **topic attribute** for each server group and set **Topic Generation Mode** to **Server Group Topic Attributes**. After you complete the configuration, apply the Logtail settings to the server groups.

If the **Server Group Topic Attributes** mode is selected, Logtail uploads the topic attribute of each server group as topics to Log Service. When you query logs, you must specify the topic of the target server group as a filter.

• **File Path RegEx**

- You can use this mode to differentiate between logs that are generated by multiple users or instances. If Log Service stores logs in different directories for different users or instances, duplicate sub-directory names or log file names may exist in these directories. As a result, Log Service cannot identify the source of logs. You can select **File Path RegEx** in the **Topic Generation Mode** field. Enter a regular expression that matches an absolute file path, and set an instance name as a topic.
- If you select **File Path RegEx**, Logtail uses an instance name as the topic of the logs that Logtail uploads to Log Service. The topic generated varies based on your directory structure and configuration. You must specify an instance name as a topic when you query logs. For example, the following directory structure includes directories that each store logs generated by different users or instances:

```

/logs
|- /userA/serviceA
|- service.log
|- /userB/serviceA
|- service.log
|- /userC/serviceA
|- service.log

```

- If you want to extract multiple separate fields from a file path, use a multi-layer extraction method of `?P<key>`. The value of the key can contain lowercase letters and digits. For example:

```

/home/admin/serviceA/userB/access.log
\home\admin\(?P<service>[^\s]+)\(?P<user>[^\s]+)\. *

```

The following custom tags are created for logs:

```

"__tag__ : service : serviceA"
"__tag__ : user : userB"

```

 **Note** Logtail 0.16.19 and later are supported.

- If you specify the `/logs` file path and the `service.log` file name in a regular expression, Logtail collects logs from the preceding directories that contain the `service.log` file and uploads the logs to Log Service. However, Log Service cannot identify the log source based on log contents. You can select **File Path RegEx** in the **Topic Generation Mode** field, and enter the `\(.*)\serviceA\.*` regular expression to extract instance names. After the configuration is complete, the following topics are generated for logs in different directories: `userA`, `userB`, and `userC`. You can specify a topic as a filter to query logs.

 **Note** You must escape the forward slashes (/) in the file path that the regular expression contains.

• **Static topic generation**

You can select **File Path RegEx** in the **Topic Generation Mode** field, and enter `customized:// + custom topic` in the **Custom RegEx** field.

 **Note** Logtail 0.16.21 and later are supported.

Set a log topic

1. Configure Logtail in the Log Service console. For more information, see [Configure text log collection](#).
To set the topic generation mode to **Server Group Topic Attributes**, go to the **Topic** section on the server group creation or modification page.
2. In the Logtail Configuration for Data Import step, click **Advanced Options** and select a **topic generation mode**.

23.3.1.5. Custom plug-ins

Context

Log Service allows you to collect text logs and system logs through Logtail. Logtail supports connections with multiple data sources, such as HTTP or MySQL query results and MySQL binary logs.

You can collect HTTP request data and upload the processing results to Log Service in real time to check service availability check and continuous availability monitoring. You can configure MySQL query results as the data source, and then synchronize incremental data based on custom IDs or time. You can also configure an SQL data source to synchronize MySQL binary logs, subscribe to database changes, and query or analyze logs in real time.

 **Note** This feature is only supported on Linux and must be used together with Logtail 0.16.0 or later versions. For more information, see [Install Logtail in Linux](#).

配置流程

1. Configure a method that is used to collect logs from the data source.
Different Logtail configurations for different data sources. Select a Logtail configuration according to your data source.
2. Configure a processing method.
Logtail provides multiple processing methods for binary logs, MySQL query results, NGINX monitoring data, and HTTP input sources. You can configure multiple processing methods for a single input source. Each input source supports all processing methods. Logtail runs the configured processing methods in sequence.
For more information, see [Configure data processing methods](#).
3. Apply the configurations to the machine group.
After you configure the collection and processing methods, apply them to the specified machine group. Then, Logtail automatically applies the configurations and starts data collection.

23.3.1.5.1. Collect MySQL binary logs

Logtail is used as a MySQL slave. It is used to collect binary logs from a MySQL master. Logtail collects binary logs by using a similar method to Alibaba Canal. This improves the efficiency of log collection.

Features

- Allows you to collect incremental data of databases in the form of binary logs to improve performance. Supports MySQL databases such as ApsaraDB RDS for MySQL.
- Supports multiple database filters, such as regular expressions.
- Allows you to set binary log file positions.
- Allows you to record synchronization statuses by using the checkpoint mechanism.

Limits

- MySQL binary logs are available only for Logtail 0.16.0 or later versions that you install on Linux. For more information about how to update Logtail and view Logtail versions, see [Install Logtail in Linux](#).

- Binary logs in the ROW format must be enabled for MySQL databases. By default, binary logs in the ROW format are enabled for RDS instances.

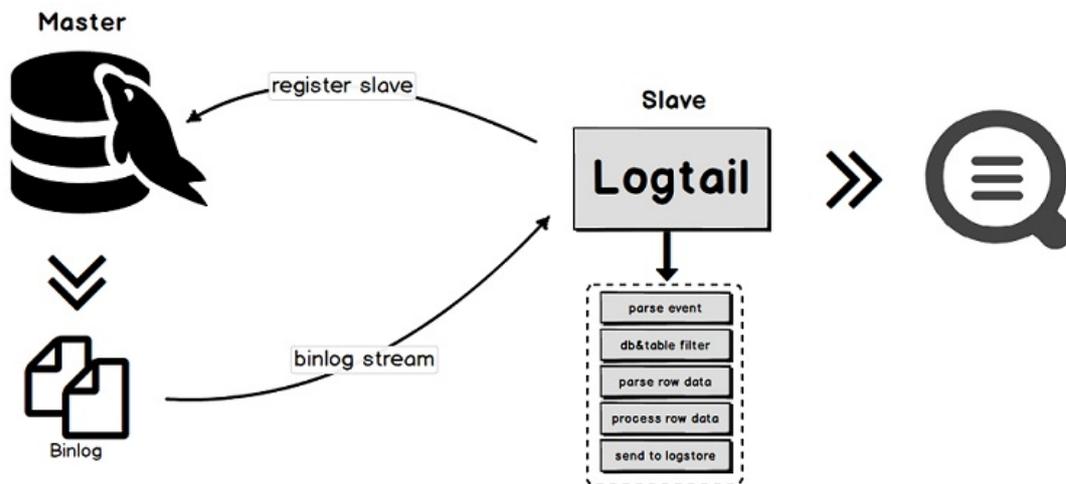
```
# Check whether binary logs are enabled.
mysql> show variables like "log_bin";
+-----+-----+
| Variable_name | Value |
+-----+-----+
| log_bin      | ON   |
+-----+-----+
1 row in set (0.02 sec)
# View the format of binary logs.
mysql> show variables like "binlog_format";
+-----+-----+
| Variable_name | Value |
+-----+-----+
| binlog_format | ROW   |
+-----+-----+
1 row in set (0.03 sec)
```

- Each server ID must be unique. Make sure that the ID of each slave to be synchronized is unique.
- Limits for RDS databases:
 - Logtail cannot be installed on an RDS instance. You must install Logtail on an ECS instance that can communicate with the destination RDS instance.
 - Secondary RDS databases cannot be used to collect binary logs. You must configure a primary RDS database to collect binary logs.

Implementation

Logtail enables communication between master and slave MySQL servers. The process of how master and slave MySQL servers communicate is provided in the following information:

1. Logtail is used as a MySQL slave. It can also be used to send dump requests to the MySQL master.
2. After the dump requests are received, the MySQL master delivers its binary logs to Logtail in real time.
3. Logtail parses and filters binary logs, and then uploads the results to Log Service.



Scenarios

The MySQL binary logging feature applies to scenarios in which you need to synchronize large amounts of data and meet high performance requirements.

- Query the incremental data of databases in real time.

- Audit operations that are performed on databases.
- Use Log Service to query data, visualize query results, transform data for stream processing, export data to MaxCompute for offline computing, and export log to Object Storage Service (OSS) for long-term storage.

Data reliability

We recommend that you enable the global transaction identifier (GTID) feature of the MySQL server and upgrade Logtail to version 0.16.15 or later. This prevents repeated data collection during a primary/secondary server switchover and ensures data reliability.

- Incomplete data collection: If the network between Logtail and the MySQL server is disconnected for a long period of time, some data may not be collected.

A MySQL binary log plug-in is used as a MySQL slave to collect binary logs from the master server. Logtail establishes a connection with the master server to obtain data from the server. If the network between Logtail and the master node is disconnected, the master node still generates new binary logs and deletes expired binary logs. After the connection is reestablished and Logtail is reconnected to the master server, Logtail uses the last checkpoint to request binary log data from the master server. However, if the network is disconnected for a long period of time, the data generated after the checkpoint may be deleted. In this case, the recovery mechanism specifies the new point at which Logtail resumes collecting binary logs. The new point is the most recent binary log file position. If the network is disconnected for a long period of time, some data generated between the checkpoint and the new data collection point may not be collected.

- Repeated data collection: If the ordinal numbers of binary logs on the master and slave servers are different and a master/slave switchover occurs, repeated data collection may occur.

If the MySQL master-slave synchronization is configured, the master server synchronizes the generated binary log data to the slave server. Then, the slave server stores the received binary log data to the local binary log file. If the ordinal numbers of binary logs on the master and slave servers are different, a master/slave switchover occurs. In this case, the mechanism that uses a binary log file name and an offset as the checkpoint causes repeated data collection.

For example, assume that a data entry ranges from (binlog.100, 4) to (binlog.105, 4) on the master server, and ranges from (binlog.1000, 4) to (binlog.1005, 4) on the slave. Logtail has obtained the data from the master server and updated the checkpoint to (binlog.105, 4). In this case, if a master/slave switchover occurs without exception, Logtail continues to obtain binary logs from the new master server based on the local checkpoint (binlog.105, 4). The new master server returns the data entries that range from (binlog.1000, 4) to (binlog.1005, 4) to Logtail. This is because the ordinal numbers of these data entries on the new master server are greater than the ordinal numbers of data entries requested by Logtail. As a result, log data is repeatedly collected.

Parameter

The type of input sources is `service_canal`.

| Parameter | Type | Required | Description |
|-----------|--------|----------|---|
| Host | string | No | The IP address of the host where the database resides. Default value: 127.0.0.1. |
| Port | int | No | The port number that you can use to connect with the database. Default value: 3306. |

| Parameter | Type | Required | Description |
|---------------|--------------|----------|---|
| User | string | No | <p>The database username. Default value: root.</p> <p>The configured user must have the read permissions on the source database and the MySQL REPLICATION permission. Example:</p> <pre>CREATE USER canal IDENTIFIED BY 'canal'; GRANT SELECT, REPLICATION SLAVE, REPLICATION CLIENT ON *.* TO 'canal'@'%'; -- GRANT ALL PRIVILEGES ON *.* TO 'canal'@'%'; FLUSH PRIVILEGES;</pre> |
| Password | string | No | <p>The database password. By default, this parameter is unspecified.</p> <p>If you require a high level of data security, we recommend that you set both the username and the password to xxx. After your configurations are synchronized to the on-premises server, find the Password parameter in the <code>/usr/local/ilogtail/user_log_config.json</code> file and modify the value.</p> <div style="background-color: #e6f2ff; padding: 10px; border: 1px solid #d9e1f2;"> <p>Note</p> <ul style="list-style-type: none"> After you modify the password, use the <code>sudo /etc/init.d/ilogtailed stop; sudo /etc/init.d/ilogtailed start</code> command to restart Logtail. If you modify the value of the Password parameter on the web and synchronize your configurations to the on-premises server, the configurations on the on-premises server are overwritten. You can change the configurations on the on-premises server later. </div> |
| ServerID | int | No | <p>The ID of a MySQL slave whose role is assumed by Logtail. Default value: 125.</p> <div style="background-color: #e6f2ff; padding: 10px; border: 1px solid #d9e1f2;"> <p>Note In a MySQL database, each ID must be unique. Otherwise, synchronization fails.</p> </div> |
| IncludeTables | String array | Yes | <p>The names of matched tables. Each value contains a database name and a table name, for example, <code>test_db.test_table</code>. You must specify a regular expression for the parameter. Logtail does not collect incremental data from tables whose names do not match the regular expression. To collect incremental data from all tables of a database, set the value of the IncludeTables parameter to <code>.*\..*</code>.</p> <div style="background-color: #e6f2ff; padding: 10px; border: 1px solid #d9e1f2;"> <p>Note If an exact match is required, make sure that the regular expression is prefixed by <code>^</code>. In this case, you must also make sure that the regular expression is suffixed by <code>\$</code>. Example: <code>^test_db\\.test_table\$</code>.</p> </div> |

| Parameter | Type | Required | Description |
|----------------|--------------|----------|---|
| ExcludeTables | String array | No | <p>The names of excluded tables expressed as a regular expression. The name of a table must include the name of the database to which the table belongs, for example, <code>test_db.test_table</code> . If a table meets one of the conditions specified in the parameter, the table is not collected. If you do not specify this parameter, incremental data from all tables is collected.</p> <p>Note If an exact match is required, make sure that the regular expression is prefixed by <code>^</code> . In this case, you must also make sure that the regular expression is suffixed by <code>\$</code> . Example: <code>^test_db\\.test_table\$</code> .</p> |
| StartBinName | string | No | <p>The name of the first binary log file that is collected by Logtail. If you do not specify this parameter, Logtail starts to collect binary log files that are generated from the current time.</p> <p>To collect data from a specific location, view the name of the current binary log file and the file offset. Then, set <code>StartBinName</code> and <code>StartBinLogPos</code> to actual values. Example:</p> <pre># Set StartBinName to mysql-bin.000063 and StartBinLogPos to 0. mysql> show binary logs; +-----+-----+ Log_name File_size +-----+-----+ mysql-bin.000063 241 mysql-bin.000064 241 mysql-bin.000065 241 mysql-bin.000066 10778 +-----+-----+ 4 rows in set (0.02 sec)</pre> <p>Note If you set the <code>StartBinName</code> parameter, a large amount of traffic is generated during the first collection.</p> |
| StartBinLogPos | int | No | The offset of the first binary log file that is collected. Default value: 0. |
| EnableGTID | bool | No | Specifies whether to add GTID . Default value: true. If the value is false, no GTID is added to uploaded data. |
| EnableInsert | bool | No | Specifies whether to collect log events triggered by INSERT operations. Default value: true. If the value is false, INSERT events are not collected. |
| EnableUpdate | bool | No | Specifies whether to collect UPDATE events. Default value: true. If the value is false, UPDATE events are not collected. |
| EnableDelete | bool | No | Specifies whether to collect DELETE events. Default value: true. If the value is false, DELETE events are not collected. |

| Parameter | Type | Required | Description |
|-----------------|--------|----------|---|
| EnableDDL | bool | No | <p>Specifies whether to collect data definition language (DDL) events. Default value: false. If the value is false, DDL events are not collected.</p> <p> Note This parameter does not support the <code>IncludeTables</code> or <code>ExcludeTables</code> filtering methods.</p> |
| Charset | string | No | The encoding method. Default value: <code>utf-8</code> . |
| TextToString | bool | No | Specifies whether to convert data of the text type into a string. Default value: false. |
| PackValues | bool | No | <p>Specifies whether to encapsulate event data into the JSON format. Default value: false. If the value is false, event data is not encapsulated. If this feature is enabled, Logtail encapsulates event data into the <code>data</code> and <code>old_data</code> fields in the JSON format. The <code>old_data</code> field is available only for <code>ROW_UPDATE</code> events.</p> <p>For example, assume that a table has three fields named <code>c1</code>, <code>c2</code>, and <code>c3</code>. If this feature is disabled, the <code>ROW_INSERT</code> event data contains three fields <code>c1</code>, <code>c2</code>, and <code>c3</code>. If this feature is enabled, <code>c1</code>, <code>c2</code>, and <code>c3</code> are encapsulated into one data field and the value is <code>{"c1": "...", "c2": "...", "c3": "..."} .</code></p> <p> Note This parameter is available only for Logtail V0.16.19 and later.</p> |
| EnableEventMeta | bool | No | <p>Specifies whether to collect event metadata. Default value: false. If the value is false, event metadata is not collected. The metadata of binary log events includes <code>event_time</code>, <code>event_log_position</code>, <code>event_size</code>, and <code>event_server_id</code>.</p> <p> Note This parameter is available only for Logtail V0.16.21 and later.</p> |

Procedure

Synchronize data from tables whose names do not end with `_inner` in the `user_info` RDS database.

1. [Log on to the Log Service console](#).

2. Select a data source.

Click **Import Data**. On the **Import Data** page, select **MYSQL BinLog**.

3. Select a Logstore, and then click Next.

Select an existing project and Logstore. You can also click **Create Now** to create a project and Logstore. For more information, see [Manage a Logstore](#).

If you enter the configuration procedure by clicking the plus sign (+) next to **Data Import** under a Logstore on the Logstores tab, the system skips this step.

4. Create a server group.

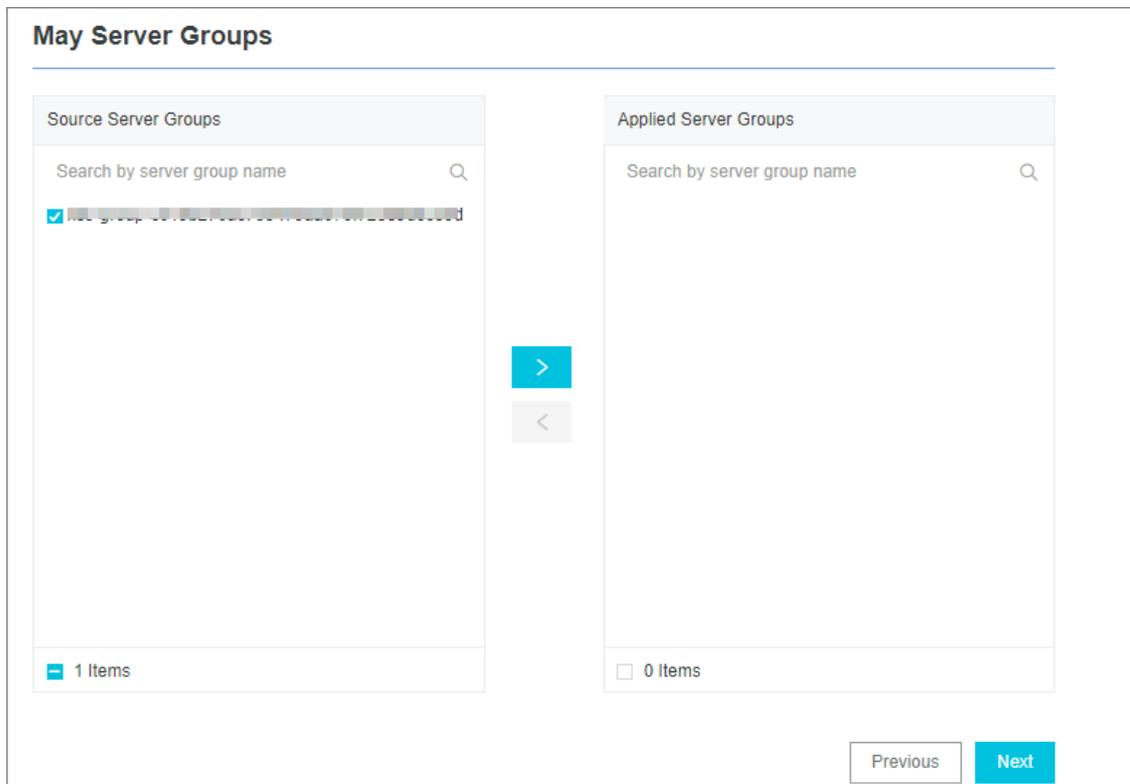
Before you create a server group, ensure that Logtail is installed.

Install Logtail as prompted. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).

After you install Logtail, click **Complete Installation** to create a server group. For more information, see [Overview](#). If you have created a server group, click **Use Existing Server Groups** to select the server group.

5. Configure the server group, and then click Next.

Select a server group and move the group from **Source Server Groups** to **Applied Server Groups**.



6. Configure the data source.

Set the Config Name and Plug-in Config fields.

In the **Plug-in Config** field, modify the parameter settings in the default configuration template based on your business requirements.

```
{
  "inputs": [
    {
      "type": "service_canal",
      "detail": {
        "Host": "*****.mysql.rds.aliyuncs.com",
        "Port": 3306,
        "User": "root",
        "ServerID": 56321,
        "Password": "*****",
        "IncludeTables": [
          "user_info\\..*"
        ],
        "ExcludeTables": [
          ".*\\.\\S+_inner"
        ],
        "TextToString": true,
        "EnableDDL": true
      }
    }
  ]
}
```

- o *inputs*: specifies the collection configurations. This parameter is required. You must configure statements to collect data based on your data source.
 - o *processors*: specifies the processing method. This parameter is optional. For more information about how to set a processing method, see [Configure data processing methods](#).
7. Configure an index.
- Configure an index based on your business requirements. For more information, see [Enable the index feature and configure indexes for a Logstore](#).
8. (Optional)Modify the configurations on the server.

If you do not enter the actual URL, username, or password in **Plug-in Config**, you must replace them with actual values after the configurations are synchronized to the server.

Note If you have entered the actual information, skip this step.

- i. Log on to the server where Logtail is installed, find the `service_canal` keyword in the `/usr/local/ilogtail/user_log_config.json` file, and then set related fields. These fields include `Host`, `User`, and `Password`.
- ii. Run the following command to restart Logtail:

```
sudo /etc/init.d/ilogtaild stop; sudo /etc/init.d/ilogtaild start
```

The configurations that are used to collect binary logs are completed. If changes are made to your database, Logtail immediately collects the updated data and uploads the data to Log Service.

Note Logtail collects incremental binary logs. If no data is collected, check whether changes are made to the table in your database after the configurations are updated.

Metadata fields

When you collect binary logs, some metadata is also uploaded. The following table lists the fields of uploaded metadata.

| Parameter | Description | Example |
|-------------------------|--|---|
| <code>_host_</code> | The name of the host where the database resides. | <code>*****.mysql.rds.aliyuncs.com</code> |
| <code>_db_</code> | The name of the RDS database. | <code>my-database</code> |
| <code>_table_</code> | The name of the table. | <code>my-table</code> |
| <code>_event_</code> | The type of the event. Valid values: | <code>row update</code> , <code>row_insert</code> , and <code>row_delete</code> |
| <code>_id_</code> | The ID of the current collection. The value starts from 0 and increments by 1 each time a binary log event is collected. | <code>1</code> |
| <code>_gtid_</code> | The GTID. | <code>7d2ea78d-b631-11e7-8afb-00163e0eef52:536</code> |
| <code>_filename_</code> | The name of the binary log file. | <code>binlog.001</code> |
| <code>_offset_</code> | The offset of the binary log file. The value is updated after each COMMIT operation. | <code>12876</code> |

Example

After you completed the preceding steps to set a processing method, perform `INSERT` , `UPDATE` , and `DELETE` operations on the `SpecialAlarm` table in the `user_info` database. The following information shows the schema, database operations, and sample logs that are collected by Logtail.

- Schema

```
CREATE TABLE `SpecialAlarm` (
  `id` int(11) unsigned NOT NULL AUTO_INCREMENT,
  `time` datetime NOT NULL,
  `alarmtype` varchar(64) NOT NULL,
  `ip` varchar(16) NOT NULL,
  `count` int(11) unsigned NOT NULL,
  PRIMARY KEY (`id`),
  KEY `time` (`time`) USING BTREE,
  KEY `alarmtype` (`alarmtype`) USING BTREE
) ENGINE=MyISAM AUTO_INCREMENT=1;
```

- Database operations

Perform the `INSERT` , `DELETE` , and `UPDATE` operations on the database.

```
insert into specialalarm (`time`, `alarmType`, `ip`, `count`) values(now(), "NO_ALARM", "10.10. **.***", 55);
delete from specialalarm where id = 4829235 ;
update specialalarm set ip = "10.11. **.***" where id = "4829234";
```

Create an index for `zc.specialalarm` .

```
ALTER TABLE `zc`.`specialalarm`
ADD INDEX `time_index` (`time` ASC);
```

- Sample logs

On the data preview or Search & Analysis page, you can view a sample log that corresponds to each operation.

◦ INSERT statement

```
__source__: 10.30.**.**
__tag__:__hostname__: iZbp145dd9fccu****
__topic__:
_db_: zc
_event_: row_insert
_gtid_: 7d2ea78d-b631-11e7-8afb-00163e0eef52:536
_host_: *****.mysql.rds.aliyuncs.com
_id_: 113
_table_: specialalarm
alarmtype: NO_ALARM
count: 55
id: 4829235
ip: 10.10.**.**
time: 2017-11-01 12:31:41
```

◦ DELETE statement

```
__source__: 10.30.**.**
__tag__:__hostname__: iZbp145dd9fccu****
__topic__:
_db_: zc
_event_: row_delete
_gtid_: 7d2ea78d-b631-11e7-8afb-00163e0eef52:537
_host_: *****.mysql.rds.aliyuncs.com
_id_: 114
_table_: specialalarm
alarmtype: NO_ALARM
count: 55
id: 4829235
ip: 10.10.**.**
time: 2017-11-01 12:31:41
```

◦ UPDATE statement

```
__source__: 10.30.**.**
__tag__:__hostname__: iZbp145dd9fccu****
__topic__:
_db_: zc
_event_: row_update
_gtid_: 7d2ea78d-b631-11e7-8afb-00163e0eef52:538
_host_: *****.mysql.rds.aliyuncs.com
_id_: 115
_old_alarmtype: NO_ALARM
_old_count: 55
_old_id: 4829234
_old_ip: 10.10.22.133
_old_time: 2017-10-31 12:04:54
_table_: specialalarm
alarmtype: NO_ALARM
count: 55
id: 4829234
ip: 10.11.**.**
time: 2017-10-31 12:04:54
```

- DDL statement

```

__source__: 10.30.**.**
__tag__:__hostname__: iZbp145dd9fccu****
__topic__:
_db_: zc
_event_: row_update
_gtid_: 7d2ea78d-b631-11e7-8afb-00163e0eef52:539
_host_: *****.mysql.rds.aliyuncs.com
ErrorCode: 0
ExecutionTime: 0
Query: ALTER TABLE `zc`.`specialalarm`
ADD INDEX `time_index` (`time` ASC)
StatusVars:

```

Usage notes

We recommend that you increase resource limits on Logtail to process traffic surges and prevent data security risks. If the limits are exceeded, Logtail may be forcibly restarted.

You can modify the resource limits in the `/usr/local/ilogtail/ilogtail_config.json` file. Then, you can run the `sudo /etc/init.d/ilogtaild stop;sudo /etc/init.d/ilogtaild start` command to restart Logtail.

The following example shows how to set the CPU limit to two and memory limit to 2,048 MB:

```

{
  ...
  "cpu_usage_limit":2,
  "mem_usage_limit":2048,
  ...
}

```

23.3.1.5.2. Collect MySQL query results

This topic describes how to configure Logtail in the Log Service console to collect MySQL query results.

Context

If you want to collect data from a MySQL database, you can install Logtail on a server and connect the server with the database. Then, you can create a Logtail configuration file by executing a custom SQL statement in the Log Service console and deliver the configuration file to Logtail. Logtail can use the custom SQL statement to collect data from the database at regular intervals.

 **Note** This feature applies only to Logtail 0.16.0 and later versions that run on Linux. For more information, see [Install Logtail in Linux](#).

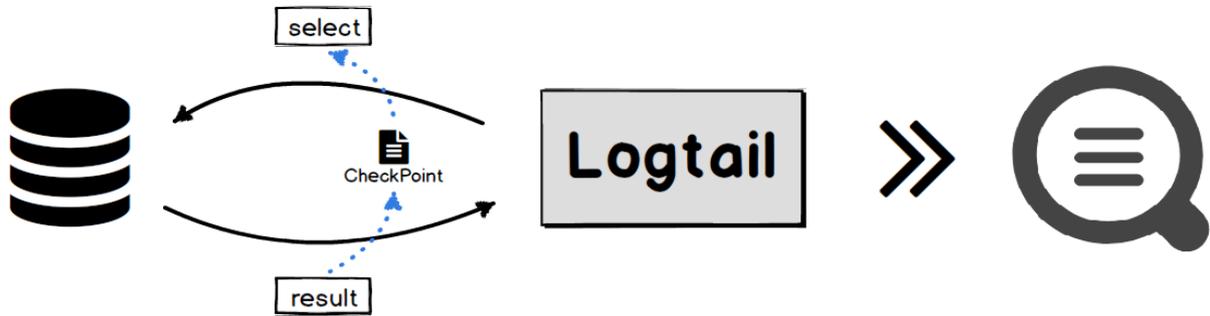
Benefits

Supports MySQL databases such as ApsaraDB RDS for MySQL. When you collect data, you can perform the following operations:

- Paginate query results.
- Set time zones.
- Set timeout periods.
- Store checkpoints.
- Transmit data over the SSL protocol.
- Set the maximum size of data that can be collected at a time.

Implementation

Logtail executes the specified SELECT statement at a regular interval based on the Logtail configurations, and then uploads the query results to Log Service.



When Logtail obtains a query result, Logtail saves the value of the CheckPoint field in the on-premises server. The next time Logtail executes the SELECT statement, Logtail adds the saved value of the CheckPoint field to the SELECT statement. This way, Logtail collects the incremental data of MySQL databases.

Scenarios

- Collect incremental data based on specific marks such as an auto-increment ID or a point in time.
- Customize data synchronization based on specified filtering conditions.

Parameters

The following table describes Logtail parameters. The type of the data source is `service_mysql`.

| Parameter | Type | Required | Description |
|---------------|--------|----------|--|
| Address | string | No | The address of the MySQL database. Default value: 127.0.0.1:3306. |
| User | string | No | The username of the MySQL database. Default value: root. |
| Password | string | No | The password of the MySQL database. By default, this parameter is unspecified. |
| DialTimeOutMs | int | No | The timeout period for the database connection. Unit: milliseconds. Default value: 5000. |
| ReadTimeOutMs | int | No | The timeout period for data reading. Unit: milliseconds. Default value: 5000. |
| StateMent | string | Yes | The SQL statement. |
| Limit | bool | No | Specifies whether to paginate query results by using the LIMIT clause. Default value: false. |

| Parameter | Type | Required | Description |
|-----------------------|--------|----------|---|
| PageSize | int | No | The number of log entries to return on each page. You must specify this parameter if you set the Limit parameter to true. |
| MaxSyncSize | int | No | The maximum number of log entries that are synchronized at a time. Default value: 0. This value indicates that no limit is set for the size of data that can be synchronized at a time. |
| CheckPoint | bool | No | Specifies whether to use checkpoints during data collection. Default value: false. |
| CheckPointColumn | string | No | The name of the checkpoint column. You must specify this parameter if you set the CheckPoint parameter to true. |
| CheckPointColumnType | string | No | The type of the checkpoint column. Valid values: <code>int</code> and <code>time</code> . |
| CheckPointStart | string | No | The initial value of the checkpoint. |
| CheckPointSavePerPage | bool | No | If this parameter is set to true, a checkpoint is saved each time query results are paginated. If this parameter is set to false, a checkpoint is saved after each synchronization. |
| IntervalMs | int | Yes | The synchronization interval. Unit: milliseconds. |

Limits

- We recommend that you paginate query results by specifying the `Limit` parameter. If you set the `Limit` parameter to true, the `LIMIT` clause is automatically appended to the SQL statement specified by the `StateMent` parameter when you run a query.
- If you set the `CheckPoint` parameter to true, the data that is selected based on the `StateMent` parameter must contain the checkpoint column. In addition, the `WHERE` clause must contain the checkpoint field. The value of the checkpoint field is a question mark (`?`).
For example, assume that the checkpoint is "id" and the value of the `StateMent` parameter is `SELECT * from ... where id > ?`.
- If you set the `CheckPoint` parameter to true, you must specify the `CheckPointColumn`, `CheckPointColumnType`

, and `CheckPointStart` parameters.

- The value of `CheckPointColumnType` can only be set to `int` or `time`. If the value is set to `int`, the int64 data type is used for internal storage. If the value is set to `time`, MySQL DATE, DATETIME, and TIME are supported.

Procedure

The following procedure describes how to synchronize incremental data from a MySQL database to Log Service. In this procedure, the `logtail.VersionOs` field is synchronized every 10 seconds. The value of the count parameter in this field is greater than 0. The value of the initial checkpoint is 2017-09-25 11:00:00. Log entries are paginated and each page contains 100 log entries. The checkpoint of each page is saved.

1. [Log on to the Log Service console.](#)

2. Select a data source.

Click **Import Data**. On the **Import Data** page, select **MYSQL Query Result -Plug-in**.

3. Select a Logstore, and then click Next.

Select an existing project and Logstore. You can also click **Create Now** to create a project and Logstore. For more information, see [Manage a Logstore](#).

If you enter the configuration procedure by clicking the plus sign (+) next to **Data Import** under a Logstore on the Logstores tab, the system skips this step.

4. Create a server group.

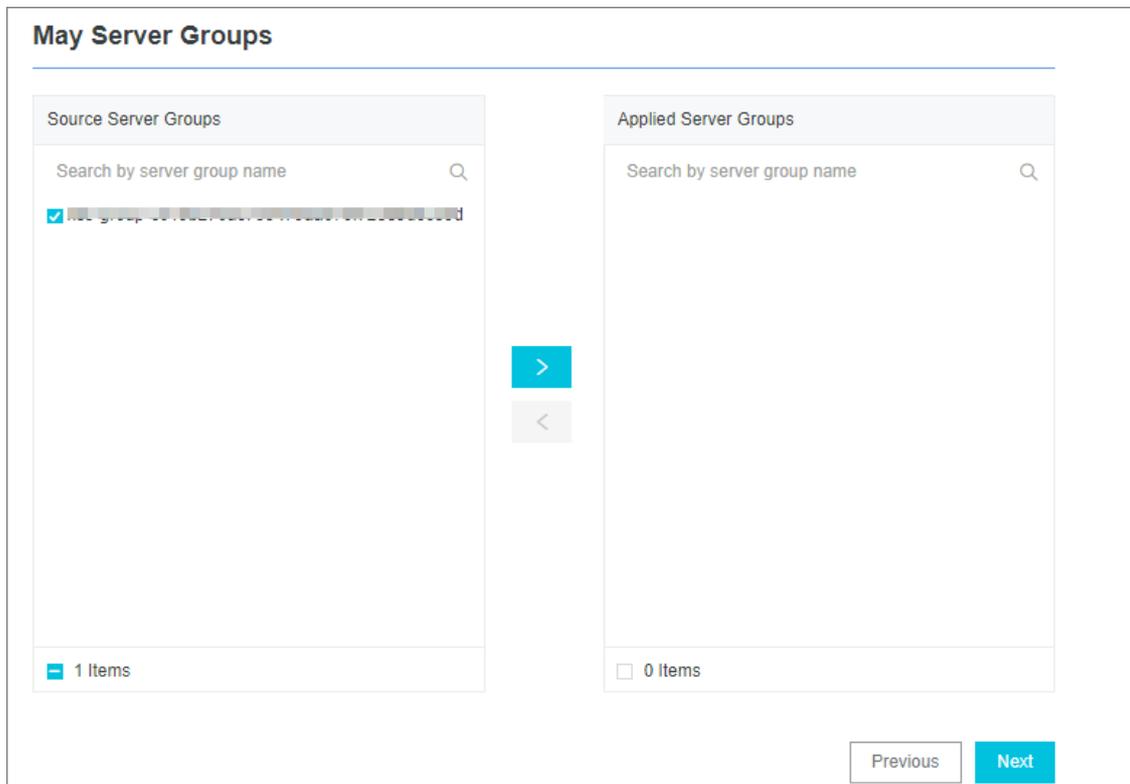
Before you create a server group, ensure that Logtail is installed.

Install Logtail as prompted. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).

After you install Logtail, click **Complete Installation** to create a server group. For more information, see [Overview](#). If you have created a server group, click **Use Existing Server Groups** to select the server group.

5. Configure the server group, and then click Next.

Select a server group and move the group from **Source Server Groups** to **Applied Server Groups**.



6. Configure the data source.

- In the **Plug-in Config** field, modify the parameter settings in the default configuration template based on your business requirements.
- **inputs** : specifies the collection configurations. This parameter is required. **processors** : specifies the processing method. This parameter is optional. You must configure statements to collect data based on your data source. You can specify one or more processing methods. For more information, see [Configure data processing methods](#).

Note If you require a high level of data security, we recommend that you set both the username and the password to xxx. After your configurations are synchronized to the on-premises server, find the Password parameter in the /usr/local/ilogtail/user_log_config.json file and modify the value.

The following example shows the configurations:

```
{
  "inputs": [
    {
      "type": "service_mysql",
      "detail": {
        "Address": "*****:3306",
        "User": "logtail",
        "Password": "*****",
        "DataBase": "logtail",
        "Limit": true,
        "PageSize": 100,
        "StateMent": "SELECT * from logtail.VersionOs where time > ?",
        "CheckPoint": true,
        "CheckPointColumn": "time",
        "CheckPointStart": "2017-09-25 11:00:00",
        "CheckPointSavePerPage": true,
        "CheckPointColumnType": "time",
        "IntervalMs": 10000
      }
    }
  ]
}
```

7. Configure an index.

Configure an index based on your business requirements. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

8. (Optional) Modify the configurations on the server.

If you do not enter the actual URL, username, or password in the **Specify Data Source** page, you must replace them with actual values after the configurations are synchronized to the server.

Note If you have entered the actual information, skip this step.

- Log on to the server where Logtail is installed, find the **service_mysql** keyword in the /usr/local/ilogtail/user_log_config.json file, and then set related fields. These fields include: **Address** , **User** , and **Password** .
- Run the following command to restart Logtail:

```
sudo /etc/init.d/ilogtaild stop; sudo /etc/init.d/ilogtaild start
```

Example

After you configure the processing method, you can view the processed data in the Log Service console. The following information shows the schema and sample logs that are collected by Logtail.

- Schema

```
CREATE TABLE `VersionOs` (  
  `id` int(11) unsigned NOT NULL AUTO_INCREMENT COMMENT 'id',  
  `time` datetime NOT NULL,  
  `version` varchar(10) NOT NULL DEFAULT "",  
  `os` varchar(10) NOT NULL,  
  `count` int(11) unsigned NOT NULL,  
  PRIMARY KEY (`id`),  
  KEY `timeindex` (`time`)  
)
```

- Sample output

```
"count": "4"  
"id": "721097"  
"os": "Windows"  
"time": "2017-08-25 13:00:00"  
"version": "1.3.0"
```

23.3.1.5.3. Collect syslogs

This topic describes how to use the syslog plug-in of Logtail to collect syslogs from a server.

Prerequisites

Logtail 0.16.13 or a later version is installed on the server.

Overview

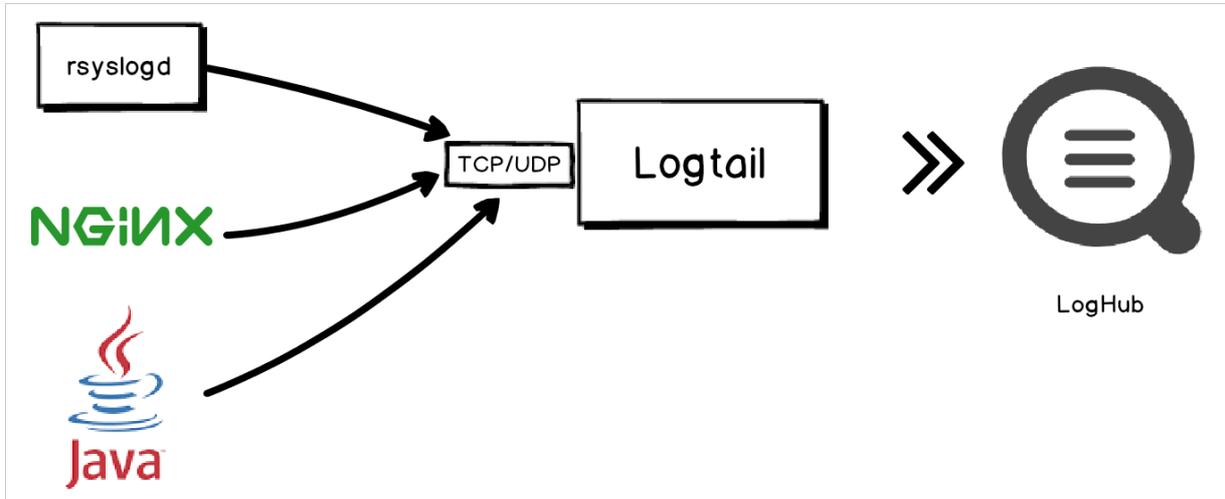
On a Linux server, local syslogs can be forwarded to the IP address and port of a specified server by using syslog agents such as rsyslog. After you create a Logtail configuration for the specified server, the syslog plug-in of Logtail receives the syslogs over TCP or UDP. In addition, the syslog plug-in parses the received syslogs and extract log fields such as facility, tag (program), severity, and content based on the specified syslog protocol. The syslog protocol can be RFC 3164 or RFC 5424.

Note

- Logtail installed on a Windows server does not support the syslog plug-in.
- You can configure multiple syslog plug-ins for Logtail. For example, you can use both TCP and UDP to listen on 127.0.0.1:9999.

Implementation

After the syslog plug-ins start to listen on a specified IP address and port, Logtail can act as a syslog server to collect syslogs from various data sources. These syslogs include system logs collected by rsyslog, access or error logs forwarded by NGINX, and logs forwarded by syslog clients in languages such as Java.



Logtail parameters

The following table describes Logtail parameters. The type of the input is `service_syslog`.

| Parameter | Type | Required | Description |
|---------------|--------|----------|---|
| Address | String | No | <p>The protocol, address, and port on which the syslog plug-in listens. The syslog plug-in obtains logs based on the value of this parameter. Format: <code>[tcp/udp]://[ip]:[port]</code>. Default value: <code>tcp://127.0.0.1:9999</code>.</p> <div style="background-color: #e0f2f7; padding: 10px; border: 1px solid #ccc;"> <p>Note</p> <ul style="list-style-type: none"> The specified protocol, address, and port must be the same as those specified for the forwarding rule in the rsyslog configuration file. If the server on which Logtail is installed has multiple IP addresses, you can set the IP address to 0.0.0.0. It means that the syslog plug-in listens on all IP addresses of the server. </div> |
| ParseProtocol | String | No | <p>The protocol that is used to parse logs. This parameter is not specified by default, indicating that logs are not parsed. Valid values:</p> <ul style="list-style-type: none"> <code>rfc3164</code>: The RFC 3164 protocol is used to parse logs. <code>rfc5424</code>: The RFC 5424 protocol is used to parse logs. <code>auto</code>: The syslog plug-in selects a protocol based on the log content. |

| Parameter | Type | Required | Description |
|--------------------|---------|----------|--|
| IgnoreParseFailure | Boolean | No | Specifies whether to ignore a parsing failure. Default value: true. Valid values: <ul style="list-style-type: none"> true: Logs that fail to be parsed are included in the returned content field. false: Logs that fail to be parsed are dropped. |

Default fields

| Field | Type | Description |
|-----------------|--------|---|
| _hostname_ | String | The hostname. If a hostname is not provided in the log, the hostname of the current host is obtained. |
| _program_ | String | The tag field in the protocol. |
| _priority_ | String | The priority field in the protocol. |
| _facility_ | String | The facility field in the protocol. |
| _severity_ | String | The severity field in the protocol. |
| _unixtimestamp_ | String | The timestamp of the log. |
| _content_ | String | The log content. If the log fails to be parsed, this field contains the complete content of the log. |
| _ip_ | String | The IP address of the current host. |

Configure the plug-in of Logtail to collect syslogs

1. Add a forwarding rule for rsyslog.

Modify the `/etc/rsyslog.conf` rsyslog configuration file on the server from which syslogs are collected. Add a forwarding rule at the end of the configuration file. Then, rsyslog forwards syslogs to the specified IP address and port.

- If you want to collect syslogs of the server by using Logtail on this server, set the forwarding address to 127.0.0.1 and the port to an idle port.
- If you want to collect syslogs of the server by using Logtail on a second server (Server B), set the forwarding address to the public IP address of the second server and port to an idle port.

For example, the following forwarding rule indicates that logs are forwarded to 127.0.0.1:9000 over TCP.

```
*.* @127.0.0.1:9000
```

2. Run the following command to restart rsyslog and validate the log forwarding rule:

```
sudo service rsyslog restart
```

3. [Log on to the Log Service console.](#)
4. Select the data source **Custom Data Plug-in**.

You can use one of the following three methods to select a data source:

- On the homepage of the Log Service console, select a data source in the **Import Data** section.
- In the **Projects** section, click a project name. On the **Overview** page, click **Import Data**, and then select a data source.
- On the Logstores tab in the left-side navigation pane, find a Logstore and click the closing angle bracket (>)

in front of the Logstore name. Click the plus sign (+) next to **Data Import**, and then select a data source.

5. Select a Logstore, and then click Next.

Select an existing project and Logstore. You can also click **Create Now** to create a project and Logstore. For more information, see [Manage a Logstore](#).

If you enter the configuration procedure by clicking the plus sign (+) next to **Data Import** under a Logstore on the Logstores tab, the system skips this step.

6. Create a server group.

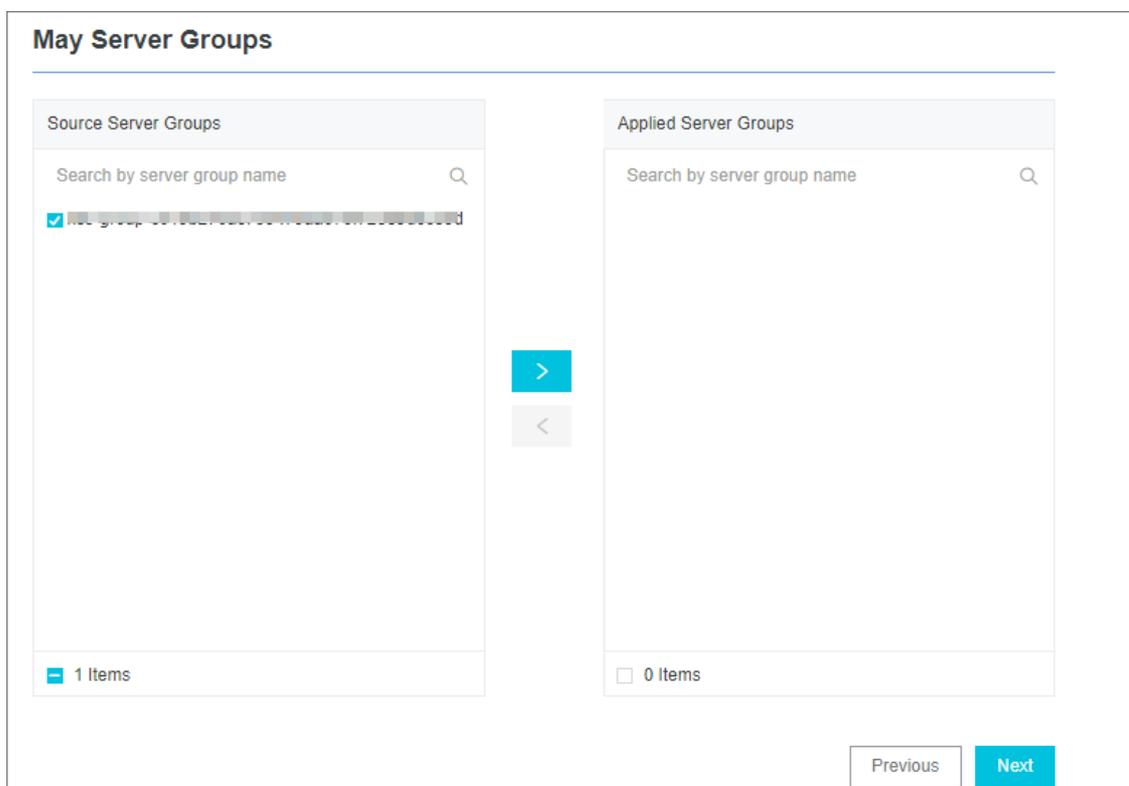
Before you create a server group, ensure that Logtail is installed.

Install Logtail as prompted. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).

After you install Logtail, click **Complete Installation** to create a server group. For more information, see [Overview](#). If you have created a server group, click **Use Existing Server Groups** to select the server group.

7. Configure the server group, and then click Next.

Select a server group and move the group from **Source Server Groups** to **Applied Server Groups**.



8. Configure the data source.

Set **Config Name** and **Plug-in Config**.

The **inputs** section is required. It specifies the collection configuration. The **processors** section is optional. It specifies the processing configuration. You must specify a collection statement for the collection configuration based on the data source. You can specify one or more processing methods for the processing configuration. For more information, see [Configure data processing methods](#).

The following sample code shows how to use UDP and TCP to listen on 127.0.0.1:9000:

```
{
  "inputs": [
    {
      "type": "service_syslog",
      "detail": {
        "Address": "tcp://127.0.0.1:9000",
        "ParseProtocol": "rfc3164"
      }
    },
    {
      "type": "service_syslog",
      "detail": {
        "Address": "udp://127.0.0.1:9001",
        "ParseProtocol": "rfc3164"
      }
    }
  ]
}
```

9. Configure an index.

Configure an index based on your business requirements. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

Configure the plug-in of Logtail to collect NGINX logs

NGINX access logs can be forwarded to specified addresses and ports over the syslog protocol. To deliver NGINX access logs as syslogs from a server to Log Service, you can create a Logtail configuration and apply it to the server group to which the server belongs.

1. Add a forwarding rule to the `nginx.conf` configuration file on the NGINX server.

For example, add the following content to the configuration file.

```
http {
  ...
  # Add this line.
  access_log syslog:server=127.0.0.1:9000,facility=local7,tag=nginx,severity=info combined;
  ...
}
```

2. Run the following command to restart the NGINX service and validate the configuration.

```
sudo service nginx restart
```

3. Create a Logtail configuration and apply it to the server group to which the server belongs.

For more information, see [Configure the plug-in of Logtail to collect syslogs](#).

4. Check whether the Logtail configuration takes effect.

Run the `curl http://127.0.0.1/test.html` command in shell to generate an access log. If the Logtail configuration takes effect, you can view the log information on the query page of the Log Service console.

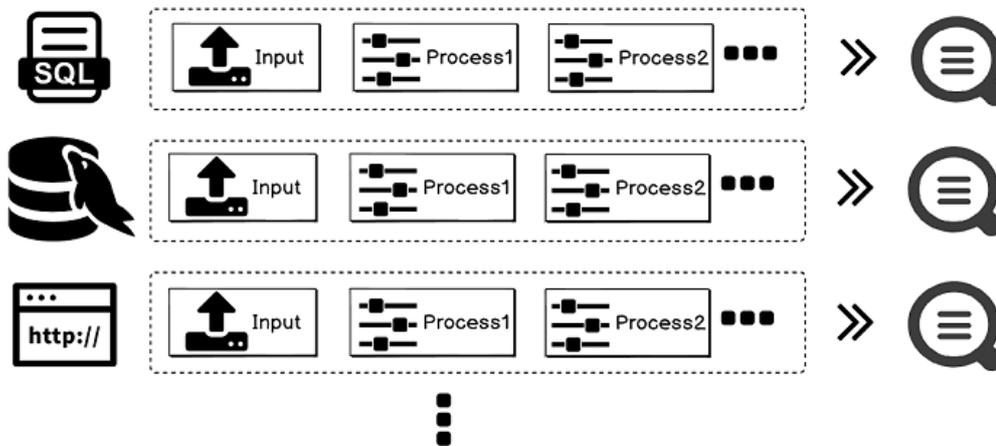
23.3.1.5.4. Configure data processing methods

Logtail allows you to configure plug-ins for data processing. Each plug-in defines a processing method. You can configure one or more processing methods for a data source. Then, Logtail executes the processing methods in sequence. This topic describes the available data processing methods and describes how to configure the methods.

Implementation

The following figure shows how collected data is processed.

Implementation



Plug-in elements

When you configure the data processing methods, you must set the key parameter to `processors` and the value parameter to an array of JSON objects. Each object contains the details of a processing method.

Each object contains the `type` and `detail` fields. The `type` field specifies the type of a processing method. The `detail` field contains configuration details.

```
"processors": [
  {
    "type": "processor_anchor",
    "detail": {
      ...
    }
  },
  {
    "type": "processor_regex",
    "detail": {
      ...
    }
  }
]
```

Processing methods

The following processing methods are supported:

- Extract log fields by using a regular expression
- Extract log fields by using start and stop keywords
- Extract log fields by using a single-character delimiter
- Extract log fields by using a multi-character delimiter
- Convert an IP address into a geographical location
- Filter log fields by using a regular expression
- Add log fields
- Remove log fields
- Extract log time (Go)
- Expand log fields (JSON)
- Combine log fields (JSON)

- [Rename fields](#)
- [Extract log time \(Strptime\)](#)

You can also create a custom method that includes several of the preceding methods. For more information, see [Custom methods](#).

Extract log fields by using a regular expression

This method extracts the fields that match a specified regular expression.

The type of the plug-in is `processor_regex`.

Parameters

| Parameter | Type | Required | Description |
|--------------|--------------|----------|---|
| SourceKey | string | Yes | The name of the field to extract by using the regular expression. |
| Regex | string | Yes | The regular expression. Enclose the value of fields to extract with parentheses <code>()</code> . |
| Keys | String array | Yes | The array of fields to extract, for example, <code>["key1", "key2" ...]</code> . |
| NoKeyError | bool | No | Specifies whether to report an error if no field matches the regular expression. Default value: false. |
| NoMatchError | bool | No | Specifies whether to report an error if the specified regular expression does not match logs. Default value: false. |
| KeepSource | bool | No | Specifies whether to return the SourceKey parameter. Default value: false. This value specifies that the SourceKey parameter is not returned. |
| FullMatch | bool | Yes | Specifies whether to extract fields that exactly match the <code>Regex</code> parameter. Default value: true. This value specifies that fields that partially match the <code>Regex</code> parameter are extracted. |

The following example shows how to extract fields from an access log.

- Input data

```
"content": "10.200. **. ** - [10/Aug/2017:14:57:51 +0800] \"POST /PutData?
Category=YunOsAccountOpLog&AccessKeyId=<yourAccessKeyId>&Date=Fri%2C%2028%20Jun%202013%2006%3A53
%3A30%20GMT&Topic=raw&Signature=<yourSignature> HTTP/1.1\" 0.024 18204 200 37 \"-\" \"aliyun-sdk-java"
```

- Configurations

```
{
  "type": "processor_regex",
  "detail": {"SourceKey": "content",
    "Regex": "([\d\.]+) \|S+ \|S+ \|([\S+ ] \|S+ \|)"(\w+) ([^\|"]*)"([\d\.]+) (\d+) (\d+) (\d+|-)"([\^\|"]*)" \|"([\^\|"]*)"
  \| (\d+)",
    "Keys": ["ip", "time", "method", "url", "request_time", "request_length", "status", "length", "ref_url", "browser"],
    "NoKeyError": true,
    "NoMatchError": true,
    "KeepSource": false
  }
}
```

- Output data

```
"ip": "10.200.**.*"
"time": "10/Aug/2017:14:57:51"
"method": "POST"
"url": "/PutData? Category=YunOsAccountOpLog&AccessKeyId=<yourAccessKeyId>&Date=Fri%2C%2028%20Jun%202013%2006%3A53%3A30%20GMT&Topic=raw&Signature=<yourSignature>"
"request_time": "0.024"
"request_length": "18204"
"status": "200"
"length": "27"
"ref_url": "-"
"browser": "aliyun-sdk-java"
```

Extract log fields by using start and stop keywords

This method uses the `start` and `stop` keywords to extract fields. You can extract fields from the JSON string between the start and stop keywords. You can also expand the JSON string into other forms.

The type of the plug-in is `processor_anchor`.

Parameters

| Parameter | Type | Required | Description |
|---------------|--------|----------|---|
| SourceKey | string | Yes | The name of the field to extract. |
| Anchors | Array | Yes | The array of anchors. For more information, see the following table. |
| NoAnchorError | bool | No | Specifies whether to report an error if no specified keyword is found. Default value: false. |
| NoKeyError | bool | No | Specifies whether to report an error if no match exists. Default value: false. |
| KeepSource | bool | No | Specifies whether to return the SourceKey parameter. Default value: false. This value specifies that the SourceKey parameter is not returned. |

Anchor fields

| Parameter | Type | Required | Description |
|-----------|------|----------|-------------|
|-----------|------|----------|-------------|

| Parameter | Type | Required | Description |
|-----------------|--------|----------|---|
| Start | string | Yes | The start keyword. If you do not specify the parameter, Logtail matches the first character of a string. |
| Stop | string | Yes | The stop keyword. If you do not specify the parameter, Logtail matches the last character of a string. |
| FieldName | string | Yes | The name of the field to extract. |
| FieldType | string | Yes | The type of the field to extract. Valid values: <code>string</code> and <code>json</code> . |
| ExpondJson | bool | No | Specifies whether to expand JSON strings. Default value: <code>false</code> . This parameter is available only if you specify <code>json</code> for the <code>FieldType</code> parameter. |
| ExpondConnector | string | No | The connector that combines separate field names into a string. Default value: <code>_</code> . |
| MaxExpondDepth | int | No | The maximum depth of JSON expansion. Default value: <code>0</code> . This indicates no limit. |

The following example shows how to use this method to process input data of multiple types.

- Input data

```
"content" : "time:2017.09.12 20:55:36\tjson:{\"key1\" : \"xx\", \"key2\" : false, \"key3\" : 123.456, \"key4\" : { \"inner1\" : 1, \"inner2\" : false}}
```

- Configurations

```
{
  "type" : "processor_anchor",
  "detail" : {"SourceKey" : "content",
    "Anchors" : [
      {
        "Start" : "time",
        "Stop" : "\t",
        "FieldName" : "time",
        "FieldType" : "string",
        "ExpondJson" : false
      },
      {
        "Start" : "json:",
        "Stop" : "",
        "FieldName" : "val",
        "FieldType" : "json",
        "ExpondJson" : true
      }
    ]
  }
}
```

- Output data

```
"time": "2017.09.12 20:55:36"
"val_key1": "xx"
"val_key2": "false"
"val_key3": "123.456"
"value_key4_inner1": "1"
"value_key4_inner2": "false"
```

Extract log fields by using a single-character delimiter

This method uses `single-character delimiters` to split logs into several fields. You can enclose delimited fields by using characters that you specify in the `Quote` parameter.

The type of the plug-in is `processor_split_char`.

Parameters

| Parameter | Type | Required | Description |
|--------------|--------------|----------|--|
| SourceKey | string | Yes | The name of the field to extract. |
| SplitSep | string | Yes | The single-character delimiter. You must specify a single character as a delimiter. You can specify a non-printable character such as <code>\u0001</code> as a delimiter. |
| SplitKeys | String array | Yes | The names of the fields that are split, for example, ["key1", "key2" ...]. |
| QuoteFlag | bool | No | Specifies whether to use the <code>Quote</code> parameter. Default value: false. |
| Quote | string | No | You must specify a single character. The parameter is available only if you specify true for the <code>QuoteFlag</code> parameter. You can specify non-printable characters, such as <code>\u0001</code> . |
| NoKeyError | bool | No | Specifies whether to report an error if no match exists. Default value: false. |
| NoMatchError | bool | No | Specifies whether to report an error if the specified delimiter is not found. Default value: false. |
| KeepSource | bool | No | Specifies whether to return the <code>SourceKey</code> parameter. Default value: false. This value specifies that the <code>SourceKey</code> parameter is not returned. |

The following example shows how to use this method to process logs.

- Input data

```
"content": "10. **. **. **|10/Aug/2017:14:57:51 +0800|POST|PutData?
Category=YunOsAccountOpLog&AccessKeyId=<yourAccessKeyId>&Date=Fri%2C%2028%20Jun%202013%2006%3A53
%3A30%20GMT&Topic=raw&Signature=<yourSignature>|0.024|18204|200|37|-|
aliyun-sdk-java"
```

- Configurations

```
{
  "type": "processor_split_char",
  "detail": {"SourceKey": "content",
    "SplitSep": "|",
    "SplitKeys": ["ip", "time", "method", "url", "request_time", "request_length", "status", "length", "ref_url", "browser"]}
}
```

- Output data

```
"ip": "10. ** . ** . **"
"time": "10/Aug/2017:14:57:51 +0800"
"method": "POST"
"url": "/PutData? Category=YunOsAccountOpLog&AccessKeyId=<yourAccessKeyId>&Date=Fri%2C%2028%20Jun%202013%2006%3A53%3A30%20GMT&Topic=raw&Signature=<yourSignature>"
"request_time": "0.024"
"request_length": "18204"
"status": "200"
"length": "27"
"ref_url": "-"
"browser": "aliyun-sdk-java"
```

Extract log fields by using a multi-character delimiter

Similar to the single-character delimiter method, this method uses multi-character delimiters to split logs into several fields. The Quote parameter is unavailable for this method.

The type of the plug-in is `processor_split_string`.

Parameters

| Parameter | Type | Required | Description |
|-----------------|--------------|----------|---|
| SourceKey | string | Yes | The name of the field to extract. |
| SplitSep | string | Yes | The multi-character delimiter. You can specify non-printable characters such as <code>\u0001\u0002</code> . |
| SplitKeys | String array | Yes | The names of the fields that are split, for example, ["key1", "key2" ...]. |
| PreserveOthers | bool | No | Specifies whether to retain excess fields when the number of split fields exceeds the number of fields defined in the <code>SplitKeys</code> parameter. Default value: false. |
| ExpandOthers | bool | No | Specifies whether to parse excess fields. Default value: false. |
| ExpandKeyPrefix | string | No | The prefix of the names of excess fields. For example, if you specify <code>expand_</code> for the parameter, excess fields are named <code>expand_1</code> and <code>expand_2</code> . |
| NoKeyError | bool | No | Specifies whether to report an error if no match exists. Default value: false. |

| Parameter | Type | Required | Description |
|--------------|------|----------|--|
| NoMatchError | bool | No | Specifies whether to report an error if no multi-character delimiter is found. Default value: false. |
| KeepSource | bool | No | Specifies whether to return the SourceKey parameter. Default value: false. The value false specifies that the SourceKey parameter is not returned. |

The following example shows how to use this method to process logs.

- Input data

```
"content": "10. **. *. **|#|10/Aug/2017:14:57:51 +0800|#|POST|#|PutData?
Category=YunOsAccountOpLog&AccessKeyId=<yourAccessKeyId>&Date=Fri%2C%2028%20Jun%202013%2006%3A53
%3A30%20GMT&Topic=raw&Signature=<yourSignature>|#|0.024|#|18204|#|200|#|37|#|-|#|
aliyun-sdk-java"
```

- Configurations

```
{
  "type": "processor_split_string",
  "detail": {"SourceKey": "content",
    "SplitSep": "|#|",
    "SplitKeys": ["ip", "time", "method", "url", "request_time", "request_length", "status"],
    "PreserveOthers": true,
    "ExpandOthers": true,
    "ExpandKeyPrefix": "expand_"
  }
}
```

- Output data

```
"ip": "10. **. *. **"
"time": "10/Aug/2017:14:57:51 +0800"
"method": "POST"
"url": "/PutData? Category=YunOsAccountOpLog&AccessKeyId=<yourAccessKeyId>&Date=Fri%2C%2028%20Jun%20
2013%2006%3A53%3A30%20GMT&Topic=raw&Signature=<yourSignature>"
"request_time": "0.024"
"request_length": "18204"
"status": "200"
"expand_1": "27"
"expand_2": "-"
"expand_3": "aliyun-sdk-java"
```

Convert an IP address into a geographical location

This method converts IP address into a geographical location, such as a country, province, city, or geographical coordinates.

 **Note** By default, the Logtail installation package does not include GeoIP databases. You must download these databases to your local host and configure the required parameters. We recommend that you download a version of the GeoIP database that includes the `City` database.

The type of the plug-in is `processor_geoip` .

Parameters

| Parameter | Type | Required | Description |
|--------------|--------|----------|--|
| SourceKey | string | Yes | The name of the field whose IP address you want to convert into a geographical location. |
| DBPath | string | Yes | The absolute path of the GeoIP database. The database format is MMDB , for example. /user/data/GeoLite2-City_20180102/GeoLite2-City.mmdb . |
| NoKeyError | bool | No | Specifies whether to report an error if no match exists. Default value: false. |
| NoMatchError | bool | No | Specifies whether to report an error if the specified IP address is invalid or does not match the IP addresses stored in the library. Default value: false. |
| KeepSource | bool | No | Specifies whether to return the SourceKey parameter. Default value: true. |
| Language | string | No | The language of the GeoIP database. Default value: zh-CN . Make sure that your GeoIP database can be displayed in a language that is suitable for your business. |

The following example shows how to use this method to convert an IP address into a geographical location.

- Input data

```
"source_ip": "***. **. **. ***"
```

Download a GeoIP database and install the database on the host where Logtail resides. We recommend that you download [MaxMind GeoLite2](#) that includes the city database.

 **Note** Make sure that the database format is MMDB.

- Configurations

```
{
  "type": "processor_geoip",
  "detail": {
    "SourceKey": "ip",
    "NoKeyError": true,
    "NoMatchError": true,
    "KeepSource": true,
    "DBPath": "/user/local/data/GeoLite2-City_20180102/GeoLite2-City.mmdb"
  }
}
```

- Output result

```
"source_ip_city_" : "***. **. **. ***"
"source_ip_province_" : "Zhejiang"
"source_ip_city_" : "Hangzhou"
"source_ip_province_code_" : "ZJ"
"source_ip_country_code_" : "CN"
"source_ip_longitude_" : "120. *****"
"source_ip_latitude_" : "30. *****"
```

Filter log fields by using a regular expression

This method uses regular expressions to filter logs. You can specify conditions in the `Include` and `Exclude` parameters.

The type of the plug-in is `processor_filter_regex`.

Parameters

| Parameter | Type | Required | Description |
|-----------|--|----------|---|
| Include | The JSON object that includes key-value pairs. | No | Each key includes a log field. Each value includes a regular expression. If the field value matches the regular expression, the log is collected. |
| Exclude | The JSON object that includes key-value pairs. | No | Each key includes a log field. Each value includes a regular expression. If the field value matches the regular expression, the log is dropped. |

Note A log is collected only if the field value matches the regular expression specified in the `Include` parameter and does not match the regular expression specified in the `Exclude` parameter. Otherwise, the log is dropped.

The following example shows how to use this method to process logs.

- Input data

- Log 1

```
"ip": "10. *. *. *"  
"method": "POST"  
...  
"browser": "aliyun-sdk-java"
```

- Log 2

```
"ip": "10. *. *. *"  
"method": "POST"  
...  
"browser": "chrome"
```

- Log 3

```
"ip": "192.168. *. *"  
"method": "POST"  
...  
"browser": "ali-sls-ilogtail"
```

- Configurations

```
{
  "type": "processor_filter_regex",
  "detail": {
    "Include": {
      "ip": "10\\.\\.\\*",
      "method": "POST"
    },
    "Exclude": {
      "browser": "aliyun.\\*"
    }
  }
}
```

• Output data

| Log | Matched | Reason |
|-------|---------|---|
| Log 1 | No | The match failed because the value of the browser field matches the regular expression specified in the Exclude parameter. |
| Log 2 | Yes | - |
| Log 3 | No | The match failed because the value of the ip field does not match the regular expression specified in the Include parameter. The regular expression matches IP addresses that start with 10 . |

Add log fields

You can use this method to add multiple fields to a log.

The type of the plug-in is `processor_add_fields` .

 **Note** This method is available for Logtail 0.16.28 or later.

Parameters

| Parameter | Type | Required | Description |
|---------------|------|----------|---|
| Fields | map | No | The JSON object that includes key-value pairs. You can specify multiple key-value pairs in the parameter. |
| IgnoreIfExist | bool | No | Specifies whether to ignore key-value pairs that have the same keys. Default value: false. |

The following example shows how to use this method to add fields to logs.

• Input data

```
"aaa1": "value1"
```

• Configurations

```
{
  "processors":[
    {
      "type":"processor_add_fields",
      "detail":{
        "Fields":{
          "aaa2": "value2",
          "aaa3": "value3"
        }
      }
    }
  ]
}
```

- Output data

```
"aaa1": "value1"
"aaa2": "value2"
"aaa3": "value3"
```

Remove log files

You can use this method to remove specific fields from logs.

The type of the plug-in is `processor_drop`.

 **Note** The method is available for Logtail 0.16.28 or later.

Parameters

| Parameter | Type | Required | Description |
|-----------|-------|----------|--|
| DropKeys | array | No | The array that includes a set of strings. You can remove multiple fields from a log. |

The following example shows how to remove the `aaa1` and `aaa2` fields from a log.

- Input data

```
"aaa1": "value1"
"aaa2": "value2"
"aaa3": "value3"
```

- Configurations

```
{
  "processors":[
    {
      "type":"processor_drop",
      "detail":{
        "DropKeys": ["aaa1","aaa2"]
      }
    }
  ]
}
```

- Output data

```
"aaa3": "value3"
```

Extract log time (Go)

You can use this method to extract time from a log field and convert the time format.

The type of the plug-in is `processor_gotime`.

 **Note** This plug-in is available for Logtail 0.16.28 or later.

Parameters

| Parameter | Type | Required | Description |
|----------------|--------|----------|---|
| SourceKey | string | Yes | The name of the log field from which you want to extract time. |
| SourceFormat | string | Yes | The time that is parsed in Go. |
| SourceLocation | int | Yes | The source time zone. If the parameter is not specified, it indicates the time zone of the local host. |
| DestKey | string | Yes | The destination field. |
| DestFormat | string | Yes | The destination time format in Go. |
| DestLocation | int | No | The destination time zone. If the parameter is not specified, it indicates the time zone of the local host. |
| SetTime | bool | No | Specifies whether to overwrite the original time. Default value: true. |
| KeepSource | bool | No | Specifies whether to return the SourceKey parameter. Default value: true. |
| NoKeyError | bool | No | Specifies whether to report an error if no match exists. Default value: true. |
| AlarmIfFail | bool | No | Specifies whether to send alerts when extraction failed. Default value: true. |

Use `2006-01-02 15:04:05 (UTC +8)` of the `s_key` field as the source time. Convert the time into `2006/01/02 15:04:05 (UTC +9)` and save the new time in the `d_key` field. The following example shows how to use this method to process logs.

- Input data

```
"s_key": "2019-07-05 19:28:01"
```

- Configurations

```
{
  "processors":[
    {
      "type":"processor_gotime",
      "detail":{
        "SourceKey": "s_key",
        "SourceFormat":"2006-01-02 15:04:05",
        "SourceLocation":8,
        "DestKey":"d_key",
        "DestFormat":"2006/01/02 15:04:05",
        "DestLocation":9,
        "SetTime": true,
        "KeepSource": true,
        "NoKeyError": true,
        "AlarmIfFail": true
      }
    }
  ]
}
```

- Output data

```
"s_key":"2019-07-05 19:28:01"
"d_key":"2019/07/05 20:28:01"
```

Expand log fields (JSON)

You can use this method to expand a log field.

The type of the plug-in is `processor_json` .

 **Note** This plug-in is available for Logtail 0.16.28 or later.

Parameters

| Parameter | Type | Required | Description |
|-----------------|--------|----------|---|
| SourceKey | string | Yes | The name of the field you want to expand. |
| NoKeyError | bool | No | Specifies whether to report an error if no match exists. Default value: true. |
| ExpandDepth | int | No | The depth of JSON expansion. The value must be a non-negative integer. Default value: 0. This indicates the depth is not limited. The value n specifies that the number of levels to expand is n. |
| ExpandConnector | string | No | The delimiter that you use to connect expanded levels. Default value: <code>_</code> . You can leave this parameter unspecified. |
| Prefix | string | No | The prefix that you want to add to the name of each new field after expansion. |
| KeepSource | bool | No | Specifies whether to return the SourceKey parameter. Default value: true. |

| Parameter | Type | Required | Description |
|----------------------|------|----------|---|
| UseSourceKeyAsPrefix | bool | No | Specifies whether to retain the name of the source field as the prefix of each new field after expansion. Default value: false. |

The following example shows how to use the field expansion (JSON) method to expand the `s_key` field.

- Input data

```
"s_key":{"k1":{"k2":{"k3":{"k4":{"k51":"51","k52":"52"},"k41":"41"}}}}
```

- Configurations

```
{
  "processors":[
    {
      "type":"processor_json",
      "detail":{
        "SourceKey": "s_key",
        "NoKeyError":true,
        "ExpandDepth":0,
        "ExpandConnector":"-",
        "Prefix":"j",
        "KeepSource": false,
        "UseSourceKeyAsPrefix": true
      }
    }
  ]
}
```

- Output data

```
"s_key":{"k1":{"k2":{"k3":{"k4":{"k51":"51","k52":"52"},"k41":"41"}}}}
"js_key-k1-k2-k3-k4-k51":"51"
"js_key-k1-k2-k3-k4-k52":"52"
"js_key-k1-k2-k3-k41":"41"
```

Combine log fields (JSON)

You can use this method to combine multiple log fields into one field.

The type of the plug-in is `processor_packjson`.

 **Note** This plug-in is available for Logtail 0.16.28 or later.

Parameters

| Parameter | Type | Required | Description |
|------------|--------|----------|---|
| SourceKeys | array | Yes | The array that includes the names of fields that you want to combine. |
| DestKey | string | No | The name of the destination field after combination. |
| KeepSource | bool | No | Specifies whether to return the SourceKey parameter. Default value: true. |

| Parameter | Type | Required | Description |
|-------------------|------|----------|--|
| AlarmIfIncomplete | bool | No | Specifies whether to send alerts if no source fields exist. Default value: true. |

The following example shows how to use the method to combine the `a` and `b` fields into the `d_key` field.

- Input data

```
"a":"1"
"b":"2"
```

- Configurations

```
{
  "processors":[
    {
      "type":"processor_packjson",
      "detail":{
        "SourceKeys":["a","b"],
        "DestKey":"d_key",
        "KeepSource":true,
        "AlarmIfEmpty":true
      }
    }
  ]
}
```

- Output data

```
"a":"1"
"b":"2"
"d_key":"{\\"a\\":\\"1\\",\\"b\\":\\"2\\"}"
```

Rename fields

You can use this method to rename multiple fields.

The type of the plug-in is `processor_rename`.

 **Note** This plug-in is available for Logtail 0.16.28 or later.

Parameters

| Parameter | Type | Required | Description |
|------------|-------|----------|---|
| NoKeyError | bool | No | Specifies whether to report an error if no match exists. Default value: true. |
| SourceKeys | array | Yes | The array that includes the names of fields that you want to rename. |
| DestKeys | array | Yes | The array that includes the new names of the fields. |

The following example shows how to use this method to rename the `aaa1` and `aaa2` fields to `bbb1` and `bbb2` fields.

- Input data

```
"aaa1": "value1"
"aaa2": "value2"
"aaa3": "value3"
```

- Configurations

```
{
  "processors": [
    {
      "type": "processor_rename",
      "detail": {
        "SourceKeys": ["aaa1", "aaa2"],
        "DestKeys": ["bbb1", "bbb2"],
        "NoKeyError": true
      }
    }
  ]
}
```

- Output data

```
"bbb1": "value1"
"bbb2": "value2"
"aaa3": "value3"
```

Extract log time (Strptime)

You can use this method to extract time from a log field and parse the time by using the Linux `strptime()` function.

The type of the plug-in is `processor_strptime`.

 **Note** This plug-in is available for Logtail 0.16.28 or later.

Parameters

| Parameter | Type | Required | Description |
|-----------------|--------|----------|--|
| SourceKey | string | Yes | The name of the field from which you want to extract time. |
| Format | string | Yes | The time format to parse the time. |
| AdjustUTCOffset | bool | No | Specifies whether to adjust the time zone. Default value: false. |
| UTCOffset | int | No | The offset that is used to adjust the time zone. Unit: seconds. For example, if you set the value to 14400, the time zone is changed to UTC+8. |
| AlarmIfFail | bool | No | Specifies whether to send alerts if a field fails to be extracted. |
| KeepSource | bool | No | Specifies whether to return the SourceKey parameter. Default value: true. |

Parse the value of the `log_time` in the `%Y/%m/%d %H:%M:%S` time format and use the time zone of your local host. The following examples show how to use this method to process logs.

- Example 1: The time zone is UTC +8.

- Input data

```
"log_time":"2016/01/02 12:59:59"
```

- Configurations

```
{
  "processors":[
    {
      "type":"processor_strptime",
      "detail":{
        "SourceKey": "log_time",
        "Format": "%Y/%m/%d %H:%M:%S"
      }
    }
  ]
}
```

- Output data

```
"log_time":"2016/01/02 12:59:59"
Log.Time = 1451710799
```

- Example 2: The time zone is UTC +7.

- Input data

```
"log_time":"2016/01/02 12:59:59"
```

- Configurations

```
{
  "processors":[
    {
      "type":"processor_strptime",
      "detail":{
        "SourceKey": "log_time",
        "Format": "%Y/%m/%d %H:%M:%S",
        "AdjustUTCOffset": true,
        "UTCOffset": 25200
      }
    }
  ]
}
```

- Output data

```
"log_time":"2016/01/02 12:59:59"
Log.Time = 1451714399
```

Custom methods

You can use a combination of multiple processing methods to process logs. The following example shows how to use a single-character delimiter to split a log into several fields and then specify anchor points to extract content from the `detail` field.

- Input data

```
"content":  
"ACCESS|QAS|11.**.**|1508729889935|52460dbed4d540b88a973cf5452b1447|1238|appKey=ba,env=pub,requestTi  
me=1508729889913,latency=22ms,  
request={appKey:ba,optional:{domains:\\daily\\,version:\\v2\\},rawQuery:\\query\\:\\The route to Location A\\,  
domain:\\Navigation\\,intent:\\navigate\\,slots:\\to_geo:level3=Location A\\,location:\\Location B\\},  
requestId:52460dbed4d540b88a973cf5452b1447},  
response={answers:[],status:SUCCESS}"
```

- Configurations

```
"processors": [
  {
    "type": "processor_split_char",
    "detail": {"SourceKey": "content",
      "SplitSep": "|",
      "SplitKeys": ["method", "type", "ip", "time", "req_id", "size", "detail"]}
  },
  {
    "type": "processor_anchor",
    "detail": {"SourceKey": "detail",
      "Anchors": [
        {
          "Start": "appKey=",
          "Stop": ",env=",
          "FieldName": "appKey",
          "FieldType": "string"
        },
        {
          "Start": ",env=",
          "Stop": ",requestTime=",
          "FieldName": "env",
          "FieldType": "string"
        },
        {
          "Start": ",requestTime=",
          "Stop": ",latency=",
          "FieldName": "requestTime",
          "FieldType": "string"
        },
        {
          "Start": ",latency=",
          "Stop": ",request=",
          "FieldName": "latency",
          "FieldType": "string"
        },
        {
          "Start": ",request=",
          "Stop": ",response=",
          "FieldName": "request",
          "FieldType": "string"
        },
        {
          "Start": ",response=",
          "Stop": "",
          "FieldName": "response",
          "FieldType": "json"
        }
      ]
    }
  }
]
```

- Output data

```
"method": "ACCESS"
"type": "QAS"
"ip": "***. **. *. **"
"time": "1508729889935"
"req_id": "52460dbed4d540b88a973cf5452b1447"
"size": "1238"
"appKey": "ba"
"env": "pub"
"requestTime": "1508729889913"
"latency": "22ms"
"request": "{appKey:nui-banma,optional:{domains:\\daily-faq\\,version:\\v2\\},rawQuery:{query:\\345\\216\\273\\344\\271\\220\\345\\261\\261\\347\\232\\204\\350\\267\\257\\347\\272\\277\\,domain:\\345\\257\\274\\350\\210\\252\\,intent\\navigate\\,slots\\to_geo:level3=\\344\\271\\220\\345\\261\\261\\,location:\\345\\214\\227\\344\\272\\254\\},requestId:52460dbed4d540b88a973cf5452b1447}"
"response_answers": "[]"
"response_status": "SUCCESS"
```

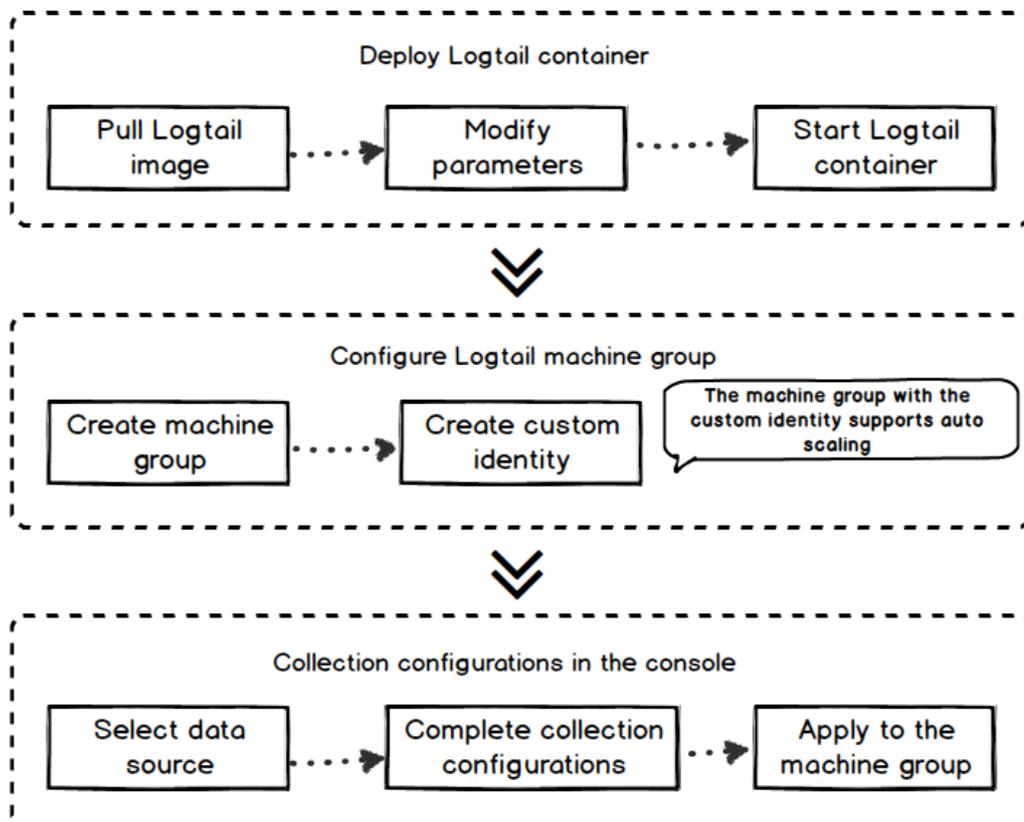
23.3.1.6. Collect container logs

23.3.1.6.1. Collect standard Docker logs

This topic describes how to use Logtail to collect standard Docker logs and upload these logs together with the container metadata to Log Service.

Procedure

Procedure



1. Deploy a Logtail container.

2. Configure a Logtail server group.

Create a server group with a custom ID in the Log Service console. The container cluster can automatically scale up or down based on data traffic.

3. Create a Logtail configuration.

Create a Logtail configuration in the Log Service console. The Logtail configuration process is completed in the Log Service console. No local configuration is needed.

Deploy a Logtail container

1. Run the following command to pull the Logtail image.

```
docker pull registry.cn-hangzhou.aliyuncs.com/log-service/logtail
```

2. Start a Logtail container.

Set the `${your_region_name}`, `${your_aliyun_user_id}`, and `${your_machine_group_user_defined_id}` parameters in the startup template.

```
docker run -d -v /:/logtail_host:ro -v /var/run:/var/run --env
ALIYUN_LOGTAIL_CONFIG=/etc/ilogtail/conf/${your_region_name}/ilogtail_config.json
--env ALIYUN_LOGTAIL_USER_ID=${your_aliyun_user_id} --env
ALIYUN_LOGTAIL_USER_DEFINED_ID=${your_machine_group_user_defined_id} registry.cn-hangzhou.aliyuncs.com/log-service/logtail
```

Notice Before you set the parameters, you must complete one of the following configurations. Otherwise, the `container text file busy` error may occur when you delete another container.

- For CentOS 7.4 and later versions, set `fs.may_detach_mounts` to 1. For more information, see [Bug 1468249](#), [Bug 1441737](#), and [Issue 34538](#).
- Grant the `privileged` permission to Logtail by adding the `--privileged` flag to the startup parameters. For more information, see [Docker run reference](#).

| Parameter | Description |
|---|--|
| <code>\${your_region_name}</code> | The region of the project. For more information, see View the information of a project . |
| <code>\${your_aliyun_user_id}</code> | The user ID. Set this parameter to the ID of your Alibaba Cloud account, which is a string. For information about how to view the ID, see Step 1 in Configure an account ID for a server . |
| <code>\${your_machine_group_user_defined_id}</code> | The custom ID of your server group. For information about how to set the custom ID, see Step 1 in Create a machine group based on a custom ID . |

After you set the parameters, run the following command to start the Logtail container.

```
docker run -d -v /:/logtail_host:ro -v /var/run:/var/run
--env ALIYUN_LOGTAIL_CONFIG=/etc/ilogtail/conf/cn_hangzhou/ilogtail_config.json --env
ALIYUN_LOGTAIL_USER_ID=1654218*****--env ALIYUN_LOGTAIL_USER_DEFINED_ID=log-docker-demo registry.cn-hangzhou.aliyuncs.com/log-service/logtail
```

Notice

You can customize the startup parameters of the Logtail container if the following conditions are met:

- The following environment variables exist before you start the Logtail container: `ALIYUN_LOGTAIL_USER_DEFINED_ID`, `ALIYUN_LOGTAIL_USER_ID`, and `ALIYUN_LOGTAIL_CONFIG`.
- The `/var/run` directory is mounted on the `/var/run` directory of the Logtail container.
- To collect container standard output, container logs, or host files, you must mount the root directory on the `/logtail_host` directory of the Logtail container.
- If an error showing *The parameter is invalid : uuid=none* occurs in the `/usr/local/ilogtail/ilogtail.LOG` Logtail log file, create a file named `product_uuid` on the host. Add a valid UUID such as `169E98C9-ABC0-4A92-B1D2-AA6239C0D261` to the file, and mount the file on the `/sys/class/dmi/id/product_uuid` directory of the Logtail container.

Configure a Logtail server group

1. [Log on to the Log Service console.](#)
2. Click a project name.
3. In the left-side navigation pane, click the **Server Groups** icon to show the server group list.
4. Click the icon next to Server Groups, and then select **Create Server Group**.

You can also create a server group when you import data to Log Service.

5. In the dialog box that appears, select **Custom ID** from the Identifier drop-down list. Enter the value of `ALIYUN_LOGTAIL_USER_DEFINED_ID` set in the previous step in the **Custom Identifier** field.

Click OK. One minute later, click the name of the server group in the **Server Groups** list. On the **Server Group Settings** page that appears, you can view the heartbeat status of the Logtail container. For more information, see [View the status of a server group.](#)

Create a Logtail configuration

Create a Logtail configuration in the console.

- For more information about Docker logs, see [Collect container text logs.](#)
- For more information about Docker standard output, see [Collect stdout and stderr logs from containers.](#)
- [Host text logs.](#)

The root directory of a host is mounted on the `/logtail_host` directory of the Logtail container by default. You must add the `/logtail_host` prefix to the log path. For example, if you want to collect data from the `/home/logs/app_log/` directory of the host, you must set the log path as `/logtail_host/home/logs/app_log/`.

What to do next

- View the status of the Logtail container.

You can run the `docker exec ${logtail_container_id} /etc/init.d/ilogtail status` command to view the status of Logtail.

- View the version number, IP address, and startup time of Logtail.

You can run the `docker exec ${logtail_container_id} cat /usr/local/ilogtail/app_info.json` command to view Logtail information.

- View the operations logs of Logtail.

The operations logs of Logtail are stored in the `ilogtail.LOG` file in the `/usr/local/ilogtail/` directory. If the log file is rotated and compressed, it is stored as a file named `ilogtail.LOG.x.gz`.

For example:

```
[root@iZbp17enxc2us3624wexh2Z ilogtail]# docker exec a287de895e40 tail -n 5 /usr/local/ilogtail/ilogtail.LOG
[2018-02-06 08:13:35.721864] [INFO] [8] [build/release64/sls/ilogtail/LogtailPlugin.cpp:104] logtail plugin Resume:
start
[2018-02-06 08:13:35.722135] [INFO] [8] [build/release64/sls/ilogtail/LogtailPlugin.cpp:106] logtail plugin Resume:
success
[2018-02-06 08:13:35.722149] [INFO] [8] [build/release64/sls/ilogtail/EventDispatcher.cpp:369] start add existed c
heck point events, size:0
[2018-02-06 08:13:35.722155] [INFO] [8] [build/release64/sls/ilogtail/EventDispatcher.cpp:511] add existed check
point events, size:0 cache size:0 event size:0 success count:0
[2018-02-06 08:13:39.725417] [INFO] [8] [build/release64/sls/ilogtail/ConfigManager.cpp:3776] check container pa
th update flag:0 size:1
```

Ignore the following standard output:

```
start umount useless mount points, /shm$/merged$/mqueue$
umount: /logtail_host/var/lib/docker/overlay2/3fd0043af174cb0273c3c7869500fbe2bdb95d13b1e110172ef57fe840c82
155/merged: must be superuser to unmount
umount: /logtail_host/var/lib/docker/overlay2/d5b10aa19399992755de1f85d25009528daa749c1bf8c16edff44beab6e6
9718/merged: must be superuser to unmount
umount: /logtail_host/var/lib/docker/overlay2/5c3125daddacedec29df72ad0c52fac800cd56c6e880dc4e8a640b1e16c2
2dbe/merged: must be superuser to unmount
...
xargs: umount: exited with status 255; aborting
umount done
start logtail
ilogtail is running
logtail status:
ilogtail is running
```

- Restart Logtail.

To restart Logtail, use the following sample code:

```
[root@iZbp17enxc2us3624wexh2Z ilogtail]# docker exec a287de895e40 /etc/init.d/ilogtaild stop
kill process Name: ilogtail pid: 7
kill process Name: ilogtail pid: 8
stop success
[root@iZbp17enxc2us3624wexh2Z ilogtail]# docker exec a287de895e40 /etc/init.d/ilogtaild start
ilogtail is running
```

23.3.1.6.2. Collect Kubernetes logs

This topic describes how to install and use Logtail to collect logs from Kubernetes clusters.

Configuration procedure

Perform the following steps to collect logs from Kubernetes clusters:

1. Install the alibaba-log-controller Helm package.
2. Use the Log Service console to manage log collection configurations.

Step 1: Install Logtail

- Install Logtail in an Alibaba Cloud Container Service for Kubernetes cluster.
 - If Log Service components are not installed in your cluster, you must manually install all the components.
 - i. Connect to the Kubernetes cluster by using CloudShell.
 - ii. Run the following command in CloudShell to obtain the ID of your Apsara Stack tenant account.

```
echo $ALIBABA_CLOUD_ACCOUNT_ID
```

- iii. After you set the `{your_k8s_cluster_id}`, `{your_ali_uid}`, and `{your_k8s_cluster_region_id}` parameters, run the following command:

```
wget https://acs-logging.oss-cn-hangzhou.aliyuncs.com/alibabacloud-k8s-log-installer.sh -O alibabacloud-k8s-log-installer.sh; chmod 744 ./alibabacloud-k8s-log-installer.sh; ./alibabacloud-k8s-log-installer.sh --cluster-id {your_k8s_cluster_id} --ali-uid {your_ali_uid} --region-id {your_k8s_cluster_region_id}
```

- Install Logtail in a user-created Kubernetes cluster.

Notice

- The version of the Kubernetes cluster must be 1.8 or later.
- Helm 2.6.4 or later must be installed.

- i. In the Log Service console, create a project whose name starts with `k8s-log-custom-`.
- ii. Replace the parameters in the following command based on your business requirements:

```
wget http://logtail-release-cn-hangzhou.oss-cn-hangzhou.aliyuncs.com/kubernetes/alibabacloud-log-k8s-custom-install.sh; chmod 744 ./alibabacloud-log-k8s-custom-install.sh; sh ./alibabacloud-log-k8s-custom-install.sh {your-project-suffix} {region-id} {aliuid} {access-key-id} {access-key-secret}
```

The following table lists the parameters in the preceding command.

| Parameter | Description |
|------------------------------------|---|
| <code>{your-project-suffix}</code> | The portion of the project name at the end of <code>k8s-log-custom-</code> . For example, if you create a project whose name is <code>k8s-log-custom-xxxx</code> , set this parameter to <code>xxxx</code> . |
| <code>{regionId}</code> | The ID of the region where the project resides. For more information, see View the information of a project . |
| <code>{aliuid}</code> | The user ID. Set this parameter to the ID of your Apsara Stack tenant account. Note The ID of an Apsara Stack tenant account is a string. For more information about how to obtain the ID, see Configure an account ID for a server . |
| <code>{access-key-id}</code> | The AccessKey ID of your Apsara Stack tenant account. |
| <code>{access-key-secret}</code> | The AccessKey secret of your Apsara Stack tenant account. |

After Logtail is installed in the Kubernetes cluster, Log Service automatically creates a machine group named `k8s-group-{your_k8s_cluster_id}` for the project.

Note

- A Logstore named `config-operation-log` is automatically created in the project. Do not delete the Logstore.
- When you install Logtail in a user-created Kubernetes cluster, Logtail is granted `privileged` permissions by default. This prevents the `container text file busy` error when you delete a pod. For more information, visit [Bug 1468249](#), [Bug 1441737](#), and [Issue 34538](#).

The following example shows a successful installation:

```
[root@iZbp1dsxxxxqfbiaZ ~]# wget http://logtail-release-cn-hangzhou.oss-cn-hangzhou.aliyuncs.com/kubernetes/al
icloud-log-k8s-custom-install.sh; chmod 744 ./alicloud-log-k8s-custom-install.sh; sh ./alicloud-log-k8s-custom-install.
sh xxxx cn-hangzhou 165xxxxxxxx050 LTAxxxxxxxx Alxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
....
....
....
NAME: alibaba-log-controller
LAST DEPLOYED: Fri May 18 16:52:38 2018
NAMESPACE: default
STATUS: DEPLOYED
RESOURCES:
==> v1beta1/ClusterRoleBinding
NAME          AGE
alibaba-log-controller 0s
==> v1beta1/DaemonSet
NAME    DESIRED CURRENT READY UP-TO-DATE AVAILABLE NODE SELECTOR AGE
logtail-ds 2    2    0 2    0    <none>    0s
==> v1beta1/Deployment
NAME          DESIRED CURRENT UP-TO-DATE AVAILABLE AGE
alibaba-log-controller 1    1    1    0    0s
==> v1/Pod(related)
NAME          READY STATUS    RESTARTS AGE
logtail-ds-7xf2d    0/1 ContainerCreating 0    0s
logtail-ds-9j4bx    0/1 ContainerCreating 0    0s
alibaba-log-controller-796f8496b6-6jxb2 0/1 ContainerCreating 0    0s
==> v1/ServiceAccount
NAME          SECRETS AGE
alibaba-log-controller 1    0s
==> v1beta1/CustomResourceDefinition
NAME          AGE
aliyunlogconfigs.log.alibabacloud.com 0s
==> v1beta1/ClusterRole
alibaba-log-controller 0s
[INFO] your k8s is using project : k8s-log-custom-xxx, region : cn-hangzhou, aliuid : *****, accessKeyId : LTA*
*****
[SUCCESS] install helm package : alibaba-log-controller success.
```

To check the status of each Log Service component in the Kubernetes cluster, run the `helm status alibaba-log-controller` command. If all pods are in the Running state, Logtail is installed.

Log on to the Log Service console to find the project. If you have multiple projects, search for the project by using the `k8s-log` keyword.

Step 2: Configure log collection

Create Logtail configurations for log collection in the console as required.

- For information about how to collect Kubernetes text logs, see [Collect container text logs](#).
- For information about how to collect Kubernetes stdout logs, see [Collect stdout and stderr logs from containers](#).
- [Host text logs](#).

By default, the root directory of a host is mounted to the `/logtail_host` directory of the Logtail container. You must add the `/logtail_host` prefix to the log path. For example, if you want to collect data from the `/home/logs/app_log/` directory of the host, you must set the log path to `/logtail_host/home/logs/app_log/`.

Other common commands

- Store logs of multiple clusters in one project

You can collect logs from multiple Kubernetes clusters. If you want to store these logs in the same project, you can specify the same cluster ID for the `${your_k8s_cluster_id}` parameter when you install Log Service components on multiple Kubernetes clusters.

For example, if you have three Kubernetes clusters whose IDs are abc001, abc002, and abc003, specify `abc001` for the `${your_k8s_cluster_id}` parameter when you install Log Service components for each Kubernetes cluster.

 **Notice** This feature is unavailable for Kubernetes clusters that reside in different regions.

- Logtail container logs

Logtail log files named `ilogtail.LOG` and `logtail_plugin.LOG` are stored in the `/usr/local/ilogtail/` directory of a Logtail container. Ignore the following standard output:

```
start umount useless mount points, /shm$/merged$/mqueue$
umount: /logtail_host/var/lib/docker/overlay2/3fd0043af174cb0273c3c7869500fbe2bdb95d13b1e110172ef57fe840c82
155/merged: must be superuser to unmount
umount: /logtail_host/var/lib/docker/overlay2/d5b10aa19399992755de1f85d25009528daa749c1bf8c16edff44beab6e6
9718/merged: must be superuser to unmount
umount: /logtail_host/var/lib/docker/overlay2/5c3125daddacedec29df72ad0c52fac800cd56c6e880dc4e8a640b1e16c2
2dbe/merged: must be superuser to unmount
.....
xargs: umount: exited with status 255; aborting
umount done
start logtail
ilogtail is running
logtail status:
ilogtail is running
```

- View the status of each Log Service component in a Kubernetes cluster

```
helm status alibaba-log-controller
```

- Troubleshoot alibaba-log-controller startup failures

Make sure that the following conditions are met:

- Log Service components are installed on the master node of the Kubernetes cluster.
- The Kubernetes cluster ID that you specified is valid when you install Log Service components.

If Log Service components fail to be installed because the preceding conditions are not met, run the `helm del --purge alibaba-log-controller` command to delete the installation package and install Log Service components again.

- View the status of Logtail DaemonSet in a Kubernetes cluster

Run the `kubectl get ds -n kube-system` command.

 **Note** The default namespace of Logtail is `kube-system`.

- View the version number, IP address, and startup time of Logtail.

Example:

```
[root@iZbp1dsu6v77zfb40qfbiaZ ~]# kubectl get po -n kube-system | grep logtail
NAME      READY   STATUS    RESTARTS   AGE
logtail-ds-gb92k 1/1   Running  0          2h
logtail-ds-wm7lw 1/1   Running  0          4d
[root@iZbp1dsu6v77zfb40qfbiaZ ~]# kubectl exec logtail-ds-gb92k -n kube-system cat /usr/local/ilogtail/app_info.json
{
  "UUID": "",
  "hostname": "logtail-ds-gb92k",
  "instance_id": "0EBB2B0E-0A3B-11E8-B0CE-0A58AC140402_172.20.4.2_1517810940",
  "ip": "172.20.4.2",
  "logtail_version": "0.16.2",
  "os": "Linux; 3.10.0-693.2.2.el7.x86_64; #1 SMP Tue Sep 12 22:26:13 UTC 2017; x86_64",
  "update_time": "2018-02-05 06:09:01"
}
```

- View the operational logs of Logtail

Logtail operational logs are stored in the `ilogtail.LOG` file in the `/usr/local/ilogtail/` directory. If the log file is rotated and compressed, it is stored as a file named `ilogtail.LOG.x.gz`.

Example:

```
[root@iZbp1dsu6v77zfb40qfbiaZ ~]# kubectl exec logtail-ds-gb92k -n kube-system tail /usr/local/ilogtail/ilogtail.LOG
[2018-02-05 06:09:02.168693] [INFO] [9] [build/release64/sls/ilogtail/LogtailPlugin.cpp:104] logtail plugin Resume:start
[2018-02-05 06:09:02.168807] [INFO] [9] [build/release64/sls/ilogtail/LogtailPlugin.cpp:106] logtail plugin Resume:success
[2018-02-05 06:09:02.168822] [INFO] [9] [build/release64/sls/ilogtail/EventDispatcher.cpp:369] start add existed check point events, size:0
[2018-02-05 06:09:02.168827] [INFO] [9] [build/release64/sls/ilogtail/EventDispatcher.cpp:511] add existed check point events, size:0 cache size:0 event size:0 success count:0
```

- Restart Logtail in a pod

Example:

```
[root@iZbp1dsu6v77zfb40qfbiaZ ~]# kubectl exec logtail-ds-gb92k -n kube-system /etc/init.d/ilogtaild stop
kill process Name: ilogtail pid: 7
kill process Name: ilogtail pid: 9
stop success
[root@iZbp1dsu6v77zfb40qfbiaZ ~]# kubectl exec logtail-ds-gb92k -n kube-system /etc/init.d/ilogtaild start
ilogtail is running
```

23.3.1.6.3. Collect container text logs

Logtail collects text logs generated in containers and uploads these logs together with the container-related metadata information to Log Service.

Features

Compared with basic log file collection, Docker file collection has the following characteristics:

- Allows you to configure the log path of a container, without the need to consider the mapping between the path and the host.
- Allows you to use labels to specify the containers whose logs are to be collected.
- Allows you to use labels to exclude specific containers.
- Allows you to use environment variables to specify the containers whose logs are to be collected.
- Allows you to use environment variables to exclude specific containers.
- Supports multiline logs such as Java Stack logs.

- Supports automatic labeling for container data.
- Supports automatic labeling for Kubernetes containers.

Limits

- **Collection stop policy:** When a container is stopped and Logtail detects the `die` event on the container, Logtail stops collecting logs of the container (with a latency of no more than 3 seconds). In this case, if a collection latency occurs, some logs generated before the stop action may be lost.
- **Docker storage driver:** Only overlay and overlay2 are supported. For other storage drivers, you must mount the log directory to the local host.
- **Logtail running mode:** Logtail must run in a container and must be deployed based on Logtail deployment solutions.
- **Label:** refers to the label information in docker inspect. It is not synonymous with labels in Kubernetes.
- **Environment:** refers to the environment information configured during container startup.

Procedure

1. Deploy and configure the Logtail container.
2. Configure log collection in Log Service.

Logtail deployment and configuration

- **Kubernetes**
For more information about Kubernetes log collection, see [Logtail deployment solution for collecting Kubernetes logs](#).
- **Management methods for other containers**
For more information about management methods for other containers, such as Swarm and Mesos, see [Common deployment solution for collecting Docker logs](#).

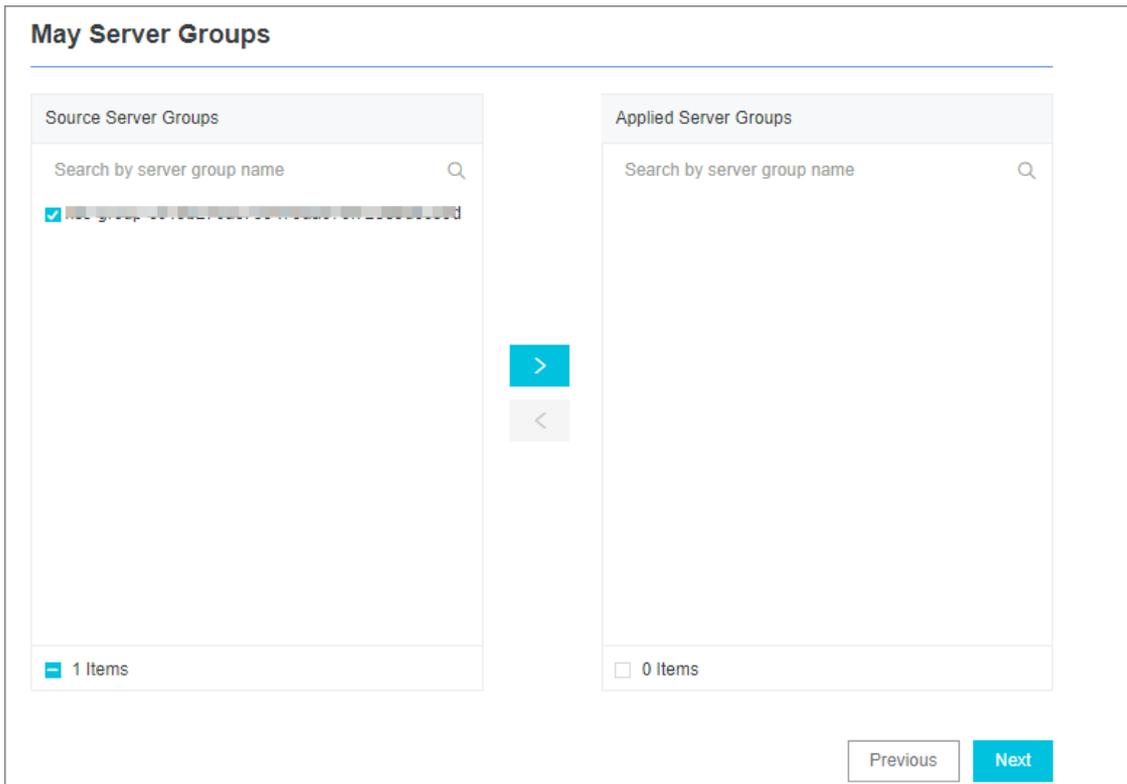
Collection configuration

1. [Log on to the Log Service console](#).
2. Click the **Import Data** button. On the **Import Data** page that appears, select **Docker File**.
3. Select a Logstore, and then click **Next**.

Select an existing project and Logstore. You can also click **Create Now** to create a project and Logstore. For more information, see [Manage a Logstore](#).

If you enter the configuration procedure by clicking the plus sign (+) next to **Data Import** under a Logstore on the Logstores tab, the system skips this step.

4. Create a server group.
Before you create a server group, ensure that Logtail is installed.
Install Logtail as prompted. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).
After you install Logtail, click **Complete Installation** to create a server group. For more information, see [Overview](#). If you have created a server group, click **Use Existing Server Groups** to select the server group.
5. Configure the server group, and then click **Next**.
Select a server group and move the group from **Source Server Groups** to **Applied Server Groups**.



6. Configure Logtail.

The following table lists data a source-specific parameters. For more information about common parameters, see [Configure text log collection](#).

| Parameter | Description |
|-----------------|---|
| Docker File | This parameter is used to check whether the collected target file is a Docker file. |
| Label Whitelist | <p>LabelKey is required. If LabelValue is not empty, only containers whose labels contain LabelKey = LabelValue are collected. If LabelValue is empty, all the containers whose labels contain LabelKey are collected.</p> <div style="background-color: #e0f2f1; padding: 10px; border: 1px solid #ccc;"> <p>Note</p> <ul style="list-style-type: none"> ◦ Key-value pairs are disjunctive with each other. If the label of a container contains one of the key-value pairs you specify, logs of the container are collected. ◦ Labels refer to Docker labels. </div> |
| Label Blacklist | <p>LabelKey is required. If LabelValue is not empty, only containers whose labels contain LabelKey = LabelValue are excluded. If LabelValue is empty, all containers whose labels contain LabelKey are excluded.</p> <div style="background-color: #e0f2f1; padding: 10px; border: 1px solid #ccc;"> <p>Note</p> <ul style="list-style-type: none"> ◦ Key-value pairs are disjunctive with each other. If the label of a container contains one of the key-value pairs you specify, the container is excluded. ◦ Labels described in this topic refer to the label information in docker inspect. </div> |

| Parameter | Description |
|--------------------------------|---|
| Environment Variable Whitelist | <p>EnvKey is required. If EnvValue is not empty, only containers whose environment variables contain EnvKey = EnvValue are collected. If EnvValue is empty, all containers whose environment variables contain EnvKey are collected.</p> <div style="background-color: #e0f2f1; padding: 10px;"> <p>Note</p> <ul style="list-style-type: none"> Key-value pairs are disjunctive with each other. If the environment variable of a container contains one of the key-value pairs you specify, logs of the container are collected. The environment variable refers to the environment information configured in container startup. </div> |
| Environment Variable Blacklist | <p>EnvKey is required. If EnvValue is not empty, only containers whose environment variables contain EnvKey = EnvValue are excluded. If EnvValue is empty, all containers whose environment variables contain EnvKey are excluded.</p> <div style="background-color: #e0f2f1; padding: 10px;"> <p>Note</p> <ul style="list-style-type: none"> Key-value pairs are disjunctive with each other. If the environment variable of a container contains one of the key-value pairs you specify, the container is excluded. The environment variable refers to the environment information configured in container startup. </div> |

Note

- Labels in whitelist and blacklist are different from those defined in Kubernetes. Labels in this topic refer to the label information in docker inspect.
- A namespace and a container name in Kubernetes can be mapped to Docker labels. LabelKey corresponding to a namespace is `io.kubernetes.pod.namespace`. LabelKey corresponding to a container name is `io.kubernetes.container.name`. For example, the namespace of the pod you created is `backend-prod` and the container name is `worker-server`. In this case, you can configure a whitelist label: `io.kubernetes.pod.namespace:backend-prod` or `io.kubernetes.container.name:worker-server`, so that only logs of the container are collected.
- In Kubernetes, we recommend that you only use the `io.kubernetes.pod.namespace` and `io.kubernetes.container.name` labels. In other cases, use an environment variable whitelist or blacklist.

7. Configure an index.

Configure an index based on your business requirements. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

Configuration examples

- Environment configuration

Collect the logs of the container whose environment variable is `NGINX_PORT_80_TCP_PORT=80` but not `POD_NAMESPACE=kube-system`. The log file path is `/var/log/nginx/access.log` and logs are parsed in simple mode.

Note The environment variable refers to the environment information configured in container startup.

Environment configuration example

```

"stdinOnce": false,
"Env": [
  "HTTP_SVC_SERVICE_PORT_HTTP=80",
  "LOG4J_APPENDER_DEMO_SPRING_BOOT_SVC_PORT=:8080",
  "LOG4J_APPENDER_DEMO_SPRING_BOOT_SVC_PORT_8080_TCP_PORT=8080",
  "HTTP_SVC_PORT_80_TCP_ADDR=",
  "NGINX_PORT_80_TCP=tcp://",
  "NGINX_PORT_80_TCP_PROTO=tcp",
  "LOG4J_APPENDER_DEMO_SPRING_BOOT_SVC_SERVICE_PORT=8080",
  "KUBERNETES_SERVICE_HOST=",
  "HTTP_SVC_SERVICE_HOST=",
  "HTTP_SVC_PORT_80_TCP_PROTO=tcp",
  "NGINX_PORT_80_TCP_ADDR=",
  "LOG4J_APPENDER_DEMO_SPRING_BOOT_SVC_PORT_8080_TCP_PROTO=tcp",
  "KUBERNETES_SERVICE_PORT_HTTPS=443",
  "KUBERNETES_PORT=tcp://:443",
  "NGINX_PORT=tcp://:80",
  "HTTP_SVC_PORT=tcp://:80",
  "HTTP_SVC_PORT_80_TCP_PORT=80",
  "NGINX_SERVICE_PORT=80",
  "KUBERNETES_PORT_443_TCP=tcp://:443",
  "KUBERNETES_PORT_443_TCP_PROTO=tcp",
  "HTTP_SVC_SERVICE_PORT=80",
  "KUBERNETES_PORT_443_TCP_ADDR=171.19.138.1",
  "HTTP_SVC_PORT_80_TCP=tcp://:80",

```

- Label configuration

Collect the logs of a container that meets the following conditions: the container label is `io.kubernetes.container.name=nginx`, the log file path is `/var/log/nginx/access.log`, and logs are parsed in simple mode.

? **Note** Labels refer to Docker labels.

Label configuration example

```

"OnBuild": null,
"Labels": {
  "annotation.io.kubernetes.container.hash": "53073f5a",
  "annotation.io.kubernetes.container.restartCount": "0",
  "annotation.io.kubernetes.container.terminationMessagePath": "/dev/termination-log",
  "annotation.io.kubernetes.container.terminationMessagePolicy": "File",
  "annotation.io.kubernetes.pod.terminationGracePeriod": "30",
  "io.kubernetes.container.logpath": "/var/log/pods/ad00a078-4182-4b1d-8000-585/nginx_0.log",
  "io.kubernetes.container.name": "nginx",
  "io.kubernetes.docker.type": "container",
  "io.kubernetes.pod.name": "example-foo-86ccd54874-r4mfh",
  "io.kubernetes.pod.namespace": "default",
  "io.kubernetes.pod.uid": "ad00a078-4182-4b1d-8000-585",
  "io.kubernetes.sandbox.id": "522643e0c6a70143da29dab0a78e0781396002040000001dfa6da112969",
  "maintainer": "NGINX Docker Maintainers <docker-maint@nginx.com>"
},
"StopSignal": "SIGTERM"

```

Default fields

Each uploaded log of a common Docker container contains the following fields.

| Field | Description: |
|-------------------------------|----------------------------|
| <code>_image_name_</code> | The name of the image. |
| <code>_container_name_</code> | The name of the container. |

| Field | Description: |
|-----------------------------|----------------------------------|
| <code>_container_ip_</code> | The IP address of the container. |

If a Kubernetes cluster is used, each uploaded log contains the following fields.

| Field | Description |
|-------------------------------|---|
| <code>_image_name_</code> | The name of the image. |
| <code>_container_name_</code> | The name of the container. |
| <code>_pod_name_</code> | The name of the pod. |
| <code>_namespace_</code> | The namespace to which the pod belongs. |
| <code>_pod_uid_</code> | The unique identifier of the pod. |
| <code>_container_ip_</code> | The IP address of the pod. |

23.3.1.6.4. Collect stdout and stderr logs from containers

This topic describes how to use Logtail to collect standard output (stdout) and standard error (stderr) logs from containers. After you collect stdout and stderr logs, you can upload the logs together with the container-related metadata to Log Service.

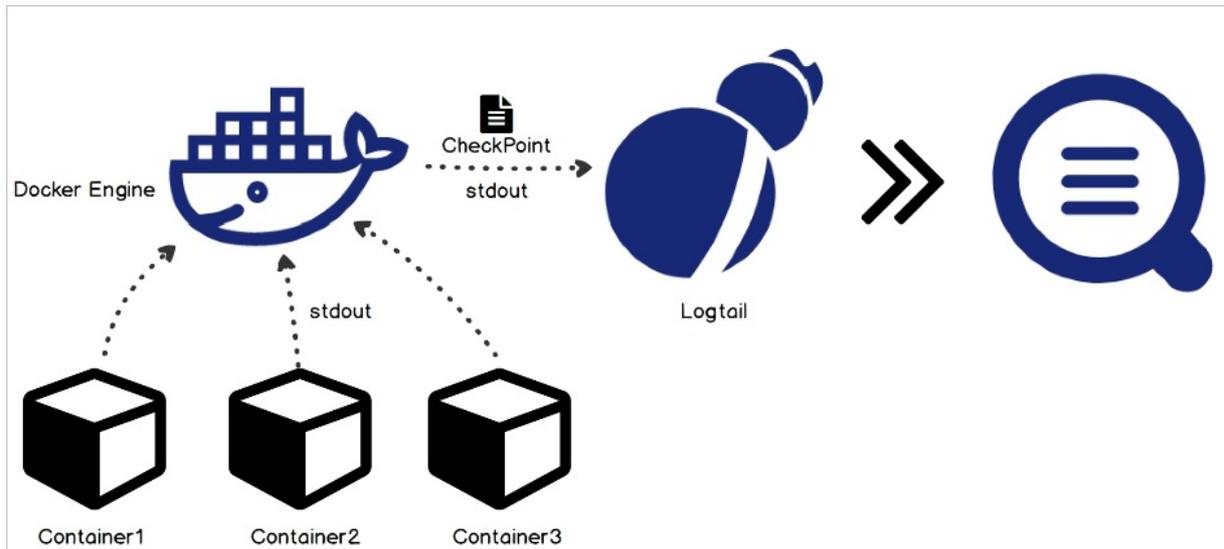
Features

- Collects stdout and xxx logs.
- Labels the containers from which you want to collect stdout and stderr logs.
- Uses tags to exclude specific containers.
- Allows you to use environment variables to specify the containers whose logs are to be collected.
- Allows you to use environment variables to exclude specific containers.
- Supports multiline logs such as Java Stack logs.
- Supports automatic labeling for container data.
- Supports automatic labeling for Kubernetes containers.

Implementation

As shown in the following figure, Logtail communicates with the domain socket of the Docker engine to query containers that run on the Docker engine. Logtail also locates containers from which you want to collect logs based on the specified labels and environment variables. Logtail then uses the docker logs command to collect logs from the specified containers.

When Logtail collects the stdout logs of a container, Logtail records information about log file positions to the checkpoint file at regular intervals. If Logtail is restarted, Logtail collects logs from the last log file position.



Limits

- This feature is available only for Logtail 0.16.0 or later that runs on Linux. For more information about Logtail versions and version updates, see [Install Logtail in Linux](#).
- The domain socket must exist and can access the Docker engine. Otherwise, Logtail cannot access the Docker engine by running the `/var/run/docker.sock` file.
- The last multiline log that you collect must be cached for at least 1,000 seconds. By default, the retention period for a multiline log is 3 seconds. You can specify the period by configuring the `BeginLineTimeoutMs` parameter. The value of the parameter must be a minimum of 1,000 ms. Otherwise, a false positive error may occur.
- Collection stop policy: When a container is stopped and Logtail detects the `die` event on the container, Logtail stops collecting stdout logs of the container. In this case, if a collection latency occurs, some stdout logs that are generated before the stop action may be lost.
- Docker log driver: Only log drivers of the JSON type are supported in the collection of stdout logs.
- Context: By default, a collection configuration is in the same context. If you need to configure different types of containers in different contexts, configure each type separately.
- Data processing: By default, collected logs start with the `content` field. You can apply standard processing methods to these logs. For more information about how to configure one or more processing methods, see [Configure data processing methods](#).
- Label: refers to the label information in `docker inspect`. It is not related to labels in Kubernetes configurations.
- Environment: refers to environment information that you specified during container startup.

Procedure

1. Deploy and configure Logtail on one or more containers.
2. Create a Logtail configuration and deliver it to Logtail.

Logtail deployment and configuration

- Kubernetes

For more information about how to collect Kubernetes logs, see [Logtail deployment solution for collecting Kubernetes logs](#).

- Configure Logtail on other containers

For more information about how to configure Logtail on other containers, such as Swarm and Mesos, see [Common Logtail deployment solution for collecting Docker logs](#).

Configure a data source

1. [Log on to the Log Service console](#).
2. Click **Import Data**. On the **Import Data** page that appears, select **Docker Standard Output**.
3. Select a Logstore, and then click **Next**.

Select an existing project and Logstore. You can also click **Create Now** to create a project and Logstore. For more information, see [Manage a Logstore](#).

If you enter the configuration procedure by clicking the plus sign (+) next to **Data Import** under a Logstore on the Logstores tab, the system skips this step.

4. Create a server group.

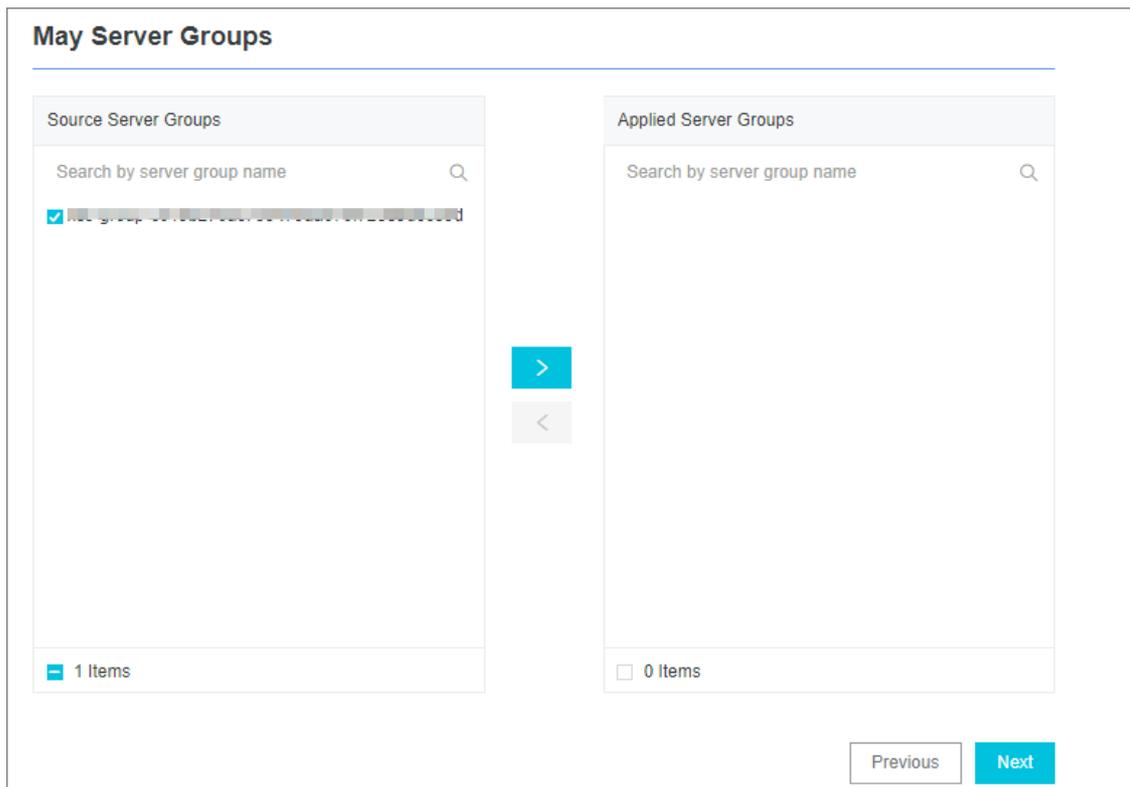
Before you create a server group, ensure that Logtail is installed.

Install Logtail as prompted. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).

After you install Logtail, click **Complete Installation** to create a server group. For more information, see [Overview](#). If you have created a server group, click **Use Existing Server Groups** to select the server group.

5. Configure the server group, and then click **Next**.

Select a server group and move the group from **Source Server Groups** to **Applied Server Groups**.



6. Configure a data source.

In the **Plug-in Config** section, set the required parameters. The following example shows how to set these parameters. For more information, see [Parameters](#).

```

{
  "inputs": [
    {
      "type": "service_docker_stdout",
      "detail": {
        "Stdout": true,
        "Stderr": true,
        "IncludeLabel": {
          "io.kubernetes.container.name": "nginx"
        },
        "ExcludeLabel": {
          "io.kubernetes.container.name": "nginx-ingress-controller"
        },
        "IncludeEnv": {
          "NGINX_SERVICE_PORT": "80"
        },
        "ExcludeEnv": {
          "POD_NAMESPACE": "kube-system"
        }
      }
    }
  ]
}

```

7. Configure an index.

Configure an index based on your business requirements. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

Parameters

The type of the input source is `service_docker_stdout`.

Note Before Logtail uploads data to Log Service, Logtail can process collected data. For more information about processing methods, see [Configure data processing methods](#).

| Parameter | Type | Required | Description |
|--------------|--|----------|---|
| IncludeLabel | JSON text. Key: JSON string. Value: JSON string. | Yes | <p>By default, this parameter is not specified. If the parameter is not specified, Logtail collects logs from all containers. If you set the key field and leave the value field empty, Logtail collects logs from containers whose labels include this key.</p> <p>Note</p> <ul style="list-style-type: none"> Key-value pairs are disjunctive with each other. If the label of a container matches one of the key-value pairs, logs of the container are collected. Labels refer to docker labels. |

| Parameter | Type | Required | Description |
|--------------|--|----------|---|
| ExcludeLabel | JSON text. Key: JSON string. Value: JSON string. | No | <p>This parameter is not specified by default. If the parameter is empty, no containers are excluded. If the key is not empty but the value is empty, the containers whose labels include this key are excluded.</p> <div style="background-color: #e1f5fe; padding: 10px;"> <p>? Note</p> <ul style="list-style-type: none"> Key-value pairs are disjunctive with each other. If the label of a container matches one of the key-value pairs, the container is excluded. Labels described in this topic refer to Docker labels. </div> |
| IncludeEnv | Map. Key: string. Value: string | No | <p>This parameter is empty by default. If the parameter is empty, logs of all containers are collected. If the key is not empty but the value is empty, logs of the containers whose environment variables include this key are collected.</p> <div style="background-color: #e1f5fe; padding: 10px;"> <p>? Note</p> <ul style="list-style-type: none"> Key-value pairs are disjunctive with each other. If the environment variable of a container includes one of the key-value pairs, the container is excluded. The environment variable refers to the environment information configured in container startup. </div> |
| ExcludeEnv | Map. Key: string. Value: string | No | <p>This parameter is empty by default. If you leave the parameter empty, no container is excluded. If you set the key and leave the value empty, the containers whose environment variables include this key are excluded.</p> <div style="background-color: #e1f5fe; padding: 10px;"> <p>? Note</p> <ul style="list-style-type: none"> Key-value pairs are disjunctive with each other. When the environment variable of a container includes one of the key-value pairs, the container is excluded. The environment variable refers to the environment information configured in container startup. </div> |
| Stdout | Boolean | No | If the value of the parameter is false, stdout data is not collected. Default value: true. |
| Stderr | Boolean | No | Default value: true. If the value of the parameter is false, stderr data is not collected. |

| Parameter | Type | Required | Description |
|----------------------|---------|----------|--|
| BeginLineRegex | String | No | This parameter is not specified by default. If the parameter is not empty, the regular expression is used to match the first line of each log. If a line matches this regular expression, this line is assumed as the start of a new log. Otherwise, this line is assumed as part of the previous log. |
| BeginLineTimeoutMs | Integer | No | The timeout period for the regular expression to match a line. Default value: 3000. Unit: ms. If no new log appears within 3 seconds, the most recent log is uploaded. |
| BeginLineCheckLength | Integer | No | The length of data for the regular expression to match. Default value: 10×1024. Unit: bytes. You can set this parameter to check whether the beginning part of a line can match the regular expression. This improves matching efficiency. |
| MaxLogSize | Integer | No | The maximum length of a log. Default value: 512×1024. Unit: bytes. If the length exceeds this value, the log data is uploaded directly without finding the first line of logs. |

Note

- Labels defined in IncludeLabel and ExcludeLabel are different from those defined in Kubernetes. Labels in this topic refer to Docker labels.
- A namespace and a container name in Kubernetes can be mapped to Docker labels. The LabelKey parameter corresponding to a namespace is `io.kubernetes.pod.namespace`. The LabelKey parameter corresponding to a container name is `io.kubernetes.container.name`. For example, the namespace of the pod you created is `backend-prod` and the container name is `worker-server`. In this case, you can configure a whitelist label: `io.kubernetes.pod.namespace:backend-prod` or `io.kubernetes.container.name:worker-server`, so that only logs of the container are collected.
- In Kubernetes, we recommend that you use the `io.kubernetes.pod.namespace` and `io.kubernetes.container.name` labels. In other cases, use IncludeEnv or ExcludeEnv.

Default fields

- Common Docker containers

Each uploaded log contains the following fields.

| Field | Description: |
|-------------------------------|---|
| <code>_time_</code> | The data upload time. Example: <code>2018-02-02T02:18:41.979147844Z</code> . |
| <code>_source_</code> | The type of the input source. Valid values: <code>stdout</code> and <code>stderr</code> . |
| <code>_image_name_</code> | The name of the image. |
| <code>_container_name_</code> | The name of the container. |
| <code>_container_ip_</code> | The IP address of the container. |

- Kubernetes containers

Each uploaded log contains the following fields.

| Field | Description |
|-------------------------------|---|
| <code>_time_</code> | The data upload time. Example: 2018-02-02T02:18:41.979147844Z . |
| <code>_source_</code> | The type of input sources. Valid values: stdout and stderr. |
| <code>_image_name_</code> | The name of the image. |
| <code>_container_name_</code> | The name of the container. |
| <code>_pod_name_</code> | The name of the pod. |
| <code>_namespace_</code> | The namespace where the pod is located. |
| <code>_pod_uid_</code> | The unique identifier of the pod. |
| <code>_container_id_</code> | The IP address of the pod. |

Common configuration examples

- Environment configuration

Collect the logs of the container whose environment variable is `NGINX_PORT_80_TCP_PORT=80` but not `POD_NAMESPACE=kube-system` .

Note The environment variable refers to the environment information configured in container startup.

Environment configuration example

```

"StdinOnce": false,
"Env": [
  "HTTP_SVC_SERVICE_PORT_HTTP=80",
  "LOG4J_APPENDER_DEMO_SPRING_BOOT_SVC_PORT=:8080",
  "LOG4J_APPENDER_DEMO_SPRING_BOOT_SVC_PORT_8080_TCP_PORT=8080",
  "HTTP_SVC_PORT_80_TCP_ADDR=",
  "NGINX_PORT_80_TCP=tcp://",
  "NGINX_PORT_80_TCP_PROTO=tcp",
  "LOG4J_APPENDER_DEMO_SPRING_BOOT_SVC_SERVICE_PORT=8080",
  "KUBERNETES_SERVICE_HOST=",
  "HTTP_SVC_SERVICE_HOST=",
  "HTTP_SVC_PORT_80_TCP_PROTO=tcp",
  "NGINX_PORT_80_TCP_ADDR=",
  "LOG4J_APPENDER_DEMO_SPRING_BOOT_SVC_PORT_8080_TCP_PROTO=tcp",
  "KUBERNETES_SERVICE_PORT_HTTPS=443",
  "KUBERNETES_PORT=tcp://:443",
  "NGINX_PORT=tcp://:80",
  "HTTP_SVC_PORT=tcp://:80",
  "HTTP_SVC_PORT_80_TCP_PORT=80",
  "NGINX_SERVICE_PORT=80",
  "KUBERNETES_PORT_443_TCP=tcp://:443",
  "KUBERNETES_PORT_443_TCP_PROTO=tcp",
  "HTTP_SVC_SERVICE_PORT=80",
  "KUBERNETES_PORT_443_TCP_ADDR=171.19.137.1",
  "HTTP_SVC_PORT_80_TCP=tcp://:80",

```

Collection configuration

```

{
  "inputs": [
    {
      "type": "service_docker_stdout",
      "detail": {
        "Stdout": true,
        "Stderr": true,
        "IncludeEnv": {
          "NGINX_PORT_80_TCP_PORT": "80"
        },
        "ExcludeEnv": {
          "POD_NAMESPACE": "kube-system"
        }
      }
    }
  ]
}

```

- Label configuration

Collect the stdout and stderr logs of the container whose label is `io.kubernetes.container.name=nginx` but not `type=pre`.

 **Note** Labels refer to Docker labels.

Label configuration example

```

"onBuild": null,
"Labels": {
  "annotation.io.kubernetes.container.hash": "53073f5a",
  "annotation.io.kubernetes.container.restartCount": "0",
  "annotation.io.kubernetes.container.terminationMessagePath": "/dev/termination-log",
  "annotation.io.kubernetes.container.terminationMessagePolicy": "File",
  "annotation.io.kubernetes.pod.terminationGracePeriod": "30",
  "io.kubernetes.container.logpath": "/var/log/pods/ad00a078-85/nginx_0.log",
  "io.kubernetes.container.name": "nginx",
  "io.kubernetes.docker.type": "container",
  "io.kubernetes.pod.name": "example-foo-86ccd54874-r4mfh",
  "io.kubernetes.pod.namespace": "default",
  "io.kubernetes.pod.uid": "ad00a07",
  "io.kubernetes.sandbox.id": "5216-a8d0b6891dfa6da112969",
  "maintainer": "NGINX Docker Maintainers <docker-maint@nginx.com>"
},
"StopSignal": "SIGTERM"

```

```

{
  "inputs": [
    {
      "type": "service_docker_stdout",
      "detail": {
        "Stdout": true,
        "Stderr": true,
        "IncludeLabel": {
          "io.kubernetes.container.name": "nginx"
        },
        "ExcludeLabel": {
          "type": "pre"
        }
      }
    }
  ]
}

```

Example of configuring multiline log collection

Configuring multiline log collection is important for the collection of Java exception stack logs. The following section introduces a standard collection configuration for Java stdout logs.

- Sample log

```
2018-02-03 14:18:41.968 INFO [spring-cloud-monitor] [nio-8080-exec-4] c.g.s.web.controller.DemoController : service start
2018-02-03 14:18:41.969 ERROR [spring-cloud-monitor] [nio-8080-exec-4] c.g.s.web.controller.DemoController : java.lang.NullPointerException
at org.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:193)
at org.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:166)
at org.apache.catalina.core.StandardWrapperValve.invoke(StandardWrapperValve.java:199)
at org.apache.catalina.core.StandardContextValve.invoke(StandardContextValve.java:96)
...
2018-02-03 14:18:41.968 INFO [spring-cloud-monitor] [nio-8080-exec-4] c.g.s.web.controller.DemoController : service start done
```

- Collection configuration

Collect logs of the container whose label is `app=monitor`. The first line of each log to be collected is of the date type. To improve matching efficiency, only the first 10 bytes of each line are checked.

```
{
  "inputs": [
    {
      "detail": {
        "BeginLineCheckLength": 10,
        "BeginLineRegex": "\\d+-\\d+-\\d+.*",
        "IncludeLabel": {
          "app": "monitor"
        }
      }
    },
    "type": "service_docker_stdout"
  ]
}
```

Process collected data

Logtail can process the collected Docker stdout logs by using a [common data processing method](#). Use a regular expression to extract the time, module, thread, class, and info fields.

- Collection configuration

Collect logs from a container whose label is `app=monitor`. The first line of each log to be collected is of the date type. To improve matching efficiency, only the first 10 bytes of each line are checked.

```
{
  "inputs": [
    {
      "detail": {
        "BeginLineCheckLength": 10,
        "BeginLineRegex": "\\d+-\\d+-\\d+.*",
        "IncludeLabel": {
          "app": "monitor"
        }
      },
      "type": "service_docker_stdout"
    }
  ],
  "processors": [
    {
      "type": "processor_regex",
      "detail": {
        "SourceKey": "content",
        "Regex": "\\d+-\\d+-\\d+ \\d+:\\d+:\\d+\\.\\.\\d+\\|s+(\\|w+\\|s+\\|[(\\^)]+\\|s+\\|[(\\^)]+\\|s+:\\|s+([\\|s\\|S]*)",
        "Keys": [
          "time",
          "module",
          "thread",
          "class",
          "info"
        ],
        "NoKeyError": true,
        "NoMatchError": true,
        "KeepSource": false
      }
    }
  ]
}
```

- Sample output

After the `2018-02-03 14:18:41.968 INFO [spring-cloud-monitor] [nio-8080-exec-4] c.g.s.web.controller.DemoController : service start done` log is processed, the following output is returned:

```
__tag__:__hostname__:logtail-dfgef
_container_name_:monitor
_image_name_:registry.cn-hangzhou.aliyuncs.aaaaaaaaaaaaaaaa
_namespace_:default
_pod_name_:monitor-6f54bd5d74-rtzc7
_pod_uid_:7f012b72-04c7-11e8-84aa-00163f00c369
_source_:stdout
_time_:2018-02-02T14:18:41.979147844Z
time:2018-02-02 02:18:41.968
level:INFO
module:spring-cloud-monitor
thread:nio-8080-exec-4
class:c.g.s.web.controller.DemoController
message:service start done
```

23.3.1.7. Limits

This topic describes the limits of Logtail. These limits apply when you collect files, manage resources, and resolve errors.

Limits on file collection

| Item | Description |
|--|--|
| File encoding | Log files encoded in UTF-8 and GBK are supported. We recommend that you use UTF-8 encoding for better processing performance. If log files are encoded in other formats, errors such as garbled characters and data loss may occur. |
| Log file size | Unlimited. |
| Log file rotation | Supported. Both <code>.log*</code> and <code>.log</code> are supported for file names. |
| Log collection behavior when log parsing is restricted | When log parsing is restricted, Logtail keeps the log file descriptor (FD) open. If log file rotation occurs multiple times during the restriction, Logtail attempts to keep the log parsing sequence of each rotation. If the number of rotated logs to be parsed exceeds 20, Logtail does not process subsequent log files. |
| Symbolic link | Monitored directories can be soft links. |
| Size of a single log entry | The maximum size of a single log entry is 512 KB. If a multi-line log entry is divided by using a regular expression to match the first line, the maximum size of each log entry after division is still 512 KB. If the size of a log entry exceeds 512 KB, the log entry is forcibly separated into multiple parts for collection. For example, if the size of a log entry is 1,025 KB, it will be split into three parts: 512 KB, 512 KB, and 1 KB. These log parts are collected in sequence. |
| Regular expression | Perl-based regular expressions can be used. |
| Multiple Logtail configuration files for the same log file | Not supported. We recommend that you collect and store log files to one Logstore, and configure multiple subscriptions. If this feature is required, configure symbolic links for log files to bypass this limit. |
| File opening behavior | When Logtail collects a log file, Logtail opens the log file. If the log file is not modified for more than 5 minutes and log rotation does not occur, Logtail closes the log file. |
| First log collection behavior | Logtail collects only incremental log files. If a log file is modified for the first time and the log file size exceeds 1 MB, Logtail collects the logs from the last 1 MB. Otherwise, Logtail collects the logs from the beginning of the log file. If the log file is not modified after the Logtail configuration file is sent to the server where the log file resides, Logtail does not collect the log file. |
| Non-standard text logs | If a log entry contains <code>\0</code> in multiple lines, the log entry is truncated at the first <code>\0</code> . |

Limits on checkpoints

| Item | Description |
|---------------------------|---|
| Checkpoint timeout period | If a log file is not modified for more than 30 days, the checkpoint of the log file is deleted. |
| Checkpoint storage policy | Checkpoints are saved every 15 minutes and are automatically saved when you exit Logtail. |
| Checkpoint storage path | By default, checkpoints are stored in the <code>/tmp/logtail_checkpoint</code> directory. For more information about how to modify the values of the related parameters, see Set Logtail startup parameters . |

Limits on configurations

| Item | Description |
|--|--|
| Configuration update | A custom configuration update requires about 30 seconds to take effect. |
| Dynamic loading of Logtail configuration files | Supported. The update of a Logtail configuration file does not affect other Logtail configuration files. |
| Number of Logtail configuration files | Unlimited. However, we recommend that you create a maximum of 100 Logtail configuration files on a server. |
| Multi-tenant isolation | Logtail configuration files for different tenants are isolated. |

Limits on resources and performance metrics

| Item | Description |
|---|---|
| Throughput for log processing | The default traffic of raw logs is limited to 2 MB/s. Data is uploaded after it is encoded and compressed. The compression ratio ranges from 5:1 to 10:1. Logs may be lost if the traffic exceeds the limit. For more information about how to modify the values of the related parameters, see Set Logtail startup parameters . |
| Maximum processing speed | Single-core processing speed: The maximum processing speed is 100 MB/s for logs in simple mode, 40 MB/s for logs in delimiter mode, and 30 MB/s for logs in JSON mode. By default, the maximum processing speed is 20 MB/s for logs in full regex mode based on the complexity of regular expressions. If multiple processing threads are enabled, the performance can be improved by 1.5 to 3 times. |
| Number of monitored directories | Logtail limits the depth of monitored directories to reduce the consumption of your resources. If the upper limit is reached, Logtail stops monitoring additional directories and log files. Logtail monitors a maximum of 3,000 directories, including subdirectories. |
| Number of monitored files | A Logtail configuration file on each server can be used to monitor a maximum of 10,000 files by default. A Logtail client on each server can monitor a maximum of 100,000 files by default. Excessive files are not monitored. If the upper limit is reached, you can perform the following operations: <ul style="list-style-type: none"> • Improve the depth of the monitored directory in each Logtail configuration file. • Modify the value of the <code>mem_usage_limit</code> parameter to increase the Logtail memory usage threshold. For more information, see Set Logtail startup parameters. <p>You can set a memory usage threshold of 2 GB for Logtail. In this case, the maximum number of files that each Logtail configuration file can be used to monitor is increased to 100,000. The maximum number of files that each Logtail client can monitor is increased to 1,000,000.</p> |
| Default resources | By default, Logtail consumes up to 40% of CPU usage and 256 MB of memory. If logs are generated at a high speed, you can modify relevant parameters. For more information, see Set Logtail startup parameters . |
| Processing policy of threshold-crossing resources | If the resources occupied by Logtail exceed the upper limit and this issue lasts for five or more minutes, Logtail is forcibly restarted. The restart may cause data loss or duplication. |

Limits on error handling

| Item | Description |
|---|--|
| Network error handling | If a network error occurs, Logtail retries and adjusts the retry interval. |
| Processing policy of threshold-crossing resources | If the data transmission speed exceeds the quota of the Logstore, Logtail restricts the log collection speed and retries the log collection. |
| Maximum retry period before timeout | If data fails to be transmitted and the failure lasts for more than six consecutive hours, Logtail discards the data. |
| Status self-check | Logtail restarts if an exception occurs, for example, an application unexpectedly exits or the resource usage exceeds the quota. |

Other limits

| Item | Description |
|------------------------|---|
| Log collection latency | A latency of less than 1 second exists between the time when a log is written to a disk and the time when Logtail collects the log. However, if the log collection speed is restricted, the latency increases. |
| Log upload policy | Before Logtail uploads logs, it aggregates the logs in the same file. The log upload starts if the number of logs exceeds 2,000, the total size of logs exceeds 2 MB, or the log collection duration exceeds 3 seconds. |

23.3.2. Other collection methods

23.3.2.1. WebTracking

This topic describes how to use WebTracking to collect logs from websites that are written in HTML or HTML5, iOS, and Android.

Context

Log Service uses WebTracking to collect logs from websites written in HTML or HTML5, iOS, and Android. You can customize dimensions and metrics.



In the preceding figure, WebTracking allows you to collect user information from various browsers, iOS apps, and Android apps (except for SDK for iOS or Android). For example:

Use SDK for Java to enable WebTracking

Example:

```
import com.aliyun.openservices.log.Client;
import com.aliyun.openservices.log.common.LogStore;
import com.aliyun.openservices.log.exception.LogException;
public class WebTracking {
    static private String accessId = "your accesskey id";
    static private String accessKey = "your accesskey";
    static private String project = "your project";
    static private String host = "log service data address";
    static private String logStore = "your logstore";
    static private Client client = new Client(host, accessId, accessKey);
    public static void main(String[] args) {
        try {
            //Enable WebTracking on an existing Logstore.
            LogStore logSt = client.GetLogStore(project, logStore).GetLogStore();
            client.UpdateLogStore(project, new LogStore(logStore, logSt.GetTtl(), logSt.GetShardCount(), true));
            //Disable WebTracking.
            //client.UpdateLogStore(project, new LogStore(logStore, logSt.GetTtl(), logSt.GetShardCount(), false));
            //Create a Logstore on which you want to enable WebTracking.
            //client.UpdateLogStore(project, new LogStore(logStore, 1, 1, true));
        }
        catch(LogException e){
            e.printStackTrace();
        }
    }
}
```

Step 2: Collect logs

After you enable WebTracking for a Logstore, you can use the following methods to upload data to the Logstore.

- Use the JavaScript SDK
 - i. Copy the *loghub-tracking.js* file to the *web* directory and add the following script to the file.
Click [here](#) to copy the script.

```
<script type="text/javascript" src="loghub-tracking.js" async></script>
```

 **Note** To ensure that a page loads, the script asynchronously sends HTTP requests. If data must be sent several times during the page loading, the newest request overwrites the previous HTTP request. A message showing WebTracking is about to exit appears in the browser. To eliminate the issue, send requests in a synchronous manner. To implement the method, perform the following step.

Original statement:

```
this.httpRequest_.open("GET", url, true)
```

Replace the original statement with the following statement:

```
this.httpRequest_.open("GET", url, false)
```

- ii. Create a tracker.

```
var logger = new window.Tracker('${host}','${project}','${logstore}');
logger.push('customer', 'zhangsan');
logger.push('product', 'iphone 6s');
logger.push('price', 5500);
logger.logger();
logger.push('customer', 'lisi');
logger.push('product', 'ipod');
logger.push('price', 3000);
logger.logger();
```

The following table lists the parameters:

| Parameter | Description |
|---------------------------|--|
| <code>\${host}</code> | The endpoint of the region where Log Service resides. For more information, see the <i>Obtain an endpoint topic in the Log Service Developer Guide</i> . |
| <code>\${project}</code> | The name of the project that you create in Log Service. |
| <code>\${logstore}</code> | The name of the Logstore in the <code>\${project}</code> . |

After you run the following code, the following logs appear in Log Service.

```
customer:zhangsan
product:iphone 6s
price:5500
```

```
customer:lisi
product:ipod
price:3000
```

- Use HTTP GET requests

```
curl --request GET 'http://${project}.${host}/logstores/${logstore}/track? APIVersion=0.6.0&key1=val1&key2=val2'
```

The following table lists the parameters.

| Parameter | Description |
|-----------------------------------|---|
| <code>\${project}</code> | The name of the project that you create in Log Service. |
| <code>\${host}</code> | The endpoint of the region where Log Service resides. |
| <code>\${logstore}</code> | The name of a Logstore that has WebTracking enabled in the <code>\${project}</code> . |
| <code>APIVersion=0.6.0</code> | (Required) A reserved parameter. |
| <code>__topic__=yourtopic</code> | (Optional) A reserved parameter that specifies the topic of the log. |
| <code>key1=val1, key2=val2</code> | The key-value pairs that you want to upload to Log Service. You can specify multiple pairs. Make sure that the length of each request URL is less than 16 KB. |

- Use HTML img tags

```
<img src='http://${project}.${host}/logstores/${logstore}/track.gif? APIVersion=0.6.0&key1=val1&key2=val2' />
<img src='http://${project}.${host}/logstores/${logstore}/track_ua.gif? APIVersion=0.6.0&key1=val1&key2=val2' />
```

The parameters that you need to specify are the same as the preceding parameters. In addition to custom parameters that are uploaded by track_ua.gif, Log Service also uses UserAgent and referer fields in the HTTP header as the fields of logs.

23.3.2.2. Use SDKs to collect logs

23.3.2.2.1. Producer Library

The Aliyun LOG Java Producer supports Java applications that run in big data processing scenarios with high concurrency. The library is easy to use and highly customizable.

For more information about the related GitHub project, visit [Aliyun LOG Java Producer](#).

23.3.2.2.2. Log4j Appender

Log4j is an open-source logging framework of Apache. You can use Log4j to write logs to the Log Service console, files, graphical user interface (GUI) components, socket servers, NT kernel loggers, or Unix syslog daemons. You can specify the output format of each log. You can also specify the severity level of each log to implement a fine-grained control on log generation.

Log4j consists of three components: loggers, appenders, and layouts.

- Loggers allow you to specify the severity level of each log.
Severity levels are sorted into ERROR, WARN, INFO, and DEBUG in descending order of severity.
- Appenders allow you to specify the destination of each log.
A destination can be the Log Service console or a file.
- Layouts allow you to specify the output format of each log.
The output format defines how logs are displayed.

To write logs to Log Service, use the Alibaba Cloud Log Log4j Appender. For information about where to download the library and how to use it, see [Log4j Appender](#).

23.3.2.2.3. Logback Appender

Logback was created by the same developer of Log4j. Logback allows you to write logs to multiple destinations. These destinations include the Log Service console, files, graphical user interface (GUI) components, socket servers, NT kernel loggers, and Unix syslog daemons. You can define the output format of each log. If you define the severity level of each log, you can implement a fine-grained control on the log generation process.

You can set the destination of logs to Log Service by using the Aliyun Log Logback Appender. The following example shows the format of logs that are uploaded to Log Service.

```
level: ERROR
location: com.aliyun.openservices.log.logback.example.LogbackAppenderExample.main(LogbackAppenderExample.java:18)
message: error log
throwable: java.lang.RuntimeException: xxx
thread: main
time: 2018-01-02T03:15+0000
log: 2018-01-02 11:15:29,682 ERROR [main] com.aliyun.openservices.log.logback.example.LogbackAppenderExample: error log
__source__: xxx
__topic__: yyy
```

For information about where to download the library and how to use it, see [Logback Appender](#).

23.3.2.2.4. Golang Producer Library

The Aliyun LOG Go Producer Library supports Go applications that run in big data processing scenarios with high concurrency. The library is easy to use and highly customizable. You can use the library to create producers that allow you to resend failed logs. Before Go applications send log data to Log Service, you can use these producers to compress the log data. This improves write performance.

For more information about the related GitHub project, visit [Aliyun Log Go Producer](#).

23.3.2.2.5. Python logging

Configurations

For more information about the configurations that are related to the Python logging module, see [Logging configuration](#).

The Python logging module allows you to configure logging by using code or a configuration file. The following example shows how to configure logging by using the `logging.conf` configuration file.

```
[loggers]
keys=root,sls
[handlers]
keys=consoleHandler, slsHandler
[formatters]
keys=simpleFormatter, rawFormatter
[logger_root]
level=DEBUG
handlers=consoleHandler
[logger_sls]
level=INFO
handlers=consoleHandler, slsHandler
qualname=sls
propagate=0
[handler_consoleHandler]
class=StreamHandler
level=DEBUG
formatter=simpleFormatter
args=(sys.stdout,)
[handler_slsHandler]
class=aliyun.log.QueuedLogHandler
level=INFO
formatter=rawFormatter
args=(os.environ.get('ALIYUN_LOG_SAMPLE_ENDPOINT', ''), os.environ.get('ALIYUN_LOG_SAMPLE_ACCESSID', ''), os.en
viron.get('ALIYUN_LOG_SAMPLE_ACCESSKEY', ''), os.environ.get('ALIYUN_LOG_SAMPLE_TMP_PROJECT', ''), "logstore")
[formatter_simpleFormatter]
format=%(asctime)s - %(name)s - %(levelname)s - %(message)s
[formatter_rawFormatter]
format=%(message)s
```

Two handlers named `root` and `sls` are created. The `sls` handler is an object of the `aliyun.log.QueuedLogHandler` class. The following shows the parameters that are specified for the `sls` handler. For more information, see [Parameters](#).

```
args=(os.environ.get('ALIYUN_LOG_SAMPLE_ENDPOINT', ''), os.environ.get('ALIYUN_LOG_SAMPLE_ACCESSID', ''), os.en
viron.get('ALIYUN_LOG_SAMPLE_ACCESSKEY', ''), os.environ.get('ALIYUN_LOG_SAMPLE_TMP_PROJECT', ''), "logstore")
```

 **Note** In this case, the `os.environ` function is used to retrieve configurations from environment variables. You can also specify values for these parameters based on your business requirements.

Upload logs

You can use the configuration file to upload logs to Log Service.

```
import logging
import logging.config
# Configurations
logging.config.fileConfig('logging.conf')
logger = logging.getLogger('sls')
# Use the logger
logger.info("test1")
try:
    1/0
except ZeroDivisionError as ex:
    logger.exception(ex)
```

Then, logs are automatically uploaded to Log Service. To use the LogSearch/Analytics feature, you must enable the index feature on the corresponding Logstore.

Configure an index for a Logstore

Enable the index feature on the Logstore that receives logs and configure an index for specific fields. We recommend that you use CLI (Command Line Interface) to configure the index as follows:

```
aliyunlog log update_index --project_name="project1" --logstore_name="logstore1" --index_detail="file:///Users/user1/loghandler_index.json"
```

For more information, see the [python_logging_handler_index.json](#) configuration file.

Specify log fields to be collected

The following table lists supported log fields that you can collect. By default, all of the fields are collected.

| Field | Description |
|-------------|---|
| message | The contents of a log. |
| record_name | The name of a handler. In the preceding example, the name is <code>sls</code> . |
| level | The output level of a logger, such as INFO and ERROR. |
| file_path | The full path of a configuration file. |
| func_name | The name of a function. |
| line_no | The number of a log line. |
| module | The name of a module where the function resides. |
| thread_id | The ID of the thread that runs the function. |
| thread_name | The name of the thread that runs the function. |
| process_id | The ID of the process that runs the function. |

| Field | Description |
|--------------|---|
| process_name | The name of the process that runs the function. |

You can specify log fields to be collected based on the `fields` parameter of the `QueuedLogHandler` class. For more information, see [aliyun.log.LogFields](#).

The following example shows how to modify the preceding configuration file and collect several fields, such as module and func_name.

```
[handler_slsHandler]
class=aliyun.log.QueuedLogHandler
level=INFO
formatter=rawFormatter
args=('cn-beijing.log.aliyuncs.com', 'ak_id', 'ak_key', 'project1', "logstore1", 'mytopic', ['level', 'func_name', 'module', 'line_no'] )
```

Note

- The message field is collected regardless of your configurations.
- To add a prefix and suffix to the names of these fields, use the `buildin_fields_prefix` and `buildin_fields_suffix` parameters. For example, `__level__`.

Configure logging by using a JSON text

You can use a JSON text to create more flexible logging configurations than code does.

```
#encoding: utf8
import logging, logging.config, os
# Configurations
conf = {'version': 1,
       'formatters': {'rawformatter': {'class': 'logging.Formatter',
                                       'format': '%(message)s'}},
       'handlers': {'sls_handler': {'():
                                   'aliyun.log.QueuedLogHandler',
                                   'level': 'INFO',
                                   'formatter': 'rawformatter',
                                   # custom args:
                                   'end_point': os.environ.get('ALIYUN_LOG_SAMPLE_ENDPOINT', ''),
                                   'access_key_id': os.environ.get('ALIYUN_LOG_SAMPLE_ACCESSID', ''),
                                   'access_key': os.environ.get('ALIYUN_LOG_SAMPLE_ACCESSKEY', ''),
                                   'project': 'project1',
                                   'log_store': "logstore1"
                                   }},
       'loggers': {'sls': {'handlers': ['sls_handler'],
                           'level': 'INFO',
                           'propagate': False}
                  }
              }
logging.config.dictConfig(conf)
# Use the logger
logger = logging.getLogger('sls')
logger.info("Hello world")
```

 **Note** To instantiate an object of the `aliyun.log.QueuedLogHandler` class, pass named parameters to the constructor. For more information, see [Parameters](#).

23.3.2.3. Common log formats

23.3.2.3.1. Log4j logs

Log Service allows you to collect Log4j logs.

Collect Log4j logs by using LogHub Log4j Appender

For more information, see [Log4j Appender](#).

Configure Logtail to collect Log4j logs

This topic describes how to configure regular expressions based on the default configuration of Log4j 1 logs. If Log4j 2 is used, you must modify the default configuration to record the complete date information. Log4j logs are sorted into Log4j 1 logs and Log4j 2 logs.

```
<Configuration status="WARN">
  <Appenders>
    <Console name="Console" target="SYSTEM_OUT">
      <PatternLayout pattern="%d{yyyy-MM-dd HH:mm:ss:SSS zzz} [%t] %-5level %logger{36} - %msg%n"/>
    </Console>
  </Appenders>
  <Loggers>
    <Logger name="com.foo.Bar" level="trace">
      <AppenderRef ref="Console"/>
    </Logger>
    <Root level="error">
      <AppenderRef ref="Console"/>
    </Root>
  </Loggers>
</Configuration>
```

For more information about how to configure Logtail to collect Log4j logs, see [Python logs](#). Configure the required parameters based on your network environment and business requirements.

The automatically generated regular expression is based on the sample log and may not be suitable for other logs. Therefore, you must make minor changes to the regular expression after it is automatically generated.

The following shows a sample log of the default Log4j format.

```
2013-12-25 19:57:06,954 [10.207.37.161] WARN impl.PermanentTairDaoImpl - Fail to Read Permanent Tair,key:e:470217319319741_1,result:com.example.tair.Result@172e3ebc[rc=code=-1, msg=connection error or timeout,value=,flag=0]
```

Regular expression that matches IP addresses that each indicate the start of a line:

```
\d+-\d+-\d+\. *
```

Regular expression used to extract log information:

```
(\d+-\d+-\d+\s\d+:\d+:\d+,\d+)\s[[([^\]]*)\]]\s(\S+)\s+(\S+)\s-\s(\. *)
```

Time conversion format:

```
%Y-%m-%d %H:%M:%S
```

The following table lists the extraction results of the sample log.

| Key | Value |
|---------|---|
| time | 2013-12-25 19:57:06,954 |
| ip | 10.207.37.161 |
| level | WARN |
| class | impl.PermanentTairDaoImpl |
| message | Fail to Read Permanent Tair,key:e:470217319319741_1,result:com.example.tair.Result@172e3ebc[rc=code=-1,msg=connection error or timeout,value=,flag=0] |

23.3.2.3.2. Python logs

The Python logging module provides a general logging system, which can be used by third-party modules or applications.

The Python logging module provides different log levels and logging methods, such as file-based, HTTP GET, HTTP POST, SMTP, and Socket logging. You can also create a custom logging method. The Python logging module works in the same way as the Log4j logging module except for some implementation details. The Python logging module includes the logger, handler, filter, and formatter objects.

Log format

A formatter specifies the output format of logs. To instantiate a formatter, pass two parameters to the constructor. One parameter includes a message format string and the other parameter includes a date format string. The parameters are optional.

Log format:

```
import logging
import logging.handlers
LOG_FILE = 'tst.log'
handler = logging.handlers.RotatingFileHandler(LOG_FILE, maxBytes = 1024*1024, backupCount = 5) # Instantiate the handler
fmt = '%(asctime)s - %(filename)s:%(lineno)s - %(name)s - %(message)s'
formatter = logging.Formatter(fmt) # Instantiate the formatter
handler.setFormatter(formatter) # Add the formatter to the handler
logger = logging.getLogger('tst') # Obtain a logger named tst
logger.addHandler(handler) # Add the handler to the logger
logger.setLevel(logging.DEBUG)
logger.info('first info message')
logger.debug('first debug message')
```

Attributes

Formatter attributes are specified in the `%(key)s` format. The following table lists the attributes.

| Format | Description |
|--------------------------|---|
| <code>%(name)s</code> | The name of a logger that generates logs. |
| <code>%(levelno)s</code> | The log output level in the numeric format. Valid values: DEBUG, INFO, WARNING, ERROR, and CRITICAL |

| Format | Description |
|---------------------|---|
| %(levelname)s | The log output level in the text format. Valid values: 'DEBUG', 'INFO', 'WARNING', 'ERROR', and 'CRITICAL'. |
| %(pathname)s | The full path of a source file that contains the logging module. |
| %(filename)s | The name of the source file. |
| %(module)s | The name of a module where the statement that you use to generate logs resides. |
| %(funcName)s | The name of the function that is used to call the log output function. |
| %(lineno)d | The number of a code line that contains the statement used to call the log output function. |
| %(created)f | The time when the log was created. The value is a UNIX timestamp representing the number of seconds that have elapsed since January 1, 1970, 00:00:00 (UTC). |
| %(relativeCreated)d | The interval between the time when a log was created and the time when the logging module was loaded. Unit: milliseconds. |
| %(asctime)s | The time when the log was created. The value of 2003-07-08 16:49:45,896 is an example of the default format. The number after the comma (,) indicates the number of milliseconds. |
| %(msecs)d | The time when the log was created. The value is a UNIX timestamp representing the number of milliseconds that have elapsed since January 1, 1970, 00:00:00 (UTC). |
| %(thread)d | The thread ID. |
| %(threadName)s | The thread name. |
| %(process)d | The process ID. |
| %(message)s | The contents of a log. |

Sample logs

Sample log:

```
2015-03-04 23:21:59,682 - log_test.py:16 - tst - first info message
2015-03-04 23:21:59,682 - log_test.py:17 - tst - first debug message
```

Common Python logs and the corresponding regular expressions:

- Sample log:

```
2016-02-19 11:03:13,410 - test.py:19 - tst - first debug message
```

Regular expression:

```
(\d+-\d+-\d+\s\S+)\s+-\s+([\^:]+):(\d+)\s+-\s+(\w+)\s+-\s+(. *)
```

- Log format

```
%(asctime)s - %(filename)s:%(lineno)s - %(levelname)s %(levelname)s %(pathname)s %(module)s %(funcName)s %(created)f %(thread)d %(threadName)s %(process)d %(name)s - %(message)s
```

Sample log:

```
2016-02-19 11:06:52,514 - test.py:19 - 10 DEBUG test.py test <module> 1455851212.514271 139865996687072 MainThread 20193 tst - first debug message
```

Regular expression:

```
(\d+-\d+-\d+\s\S+)\s-\s([\^:]+):(\d+)\s+-\s+(\d+)\s+(\w+)\s+(\S+)\s+(\w+)\s+(\S+)\s+(\S+)\s+(\d+)\s+(\w+)\s+(\d+)\s+(\w+)\s+-\s+(. *)
```

Configure Logtail to collect Python logs

1. [Log on to the Log Service console](#).
2. Click **Import Data**. On the **Import Data** page that appears, select **RegEx-Text Log**.
3. Select a Logstore, and then click **Next**.

Select an existing project and Logstore. You can also click **Create Now** to create a project and Logstore. For more information, see [Manage a Logstore](#).

If you enter the configuration procedure by clicking the plus sign (+) next to **Data Import** under a Logstore on the Logstores tab, the system skips this step.

4. Create a server group.

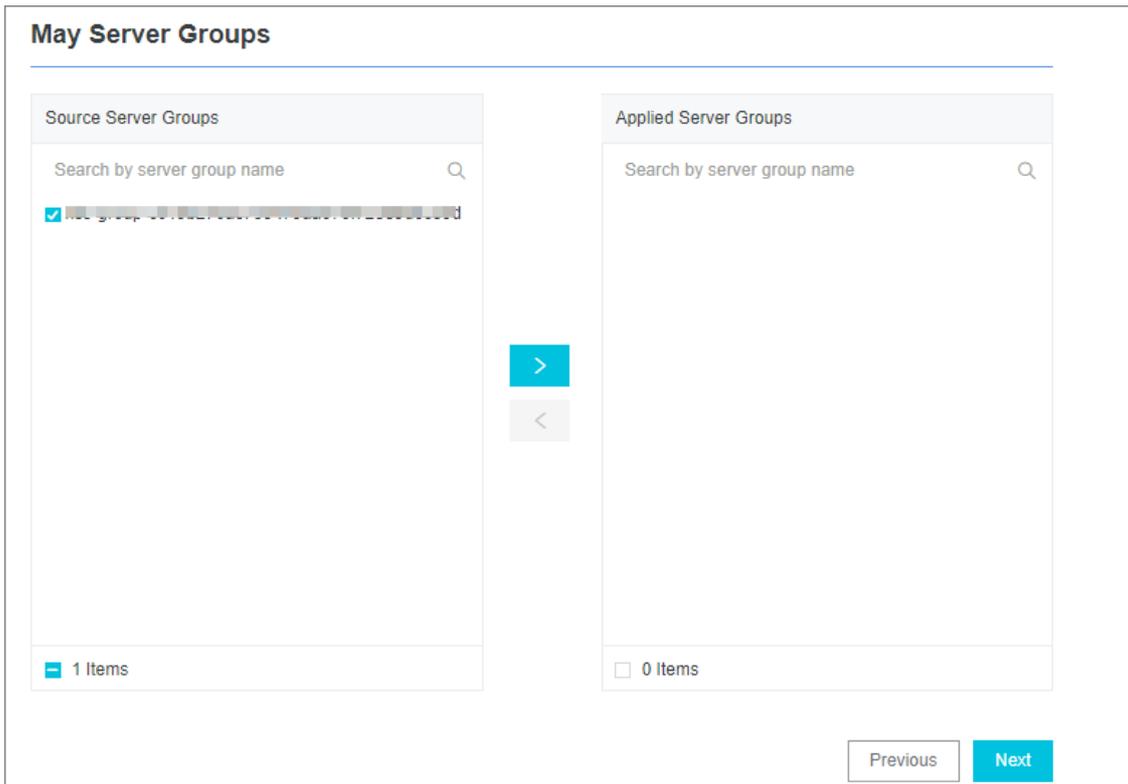
Before you create a server group, ensure that Logtail is installed.

Install Logtail as prompted. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).

After you install Logtail, click **Complete Installation** to create a server group. For more information, see [Overview](#). If you have created a server group, click **Use Existing Server Groups** to select the server group.

5. Configure the server group, and then click **Next**.

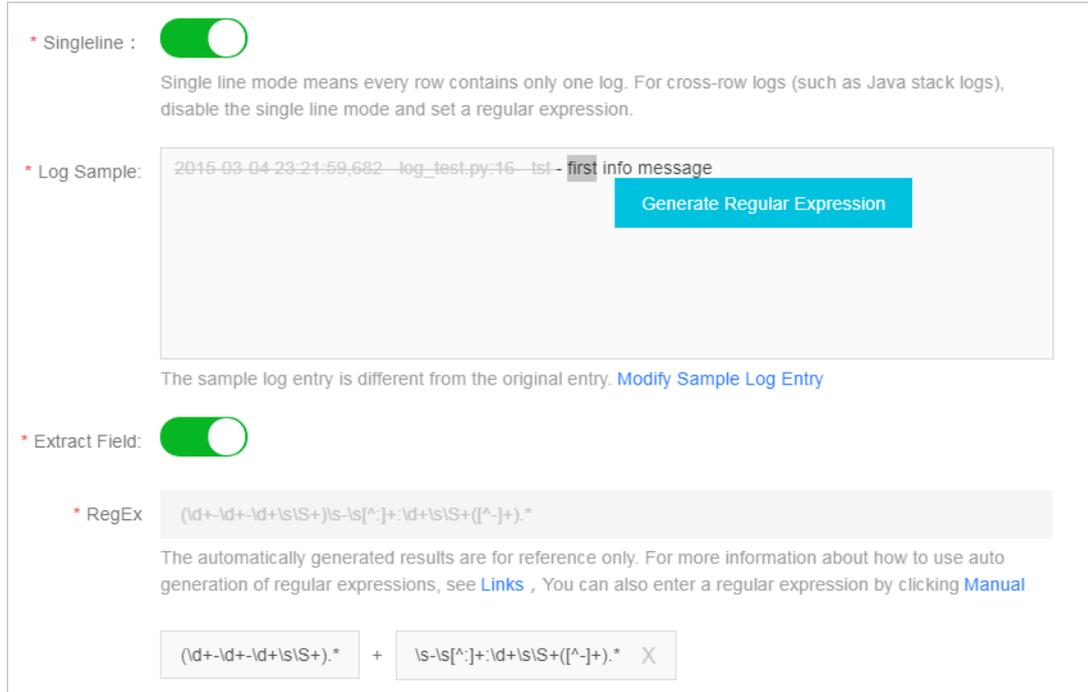
Select a server group and move the group from **Source Server Groups** to **Applied Server Groups**.



6. Create a Logtail configuration.
 - i. Enter a **configuration name** and **log path** and select **Full Regex Mode** in the Mode field.
 - ii. Turn on **Singleline**.
 - iii. Enter a snippet in the **Log Sample** field.
 - iv. Turn on **Extract Field**.

- v. Set a regular expression in the **RegEx** field.
 - a. Select fields to generate a regular expression

If the regular expression that is automatically generated does not match your sample log, you can select fields in the sample log to generate a regular expression. Log Service can automatically parse the highlighted fields of the sample log to generate a regular expression. In the **Log Sample** field, select the required fields, and click **Generate Regular Expression**. The regular expression of the selected field is displayed in the **RegEx** field. To obtain a full regular expression for the sample log, generate regular expressions for each log field.



- b. Modify the regular expression
- Actual data formats may vary. In this case, click **Manual** under the RegEx field to adjust the regular expression that is automatically generated based on your business requirements. This ensures that the regular expression is suitable for all formats of the collected logs.
- c. Verify the regular expression
- After you modify the regular expression, click **Validate** next to the RegEx field. If the regular expression is valid, the extraction results are displayed. If the regular expression is invalid, modify the regular expression again.

vi. Confirm the extraction results of log fields.

View the extraction results of log fields and specify keys for the extracted fields.

Specify an informative name for each log field in the extraction results. For example, time as the name for a time field. If you do not use the system time, you must specify the name of a time field in the Value fields and time in the Key field.

* Extracted Content:

| Key | Value |
|----------|---------------------|
| asctime | 2015-03-04 23:21:59 |
| filename | 682 - log_test.py |
| lineno | 16 |
| name | tst |
| message | first info message |

When you use a regular expression to generate key/value pairs, you can specify the key name in each pair. If you do not specify system time, you must specify a pair that uses "time" as the key name.

7. (Optional)Specify **Advanced Options** and click **Next**.

Specify **Advanced Options** based on your business requirements. We recommend that you do not modify the default settings unless otherwise required.

| Parameter | Description |
|---------------------------|---|
| Enable Plug-in Processing | Specifies whether to enable plug-in processing. If you turn on this switch, you can use the plug-in of Logtail to process text logs. |
| Upload Raw Log | Specifies whether to upload raw logs. If you turn on this switch, raw logs are written to the <code>__raw__</code> field and uploaded with the parsed logs. |
| Topic Generation Mode | <ul style="list-style-type: none"> ◦ Null - Do not generate topic: This mode is selected by default. In this mode, the topic is set to an empty string and you can query logs without the need to enter a topic. ◦ Server Group Topic Attributes: This mode is used to differentiate log data that is generated by different frontend servers. ◦ File Path RegEx: If you select this mode, you must enter a value in the Custom RegEx field to extract part of the path as the topic. This mode is used to differentiate log data that is generated by users or instances. |
| Custom RegEx | Specifies a custom regular expression. If you select File Path RegEx for Topic Generation Mode , you must enter a custom regular expression. |
| Log File Encoding | <ul style="list-style-type: none"> ◦ <code>utf8</code>: indicates UTF-8 encoding. ◦ <code>gbk</code>: indicates GBK encoding. |
| Timezone | <p>Specifies the time zone where logs are collected.</p> <ul style="list-style-type: none"> ◦ System Timezone: This option is selected by default. It indicates that the time zone where logs are collected is the same as the time zone to which the server belongs. ◦ Custom: Select a time zone. |

| Parameter | Description |
|----------------------|---|
| Timeout | <p>If a log file is not updated within a specific period of time, Logtail considers the file to be timed out.</p> <ul style="list-style-type: none"> Never: All log files are continuously monitored and never time out. 30 Minute Timeout: If a log file is not updated within 30 minutes, Logtail considers the log file to be timed out and no longer monitors the file. |
| Filter Configuration | <p>Only logs that meet all filter conditions are collected.</p> <p>Examples:</p> <ul style="list-style-type: none"> Collect logs that meet a condition: Set the condition to <code>Key:level Regex:WARNING ERROR</code>. It indicates that only logs with the severity level of WARNING or ERROR are collected. Filter logs that do not meet a condition: <ul style="list-style-type: none"> Set the condition to <code>Key:level Regex:^(?!.*(INFO DEBUG)).*</code>. It indicates that logs with the severity level of INFO or DEBUG are not collected. Set the condition to <code>Key:url Regex:^(?!.*(healthcheck)).*</code>. It indicates that logs whose URL contains the keyword healthcheck are not collected. For example, logs in which the key is url and the value is <code>/inner/healthcheck/jiankong.html</code> are not collected. |

8. Configure an index.

Configure an index based on your business requirements. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

After the configuration is complete, apply the settings to the server group to collect Python logs.

23.3.2.3.3. Node.js logs

Node.js logs are displayed in the Log Service console by default. This impedes data collection and troubleshooting. You can use the log4js function to write logs into files and customize the log format. This facilitates data collection and consolidation.

Example:

```
var log4js = require('log4js');
log4js.configure({
  appenders: [
    {
      type: 'file', //Output to a file
      filename: 'logs/access.log',
      maxLogSize: 1024,
      backups: 3,
      category: 'normal'
    }
  ]
});
var logger = log4js.getLogger('normal');
logger.setLevel('INFO');
logger.info("this is a info msg");
logger.error("this is a err msg");
```

Log format

After logs are written to text files by using the log4js function, these logs are displayed in the following format :

```
[2016-02-24 17:42:38.946] [INFO] normal - this is a info msg
[2016-02-24 17:42:38.951] [ERROR] normal - this is a err msg
```

The log4js function defines six log severity levels. They are TRACE, DEBUG, INFO, WARN, ERROR, and FATAL in ascending order of severity.

Use Logtail to collect Node.js logs

For more information about how to configure Logtail to collect Python logs, see [Python logs](#). Use configurations based on your network environment and business requirements.

The automatically generated regular expression is based on the sample log and may not apply to other logs. Therefore, you must make minor changes to the regular expression after it is generated. You can use the following sample Node.js logs to configure appropriate regular expressions for your logs.

Sample Node.js logs:

- Example 1

- Sample log

```
[2016-02-24 17:42:38.946] [INFO] normal - this is a info msg
```

- Regular expression:

```
\\([\\^]+)\\s\\([\\^]+)\\s\\(w+\\)\\s-(. *)
```

- Extracted fields:

```
time , level , loggerName , and message
```

- Example 2

- Sample log

```
[2016-01-31 12:02:25.844] [INFO] access - 42.120.73.203 - - "GET /user/projects/ali_sls_log? ignoreError=true HTTP/1.1" 304 - "http://aliyun.com/" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_10_3) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/48.0.2564.97 Safari/537.36"
```

- Regular expression:

```
\\([\\^]+)\\s\\(\\(w+\\)\\s\\(w+\\)\\s-\\s\\(S+\\)\\s-\\s-\\s"([\\^]+)"\\s\\(d+\\)[\\^]+("([\\^]+)"\\s"([\\^]+). *
```

- Extracted fields:

```
time , level , loggerName , ip , request , status , referer , and user_agent
```

23.3.2.3.4. WordPress logs

This topic describes the format of WordPress logs and extraction results of a sample log.

Log format

Sample log:

```
172.64.0.2 - - [07/Jan/2016:21:06:39 +0800] "GET /wp-admin/js/password-strength-meter.min.js? ver=4.4 HTTP/1.0" 200 776 "http://wordpress.c4a1a0aecdb1943169555231dcc4adfb7.cn-hangzhou.alicontainer.com/wp-admin/install.php" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_10_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/47.0.2526.106 Safari/537.36"
```

Configure Logtail to collect WordPress logs

Configurations required to collect WordPress logs:

- Regular expression that matches IP addresses that each indicate the start of a line

```
\d+\.\d+\.\d+\.\d+|s-|s.*
```

- Regular expression used to extract information from the log:

```
(S+) -- \[[^]]*\] "(S+) ([^"]+)" (S+) (\S+) "[^"]+" "[^"]+"
```

- Time conversion format:

```
%d/%b/%Y:%H:%M:%S
```

- Results after Logtail extracts information from the sample log

| Key | Value |
|------------|---|
| ip | 10.10.10.1 |
| time | 07/Jan/2016:21:06:39 +0800 |
| method | GET |
| url | /wp-admin/js/password-strength-meter.min.js? ver=4.4 HTTP/1.0 |
| status | 200 |
| length | 776 |
| ref | http://wordpress.c4a1a0aecdb1943169555231dcc4adfb7.cn-hangzhou.alicdn.com/wp-admin/install.php |
| user-agent | Mozilla/5.0 (Macintosh; Intel Mac OS X 10_10_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/47.0.2526.106 Safari/537.36 |

23.3.2.3.5. Unity3D logs

Log Service can use the WebTracking feature to collect Unity3D logs. The following example shows how to collect *Unity logs* of the debug type.

Context

Unity3D is a cross-platform game engine developed by Unity Technologies. The engine allows you to create 3D video games, VR buildings, real-time 3D animation, and other interactive content.

Procedure

1. Enable the WebTracking feature.

For more information about how to enable this feature, see [WebTracking](#).

2. Create a Unity3D LogHandler.

In a Unity editor, create a C# file named *LogOutputHandler.cs*, add the following code, and modify the following variables.

- project: specifies the name of the project.
- logstore: specifies the name of the Logstore.
- serviceAddr: specifies the endpoint of the project.

For more information about serviceAddr, see [View the information of a project](#).

```
using UnityEngine;
using System.Collections;
public class LogOutputHandler : MonoBehaviour
{
    //Register the HandleLog function on scene start to fire on debug.log events
    public void OnEnable()
    {
        Application.logMessageReceived += HandleLog;
    }
    //Remove callback when object goes out of scope
    public void OnDisable()
    {
        Application.logMessageReceived -= HandleLog;
    }
    string project = "your project name";
    string logstore = "your logstore name";
    string serviceAddr = "http address of your log service project";
    //Capture debug.log output, send logs to Loggly
    public void HandleLog(string logString, string stackTrace, LogType type)
    {
        string parameters = "";
        parameters += "Level=" + WWW.EscapeURL(type.ToString());
        parameters += "&";
        parameters += "Message=" + WWW.EscapeURL(logString);
        parameters += "&";
        parameters += "Stack_Trace=" + WWW.EscapeURL(stackTrace);
        parameters += "&";
        //Add any User, Game, or Device MetaData that would be useful to finding issues later
        parameters += "Device_Model=" + WWW.EscapeURL(SystemInfo.deviceModel);
        string url = "http://" + project + "." + serviceAddr + "/logstores/" + logstore + "/track? APIVersion=0.6.0&" + parameters;
        StartCoroutine(SendData(url));
    }
    public IEnumerator SendData(string url)
    {
        WWW sendLog = new WWW(url);
        yield return sendLog;
    }
}
```

The preceding code allows you to send logs to Log Service in an asynchronous manner. In the code, you can specify more fields you want to collect.

3. Generate Unity logs.

In the project, create a C# file named *LogglyTest.cs* and add the following code.

```
using UnityEngine;
using System.Collections.Generic;
public class LogglyTest : MonoBehaviour {
    void Start () {
        Debug.Log ("Hello world");
    }
}
```

4. View logs in the console.

After you complete the preceding steps, run the Unity application. In the Log Service console, view logs that are sent to Log Service.

The preceding code shows how to use `Debug.Log`, `Debug.LogError`, and `Debug.LogException` methods to collect logs. Unity provides Component Object Model (COM)-based exception handling and log handling APIs. These APIs allow you to easily collect device details of clients.

23.4. Query and analysis

23.4.1. Overview

Log Service provides the LogSearch/Analytics feature that you can use to query and analyze a large number of logs. If you do not enable indexes, raw data is consumed in sequence based on shards. The procedure is similar to the sequential consumption of Kafka messages. If you enable indexes, you can query logs and perform statistical analysis on query results in addition to consuming logs in sequence.

Benefits

- **Real-time:** Logs can be analyzed immediately after they are written.
- **Fast:**
 - **Query:** Billions of data records can be processed and queried within one second. Each search statement has a maximum of five conditions specified.
 - **Analysis:** Hundreds of millions of data records can be aggregated and analyzed within one second. Each query has a maximum of five aggregate functions and a GROUP BY clause specified.
- **Flexible:** Query and analysis conditions can be changed as required and the results are returned in real time.
- **All-in-one:** Reports and dashboards are available in the console for quick analysis. In addition to these features, Log Service can work together with Grafana, DataV, Jaeger, and other services. It also supports RESTful APIs, Java Database Connectivity (JDBC) APIs, and other APIs.

Indexing

Indexes refer to a data structure that you can use to sort the values of one or more columns of logs. Indexes allow you to obtain the required information in a timely manner from logs that Log Service collects. Before you use the LogSearch/Analytics feature, you must collect logs and [Enable the index feature and configure indexes for a Logstore](#) on the collected logs.

In Log Service, indexes are sorted into **full-text indexes** and **field-specific indexes**.

- **Full-text index:** Indexing is enabled for the full contents of a log. The values of all fields in a log are queried by default. The log can be queried if one of the fields matches the search term.
- **Field-specific index:** You can configure a field-specific index for a key. Then, you can query logs based on specific keys to narrow the query scope.

To use **field-specific indexes**, you must specify the data type for a field. Available data types for fields in Log Service include [Text](#), [JSON](#), [Long](#), and [Double](#). For more information about [Overview](#).

Query methods

- **Console**

In the Log Service console, you can query logs by specifying time ranges and search statements. For more information about the procedure and search statements, see [Query logs](#) and [Query syntax](#).

- **API**

To query logs, you can call the GetLogs and GetHistograms API operations of the Log Service API.

 **Note** Before you query logs, make sure that you collect logs and [Enable the index feature and configure indexes for a Logstore](#).

Search and analysis statements

To apply real-time LogSearch/Analytics to collected logs, you must specify query statements. Each query statement includes the search section and the analytics section. Separate the sections with a vertical bar (|).

`$Search|$Analytics`

| Statement | Required | Description |
|-----------|----------|---|
| Search | No | A search statement contains search conditions. These conditions include keywords, fuzzy keywords, values, ranges, and combined conditions. If you leave the statement empty or specify an asterisk (*) for the statement, it indicates that no condition is specified and all data is returned. For more information, see Query syntax . |
| Analytics | No | You can use an analytics statement to aggregate or analyze data based on query results. If you leave the statement empty, it indicates that no analytics is required and all query results are returned. For more information, see Real-time analysis . |

Precautions

You may query a large number of logs. For example, if the number of logs to be queried is more than 1,000,000,000, Log Service may fail to return all results. Log Service returns a partial result set and notifies you that the returned data set includes partial results.

Query results are cached every 15 minutes. If a partial result set is matched in the cache, Log Service continues to scan logs that are not cached. Log Service combines query results of the current query with results of cached results.

Therefore, Log Service enables you to obtain results by calling the API operation multiple times with the same parameters.

23.4.2. Real-time analysis

Log Service supports SQL-like aggregate calculation. This feature integrates search statements with SQL aggregate functions to calculate query results.

Sample statement:

```
status>200 |select avg(latency),max(latency) ,count(1) as c GROUP BY method ORDER BY c DESC LIMIT 20
```

Basic syntax:

```
[search query] | [sql query]
```

Separate a search statement and a calculation statement with a vertical bar (|). You can use the search statement to query logs and obtain the required results. Then, use the calculation statement for further aggregation. The search query syntax is specific to Log Service. For more information, see [Query syntax](#).

Prerequisites

To use the statistical analysis feature, click **Index Attributes**. Turn on the **Enable Analytics** switch for the required field. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

- If you turn off the switch, Log Service calculates up to 10,000 rows of data in each shard along with high latency

of calculation.

- If you turn on the switch, Log Service analyzes data in seconds.
- The update applies only to new data.
- No extra fee is incurred for the update.

Supported SQL syntax

Log Service supports the following SQL syntax. For more information about a specific topic, click the corresponding link.

- Aggregate functions that are available for SELECT statements include:
 - [General aggregate functions](#)
 - [Security check functions](#)
 - [Map functions](#)
 - [Approximate functions](#)
 - [Mathematical statistics functions](#)
 - [Mathematical calculation functions](#)
 - [String functions](#)
 - [Date and time functions](#)
 - [URL functions](#)
 - [Regular expression functions](#)
 - [JSON functions](#)
 - [Type conversion functions](#)
 - [IP functions](#)
 - [Array functions](#)
 - [Binary string functions](#)
 - [Bitwise operations](#)
 - [Interval-valued comparison and periodicity-valued comparison functions](#)
 - [Comparison functions and operators](#)
 - [Lambda functions](#)
 - [Logical functions](#)
 - [Geospatial functions](#)
 - [Geography functions](#)
 - [Machine learning functions](#)
- [GROUP BY syntax](#)
- [Window functions](#)
- [HAVING syntax](#)
- [ORDER BY syntax](#)
- [LIMIT syntax](#)
- [Syntax for CASE statements and if\(\) functions](#)
- [UNNEST function](#)
- [Field aliases](#)
- [Nested subqueries](#)

Precautions

Before you use SQL statements, note the following items:

- You do not need to specify FROM and WHERE clauses for SQL statements. By default, Log Service queries logs

from the current Logstore, and each WHERE clause is specified in the [search query] section.

- The supported clauses include SELECT, GROUP BY, ORDER BY [ASC,DESC], LIMIT, and HAVING.

Note By default, only the first 10 results are returned. If you want to return more results, add a LIMIT clause to the statement. For example, `* | select count(1) as c, ip group by ip order by c desc limit 100`.

Built-in fields

Log Service has multiple built-in fields for statistical analysis. A built-in field is automatically added to a valid column that you create.

| Field | Type | Description |
|-------------------------|---------|---|
| <code>__time__</code> | Bigint | The time when a log was created. |
| <code>__source__</code> | Varchar | The source IP address of a log. When you query logs, the name of the field is source. If you specify the field for an SQL statement, you must add two underscores (__) at both the start and end of source. |
| <code>__topic__</code> | Varchar | The topic of a log. |

Limits

- Maximum number of Logstores from which you can query logs at the same time is 15.
- Maximum length for a field value of the varchar type is 2,048. Extra data will be truncated.
- By default, up to 100 lines of a log file are returned and pagination is not supported. If you want to return more lines, use [LIMIT syntax](#).

Example

Calculate the hourly PV and UV, and the user request of the highest latency.

```
*|select date_trunc('hour',from_unixtime(__time__)) as time,
count(1) as pv,
approx_distinct(userid) as uv,
max_by(url,latency) as top_latency_url,
max(latency,10) as top_10_latency
group by 1
order by time
```

23.4.3. Enable the indexing feature and configure indexes for a Logstore

This topic describes how to enable the indexing feature and configure indexes for a Logstore.

Context

Before you can query logs that are stored in a Logstore, you must enable the indexing feature and configure indexes for the Logstore. We recommend that you configure indexes for your Logstores based on your business requirements.

Note After you enable the indexing feature, the indexes occupy extra storage space and transferring the indexes occupies extra bandwidth.

When you collect a log entry, Log Service adds the relevant information (such as the source and time fields) to the log entry as key-value pairs. These fields are reserved in Log Service. If you enable the indexing feature for a Logstore and configure indexes for fields in the Logstore, the indexing and analytics features are automatically enabled for these fields.

Reserved fields in Log Service

| Field | Description |
|-------------------------|---|
| <code>__topic__</code> | The topic of a log entry. If you specify a topic for a log entry, Log Service adds the topic field to the log entry. The key of the field is <code>_topic</code> and the value of the field is the log topic. For more information, see Specify a log topic . |
| <code>__source__</code> | The source of a log entry. The source device that generates the log entry. |
| <code>__time__</code> | The time when a log entry is written to the Logstore by using an SDK. |

Note If the values of the `__topic__` and `__source__` fields are null, the keywords that you use to query the two fields must exactly match the field values.

Procedure

1. [Log on to the Log Service console](#).
2. In the Projects section, click the target project.
3. Find the target Logstore, and choose  > **Search & Analysis**.
4. On the page that appears, click **Enable** in the upper-right corner.

Note If you have created indexes for the Logstore, choose **Index Attributes > Modify** to modify the indexes.

5. Configure the indexes.

Note If you enable a full-text index and a field-specific index at the same time, the field-specific index takes precedence over the full-text index.

Index types

| Index type | Description |
|-----------------|--|
| Full text index | An index is created in the text format for all fields. You can search for key-value pairs that are included in these fields. For fields of the LONG type, you must specify the key name of a field when you query a value of the field. For fields of the other types, you do not need to specify a key name in queries. |

| Index type | Description |
|----------------------|---|
| Field-specific index | <p>After you configure a field-specific index, you must specify the name of a key when you query logs. If you configure the field-specific index on a field, the field-specific index takes effect when you query logs. The full-text index does not take effect.</p> <p>Available data types that you can specify for fields include:</p> <ul style="list-style-type: none"> Query text data JSON indexes Numeric (LONG and DOUBLE) |

- Configure a full-text index.

After you configure a full-text index for a Logstore, the values of all fields in the Logstore are queried by default.

| Parameter | Description | Example value |
|-----------------|--|--|
| Full Text Index | If you turn on the switch, Log Service traverses the values of all fields in a log entry. If the value of one of the fields matches the keyword, the log entry is returned. | - |
| Case Sensitive | <p>Specifies whether queries are case-sensitive.</p> <ul style="list-style-type: none"> If you turn off the switch, queries are not case-sensitive. For example, if you search for <code>internalError</code>, you can use either <code>INTERNALERROR</code> or <code>internalerror</code> as the keyword. If you turn on the switch, queries are case-sensitive. For example, if you search for <code>internalError</code>, you can use only <code>internalError</code> as the keyword. | - |
| Include Chinese | <p>Specifies whether to differentiate the Chinese content and English content.</p> <ul style="list-style-type: none"> If you turn on the switch, Log Service separates the Chinese content based on the Chinese semantics and English content based on the specified delimiters. If you turn off the switch, the content of a log entry is separated by the specified delimiters. | - |
| Delimiter | <p>The delimiters that you use to separate the content of a log entry into multiple keywords.</p> <p>For example, the content of a log entry is <code>a,b;c;D-F</code>. You can specify commas (,), semicolons (;), and hyphens (-) as delimiters to delimit the log content. Then you can use the five letters a, b, c, D, and F as keywords to match the log entry.</p> | <code>,";=()[]{}? @&<>:\n\t</code> |

- Configure a field-specific index.

You can specify fields to be indexed. Field-specific indexes allow you to query log data based on the values of specific fields. This narrows down the query scope.

| Parameter | Description | Example value |
|----------------|---|---|
| Key Name | <p>The name of a log field.</p> <div style="background-color: #e6f2ff; padding: 10px; border: 1px solid #d9e1f2;"> <p>Note</p> <ul style="list-style-type: none"> If you want to configure an index for fields of the tag type, such as fields that include public IP addresses or UNIX timestamps, you must set the value of the Key Name parameter in the <code>__tag__:key</code> format, for example, <code>__tag__:__receive_time__</code>. Indexes of the numeric types are unavailable for tag fields. You must select text in the Type field for all tag fields. </div> | <code>_address_</code> |
| Type | <p>The type of a field. Valid values:</p> <ul style="list-style-type: none"> text: The data type of the field is TEXT. long: The data type of the field is LONG. You must specify a numeric range to query log data. double: The data type of the field is DOUBLE. You must specify a numeric range to query log data. json: The data type of the field is JSON. <div style="background-color: #e6f2ff; padding: 10px; border: 1px solid #d9e1f2;"> <p>Note The Case Sensitive, Include Chinese, and Delimiter parameters are unavailable for fields of numeric types (LONG and DOUBLE).</p> </div> | - |
| Alias | <p>The alias of a column.</p> <p>Aliases are applied only to SQL statistics. Original names are applied when you store and query log data in Log Service. For more information, see Field aliases.</p> | <code>address</code> |
| Case Sensitive | <p>Specifies whether queries are case-sensitive.</p> <ul style="list-style-type: none"> If you turn off the switch, queries are not case-sensitive. For example, if you search for <code>internalError</code>, you can use either <code>INTERNALERROR</code> or <code>internalerror</code> as the keyword. If you turn on the switch, queries are case-sensitive. For example, if you search for <code>internalError</code>, you can use only <code>internalError</code> as the keyword. | - |
| Delimiter | <p>The delimiters that you use to separate the content of a log entry into multiple keywords.</p> <p>For example, the content of a log entry is <code>a,b;c;D-F</code>. You can specify commas (,), semicolons (;), and hyphens (-) as delimiters to delimit the log content. Then you can use the five letters a, b, c, D, and F as keywords to match the log entry.</p> | <code>,";=()[]{}? @&<>:/\n\t</code> |

| Parameter | Description | Example value |
|------------------|---|---------------|
| Include Chinese | <p>Specifies whether to differentiate the Chinese content and English content.</p> <ul style="list-style-type: none"> ▪ If you turn on the switch, Log Service separates the Chinese content based on the Chinese semantics and English content based on the specified delimiters. ▪ If you turn off the switch, the content of a log entry is separated by the specified delimiters. | - |
| Enable Analytics | <p>Specifies whether to enable the analytics feature. The switch is turned on by default.</p> <p>After you turn on the switch, you can use search and analytic statements to obtain statistical results.</p> | - |

6. Click **OK**.

 **Note**

- The index configurations take effect within one minute.
- After an index is enabled or modified, the updates on the index apply only to new data that is written to Log Service.

23.4.4. Query logs

This topic describes how to query logs in a Logstore. After you enable and configure the index of the Logstore, you can query and analyze logs in a Logstore in real time.

Prerequisites

- Logs are collected and stored in a Logstore.
- Indexes are enabled and configured. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

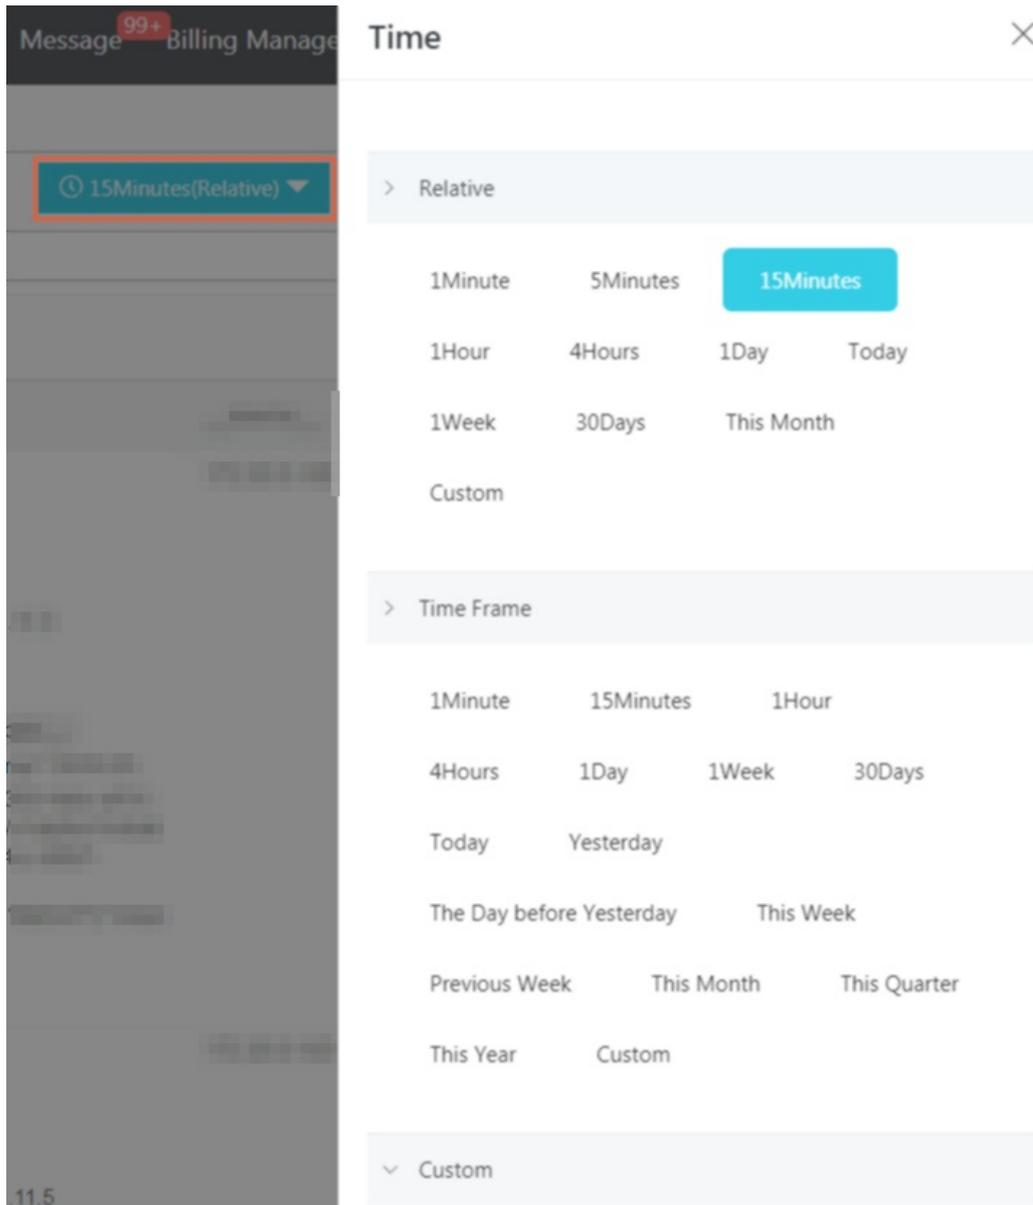
Procedure

1. [Log on to the Log Service console](#).
2. Click the  icon next to the name of the Logstore, and then select **Search & Analysis**.
3. Enter a query statement in the search box.

A query statement consists of a search statement and an analytic statement. The syntax is `search statement|analytic statement`. For more information, see [Search and analysis statements](#).

4. On the Search & Analysis page, select **15 Minutes** in the Relative section to set the time range for the query. You can select a relative time, time frame, or a custom time range.

 **Note** The query results may contain logs that are generated 1 minute earlier or later than the specified time period.



5. Click **Query & Analyze** to view the query results.

Log Service illustrates query results on log distribution histograms, Raw Logs tab, or statistical graphs.

Note 100 results are returned by default. For information about how to retrieve more results, see [LIMIT syntax](#).

- o Log distribution histogram

The log distribution histogram shows the distribution of query results across different time ranges.

- Move the pointer over a green block to view a time range and the number of logs obtained within the time range.
- Click a data block to view finer-grained log distribution. You can also view the query results on the **Raw Logs** tab.



o Raw Logs tab

On the **Raw Logs** tab, view logs that match your search conditions.

- **Quick analysis:** Use this feature to analyze the distribution of values for a specific field within a period of time. For more information, see [Quick analysis](#).
- **Log download:** Click the download icon in the upper-right corner of the tab, select a time range, and then click **OK**.
- **Column settings:** Click **Column Settings** in the upper-right corner of the tab, select the required fields and click **Add** to add the fields. Then, the columns that correspond to the fields appear on the tab. The field names are also column names. The columns list the field values.

? **Note** To view the log contents on the tab, select **Content**.

- **Content column settings:** If the content of a field exceeds 3,000 characters, extra characters will be hidden. In this case, the message "The character string is too long and has been truncated" will be displayed before the Key field. Click **Display Content Column**. In the dialog box that appears, set the **Key-Value Arrangement** and **Truncate Character String** parameters.

? **Note** If the content limit is set to 10,000 characters, no delimiter will be specified for extra characters.

| Parameter | | Description |
|-----------------------------------|----------------------|---|
| Key-Value Pair Arrangement | | You can set this parameter to New Line or Full Line . |
| Truncate Character String | Key | If a field value contains more than 3,000 characters, the field value is truncated. However, this parameter remains unspecified if no field value exceeds 3,000 characters. The value of this parameter is the key of the truncated value. |
| | Status | This parameter determines whether to enable the value truncation feature. By default, the feature is enabled. <ul style="list-style-type: none"> ■ Enable: If the value in a key-value pair exceeds the specified Truncate Step, extra characters will be truncated. You can click the Show button at the end of the value to show the truncated characters. The increment per click is the specified truncate step. ■ Disable: If the value in the key-value pair exceeds the specified Truncate Step, extra characters will not be truncated. |
| | Truncate Step | This parameter specifies the maximum number of characters that a field value shows by default. The parameter also specifies the number of extra characters that you displayed each time you click the Show button. Valid values: 500 to 10000. Default value: 3000. |

o Graph

If you enable the Analytics feature on the Search & Analysis page and use search and analytic statements to query logs, you can view the analytical results on the **Graph** tab.

- Graphs of multiple types are provided in Log Service, including tables, line charts, and bar charts. You can select a graph to show the required analytical results. For more information, see [Graphs](#).
- Log Service allows you to create dashboards for real-time data analysis. For more information, see [Create and delete a dashboard](#). Click **Add to New Dashboard** to save a common chart as query statements to a dashboard.
- Drill-down analysis allows you to move to deeper data layers, which reveals more detailed information. You can set the drill-down parameters and add the chart to the dashboard. Then, you can click the values in the chart to view the analysis results from more dimensions. For more information, see [Drill-down analysis](#).

You can also click **Save Search** or **Save as Alarm** on the Search & Analysis page to use the saved search and alarm features. For more information, see [Save a query statement as a search](#) and [Configure an alert](#).

23.4.5. Export logs

You can export logs on the current page to a CSV file and save the file to your localhost.

Procedure

1. [Log on to the Log Service console](#).
2. Click a project name.
3. Click the  icon next to the name of the Logstore, and then select **Search & Analysis**.
4. Click the  icon next to the **Raw Logs** tab.
5. In the **Download Log** dialog box, select **Download Log in Current Page**.
6. Click **OK** to export logs of the current page to a .CSV file and save the file to the localhost.

23.4.6. Index data type

23.4.6.1. Overview

Log Service allows you to use full-text indexes or field-specific indexes to query collected logs. If you set a full-text index for a log, the value is the entire log. If you set a field-specific index for a log, you can specify a data type for each key.

Date types

The following table lists the supported data types.

| Query type | Data type (index) | Description | Example |
|-------------|-------------------|--|---|
| Basic query | Text | The text type. You can use keywords and fuzzy matches to query logs. | <code>uri:"login*" method:"post"</code> |
| | Long | The numeric type. You can specify numeric ranges to query logs. | <code>status>200 and status in [200, 500]</code> |
| | Double | The floating-point type. | <code>price>28.95 and t in [20.0, 37]</code> |

| Query type | Data type (index) | Description | Example |
|----------------|-------------------|---|---|
| Combined query | JSON | Indicates that the index is a JSON field that supports nested queries. By default, the data type of the field is text. You can set indexes of the Text, Long, and Double types for the elements at layer a in the a.b path format. The fields adopt the configured types. | <code>level0.kev>29.95</code> <code>level0.key2:"action"</code> |
| | Text | Indicates that the full contents of the log are queried as text. | <code>error and "login fail"</code> |

Query examples

The following table lists the keys included in the sample log.

| No. | Key | Type |
|-----|---------|--------|
| 0 | time | N/A |
| 1 | class | text |
| 2 | status | long |
| 3 | latency | double |
| 4 | message | json |

```

0. time:2018-01-01 12:00:00
1. class:central-log
2. status:200
3. latency:68.75
4. message:
{
  "methodName": "getProjectInfo",
  "success": true,
  "remoteAddress": "1.1.1.1:11111",
  "usedTime": 48,
  "param": {
    "projectName": "ali-log-test-project",
    "requestId": "d3f0c96a-51b0-4166-a850-f4175dde7323"
  },
  "result": {
    "message": "successful",
    "code": "200",
    "data": {
      "clusterRegion": "ap-southeast-1",
      "ProjectName": "ali-log-test-project",
      "CreateTime": "2017-06-08 20:22:41"
    },
    "success": true
  }
}

```

You can set an index as follows.

Set an index

| Field Search | | | | | | | Automatic Index Generation | |
|-------------------|------|-------|--------------------------|--|---------------------|--------------------------|-------------------------------------|--------|
| Key Name | Type | Alias | Enable Search | | Delimiter: ? | Include Chinese | Enable Analytics | Delete |
| | | | Case Sensitive | | | | | |
| class | text | | <input type="checkbox"/> | | , "';=000?@&<>/\n\r | <input type="checkbox"/> | <input checked="" type="checkbox"/> | × |
| info | json | | <input type="checkbox"/> | | , "';=000?@&<>/\n\r | <input type="checkbox"/> | <input type="checkbox"/> | × |
| methodName | text | | | | | | <input checked="" type="checkbox"/> | × |
| param.projectName | text | | | | | | <input checked="" type="checkbox"/> | × |
| param.requestId | text | | | | | | <input checked="" type="checkbox"/> | × |
| result.code | long | | | | | | <input checked="" type="checkbox"/> | × |
| result.message | text | | | | | | <input checked="" type="checkbox"/> | × |
| success | text | | | | | | <input checked="" type="checkbox"/> | × |
| usedTime | long | | | | | | <input checked="" type="checkbox"/> | × |
| + | | | | | | | | |
| latency | long | | | | | | <input checked="" type="checkbox"/> | × |
| status | long | | | | | | <input checked="" type="checkbox"/> | × |

In the preceding figure,

- ① specifies that Log Service can query data of the string and Boolean in JSON fields.
- ② specifies that Log Service can query data of the long type.
- ③ enables SQL analysis for specified fields.

Example

1. Query data of the string and Boolean types

- You do not need to configure JSON fields.
- JSON maps and arrays are automatically expanded and can contain nested fields. Separate multiple levels with periods (.).

```
class : cental*
message.traceInfo.requestId : 92.137_1518139699935_5599
message.param.projectName : ali-log-test-project
message.success : true
```

2. Query data of the double and long types

Each JSON field must be specified separately and cannot be contained in a JSON array.

```
latency>40
message.usedTime > 40
```

3. Combined query

```
class : cental* and message.usedTime > 40 not message.param.projectName:ali-log-test-project
```

23.4.6.2. Query text data

This topic describes how to query text data.

Similar to search engines, Log Service queries text data based on terms. Therefore, you must set the Delimiter, Case Sensitive fields.

Configurations

- Case sensitivity

You can specify whether log queries are case-sensitive. For example, you want to query logs by using a search term named `internalError`.

- o `false` specifies a case-insensitive query. Both `INTERNALERROR` and `internalerror` can be the keywords.
- o `true` specifies a case-sensitive query. Only the `internalError` can be the keyword.

• Delimiter

You can use delimiters to split a search term into multiple keywords.

For example, you want to query logs by using the following search term.

```
/url/pic/abc.gif
```

- o If no delimiter is set, the entire `/url/pic/abc.gif` string is treated as a keyword. You must use the entire string as a keyword or a fuzzy string named `/url/pic/*` to query logs.
- o If the delimiter is set to `/`, the search term is split into three words: `url`, `pic`, and `abc.gif`. You can use one of these words or a fuzzy word to query logs. For example, `url`, `abc.gif`, or `pic*`. You can also use the `/url/pic/abc.gif` string as a search term to query logs. However, the search term is split into three keywords named `url`, `pic`, and `abc.gif`.
- o If the delimiter is set to `./`, the search term is split into four keywords named `url`, `pic`, `abc`, and `gif`.

 **Note** You can extend query ranges by setting appropriate delimiters.

• Full-text index

By default, full-text indexes treat each log except for the time field as text data. You do not need to specify any keys for a full-text index. For example, the following log includes the time field, status field, level field, and message field.

```
[20180102 12:00:00] 200,error,some thing is error in this field
```

- o `time:2018-01-02 12:00:00`
- o `level:"error"`
- o `status:200`
- o `message:" some thing is error in this field"`

 **Note**

- o Prefixes are not required for full-text indexes. If you set the search term to `error`, the level and message fields that include error match the search term.
- o You must set delimiters for full-text indexes. For example, if you set a space () as a delimiter, the `status:200` string is a search term. If you set a colon (:) as a delimiter, the search term is split into two keywords named `status` and `200`.
- o Numbers are treated as text data. For example, you can use `200` to query logs. Values in the time field are not treated as text data.
- o You can query logs by using keys such as `status`.

23.4.6.3. Numeric type

When you configure indexes, you can set the data type of a key to number. To query logs, you can specify a numeric range for the key.

Configurations

Supported types: `long` (long integers) and `double` (decimals). After you set the data type of a key to number, you must specify a numeric range for the key to query logs.

Query examples

To specify a numeric range from 1000 to 2000 (excluding 1000) for a key of the long type, you can use the following methods:

- Query syntax for numbers. For example:

```
longKey > 1000 and longKey <= 2000
```

- Query grammar for numeric ranges. For example:

```
longKey in (1000 2000]
```

For more information about query syntax, see [Query syntax](#).

23.4.6.4. JSON indexes

Log Service can query and analyze logs in the JSON format. You can set the data type of indexes to JSON.

JSON texts include data of multiple types, including string, Boolean, number, array, and map. JSON-formatted data is self-parsed and flexible. You can use JSON-formatted data in various scenarios. In most cases, variable log fields are recorded in the JSON format. For example, HTTP request and response parameters are recorded in a log in the JSON format.

Log Service allows you to set the data type of index fields to JSON so that you can query and analyze logs in the JSON format.

Configurations

- Log Service can parse JSON-formatted fields and generate indexes for all the fields of the text and Boolean types.

```
json_string.key_map.key_text : test_value
json_string.key_map.key_bool : true
```

- To query fields of the double or long type that is not in a JSON array, you can specify a JSON path.

```
Set the data type of the key_map.key_long field to long.
Search condition: json_string.key_map.key_long > 50
```

- To query fields of the text, double, or long type that is not in a JSON array, you can enable the Analytics feature and use SQL statements to analyze these fields.

```
json_string.key_map.key_long > 10 | select count(*) as c,
"json_string.key_map.key_text" group by
"json_string.key_map.key_text"
```

Note

- JSON objects and JSON arrays are not supported.
- Fields cannot be contained in JSON arrays.
- Fields of the Boolean type can be converted into the text type.
- To query and analyze logs, JSON-formatted fields must be enclosed with double quotation marks ("").

- Log Service cannot parse invalid JSON-formatted data.

Log Service does not stop parsing logs until it detects an invalid field.

In the following example, data after the key_3 field is truncated and lost in the following text. Log Service can parse the json_string.key_map.key_2 field and the contents before this field.

```
"json_string":
{
  "key_1": "value_1",
  "key_map":
  {
    "key_2": "value_2",
    "key_3": "valu
```

Query syntax

To query a specific key, you must add the JSON parent path to the query statement as the prefix of the key. The query syntax for the fields of the text and numeric types is the same for both JSON-formatted data and other data. For more information, see [Query syntax](#).

Query example

The following table lists the keys included in the sample log. The data type of the `message` key is JSON.

| Number | Key | Type |
|--------|---------|--------|
| 0 | time | N/A |
| 1 | class | text |
| 2 | status | long |
| 3 | latency | double |
| 4 | message | json |

```
0. time:2018-01-01 12:00:00
1. class:central-log
2. status:200
3. latency:68.75
4. message:
{
  "methodName": "getProjectInfo",
  "success": true,
  "remoteAddress": "1.1.1.1:11111",
  "usedTime": 48,
  "param": {
    "projectName": "ali-log-test-project",
    "requestId": "d3f0c96a-51b0-4166-a850-f4175dde7323"
  },
  "result": {
    "message": "successful",
    "code": "200",
    "data": {
      "clusterRegion": "ap-southeast-1",
      "ProjectName": "ali-log-test-project",
      "CreateTime": "2017-06-08 20:22:41"
    },
    "success": true
  }
}
```

You can set indexes for the log as follows:

Set an index

| Field Search | | | | | | | Automatic Index Generation |
|-------------------|---------------|-------|--------------------------|-------------------|--------------------------|-------------------------------------|----------------------------|
| Key Name | Enable Search | | | | Include Chinese | Enable Analytics | Delete |
| | Type | Alias | Case Sensitive | Delimiter: ? | | | |
| class | text | | <input type="checkbox"/> | , "=000?@&<>/\n\r | <input type="checkbox"/> | <input checked="" type="checkbox"/> | × |
| info | json | | <input type="checkbox"/> | , "=000?@&<>/\n\r | <input type="checkbox"/> | <input type="checkbox"/> | × |
| methodName | text | | | | | <input checked="" type="checkbox"/> | × |
| param.projectName | text | | | | | <input checked="" type="checkbox"/> | × |
| param.requestId | text | | | | | <input checked="" type="checkbox"/> | × |
| result.code | long | | | | | <input checked="" type="checkbox"/> | × |
| result.message | text | | | | | <input checked="" type="checkbox"/> | × |
| success | text | | | | | <input checked="" type="checkbox"/> | × |
| usedTime | long | | | | | <input checked="" type="checkbox"/> | × |
| latency | long | | | | | <input checked="" type="checkbox"/> | × |
| status | long | | | | | <input checked="" type="checkbox"/> | × |

In the preceding figure:

- ① specifies that Log Service can query data of the string and Boolean types in JSON fields.
- ② specifies that Log Service can query data of the long type.
- ③ enables SQL analysis for specified fields.

Examples

1. Query data of the string and Boolean types

Note

- You do not need to configure JSON fields.
- JSON maps and arrays are automatically expanded and can include hierarchical levels. Separate multiple levels with periods (.).

```
message.traceInfo.requestId : 92.137_1518139699935_5599
message.param.projectName : ali-log-test-project
message.success : true
message.result.data.ProjectStatus : Normal
```

2. Query fields of the double and long types

Note Each JSON field must be configured and cannot be contained in an array.

```
message.usedTime > 40
```

3. Use SQL statements to analyze fields

Note

- Each JSON field must be configured and cannot be contained in an array.
- Each field to be queried must be enclosed with double quotation marks (" ") or be configured with an alias.

```
* | select avg("message.usedTime") as avg_time ,
"message.methodName" group by "message.methodName"
```

23.4.7. Query syntax and functions

23.4.7.1. Search syntax

This topic describes the search syntax that is used in Log Service.

Search types

After you [enable and configure the index feature](#) of a Logstore, you can enter a search statement on the search and analysis page to [query logs](#).

A query statement consists of two sub-statements in sequence: a search statement and an analytic statement. A search statement specifies one or more search conditions and returns the log entries that match the search conditions. You can execute a search statement to perform a full-text search or field-specific search.

- Full-text search

During full-text search, a log entry is considered a key-value pair. The value in the key-value pair indicates the content of the log entry. A full-text search statement returns the log entries that include or exclude the specified keywords.

Full-text search is divided into basic full-text search, phrase search, and wildcard-based search.

- Basic full-text search: You can specify keywords and operators in search conditions of a search statement. You can then execute the search statement to query the log entries that match the search conditions.

For example, the `a and b` statement returns the log entries that include the `a` and `b` keywords.

- Phrase search: A phrase is a string that is enclosed in double quotation marks (""). Substrings in a phrase are separated by space characters. Each substring is a keyword.

For example, the `"http error"` statement returns the log entries that contain the `http` and `error` keywords. This statement is equivalent to `http and error`.

- Wildcard-based search: You can use an asterisk (`*`) or a question mark (`?`) as a wildcard character in a keyword. Each keyword that includes wildcards can contain 1 to 64 characters in length and cannot start with a wildcard character. If a search condition contains a keyword that includes a wildcard character, Log Service returns a maximum of 100 log entries and each log entry contains a word that matches the keyword pattern.

For example, if you execute the `addr?` statement, Log Service returns a maximum of 100 log entries and each log entry contains a word that is prefixed with `addr`.

When you use wildcard-based search, note the following information:

- A keyword cannot start with an asterisk (`*`) or a question mark (`?`).
- The more accurate the keyword is, the more accurate the search results will be.
- Wildcard-based search is not supported for a keyword that contains more than 64 characters in length.
- A search statement returns a maximum of 100 log entries that match the search conditions.

- Field-specific search

After you configure the field index, you can search log entries based on the keys and values of the fields in the field index. For a field of the DOUBLE or LONG type, you can specify a value range for search. For example, the `Latency>5000 and Method:Get* and not Status:200` statement returns the log entries that meet the following conditions: The value of the `Latency` field is greater than 5000, the value of the `Method` field is prefixed with `Get`, and the value of the `Status` field is not 200.

You can perform a basic query or combined query, depending on the data types of the fields in the field index. For more information, see [Overview](#).

Additional considerations

- If you execute a search statement to perform both full-text search and field-specific search and you set different delimiters for the two search types, the delimiter that is set for field-specific search is used.

- You must set the data type of a field to `DOUBLE` or `LONG` before you specify a value range to search the field. If the data type of a field is not `DOUBLE` or `LONG` or the value range syntax is incorrect, the field-specific search condition is considered a full-text search condition. In this case, unexpected search results may be returned.
- If you change the data type of a field from `TEXT` to `DOUBLE` or `LONG`, only the equal-to operator (`=`) can be used to search for the log entries that are collected before the change.

Operators

The following table lists the operators that are supported by search statements.

| Operator | Description |
|--------------------|---|
| <code>and</code> | A binary operator. The syntax is <code>query1 and query2</code> . It indicates the intersection of the search results of <code>query1</code> and <code>query2</code> . The default operator between keywords is <code>and</code> . |
| <code>or</code> | A binary operator. The syntax is <code>query1 or query2</code> . It indicates the union of the search results of <code>query1</code> and <code>query2</code> . |
| <code>not</code> | A binary operator. The syntax is <code>query1 not query2</code> . It indicates that the log entries that match <code>query1</code> but do not match <code>query2</code> are returned. The syntax is equivalent to <code>query1-query2</code> . You can also use the <code>not query1</code> syntax. It indicates that the log entries that do not match <code>query1</code> are returned. |
| <code>(,)</code> | The operator that merges one or more sub-conditions into one search condition. The search based on a sub-condition that is enclosed in parentheses (<code>()</code>) is performed first. |
| <code>:</code> | The operator that is used to specify a pattern of key-value pairs. The syntax is <code>term1:term2</code> . If the key or value contains reserved characters such as spaces and colons (<code>:</code>), use double quotation marks (<code>""</code>) to enclose the entire key or value. |
| <code>"</code> | The operator that converts another operator into a common character. All terms enclosed in double quotation marks (<code>""</code>) are considered keywords rather than operators. In a field-specific search statement, you can enclose the entire key or value in double quotation marks. |
| <code>\</code> | The operator that escapes a double quotation mark. The escaped double quotation mark is considered a symbol instead of an operator. Example: <code>"\"</code> . |
| <code> </code> | The pipeline operator that is used to chain a search statement and an analytic statement. The analytic statement that follows the pipeline operator is executed based on the result of the search statement that the pipeline operator follows. Example: <code>query1 select count(1)</code> . |
| <code>count</code> | The count operator that is used to summarize the number of log entries. |
| <code>*</code> | The wildcard character that is used to replace zero or more characters. For example, the <code>que*</code> statement returns the log entries with a word that is prefixed with <code>que</code> . Note A wildcard-based search statement returns a maximum of 100 log entries that match the search condition. |
| <code>?</code> | The wildcard character that replaces a single character. The <code>qu?ry</code> statement returns the log entries with a word that is prefixed with <code>qu</code> , is suffixed with <code>ry</code> , and contains a character in between. |

| Operator | Description |
|------------------------|--|
| <code>__topic__</code> | The operator that specifies zero or more topics from which to query log entries. Example: <code>__topic__:mytopicname</code> . |
| <code>__tag__</code> | The operator that specifies a tag value of a tag key to query. Example: <code>__tag__:tagkey:tagvalue</code> . |
| <code>source</code> | The operator that specifies the IP address of a log source whose log entries you want to query. Example: <code>source:127.0.0.1</code> . |
| <code>></code> | The greater-than operator. You can use this operator to query the log entries whose value of a field is greater than a specified number. Example: <code>latency > 100</code> . |
| <code>>=</code> | The greater-than-or-equal-to operator. You can use this operator to query the log entries whose value of a field is greater than or equal to a specified number. Example: <code>latency >= 100</code> . |
| <code><</code> | The less-than operator. You can use this operator to query the log entries whose value of a field is less than a specified number. Example: <code>latency < 100</code> . |
| <code><=</code> | The less-than-or-equal-to operator. You can use this operator to query the log entries whose value of a field is less than or equal to a specified number. Example: <code>latency <= 100</code> . |
| <code>=</code> | The equal-to operator. You can use this operator to query the log entries whose value of a field is equal to a specified number. Example: <code>latency = 100</code> . |
| <code>in</code> | The operator that is used to query the log entries whose value of a field falls in a specified range. Brackets [] indicate closed intervals and parentheses () indicate open intervals. The beginning number and ending number of the range are enclosed in brackets or parentheses and separated by one or more space characters. The in operator must be in lowercase. Example: <code>latency in [100 200]</code> or <code>latency in (100 200)</code> . |

Note

- All operators except the in operator are case-insensitive.
- You can use the following operators, which are sorted in descending order of precedence: `:` , `"` , `()` , `and` , `not` , and `or` .
- Log Service uses the following operators: `sort` , `asc` , `desc` , `group by` , `avg` , `sum` , `min` , `max` , and `limit` . If you need to use these operators as keywords, enclose them in double quotation marks ("").

Search statement examples

| Expected search result | Search statement |
|---|--|
| Log entries that contain a and b | <code>a and b</code> or <code>a b</code> |
| Log entries that contain a or b | <code>a or b</code> |
| Log entries that contain a but do not contain b | <code>a not b</code> |
| Log entries that do not contain a | <code>not a</code> |
| Log entries that contain a and b but do not contain c | <code>a and b not c</code> |

| Expected search result | Search statement |
|---|---|
| Log entries that contain a or b and contain c | <code>(a or b) and c</code> |
| Log entries that contain a or b but do not contain c | <code>(a or b) not c</code> |
| Log entries that contain a and b and may contain c | <code>a and b or c</code> |
| Log entries whose FILE field contains apsara | <code>FILE:apsara</code> |
| Log entries whose FILE field contains apsara and shennong | <code>FILE:"apsara shennong" , FILE:apsara FILE: shennong , or FILE:apsara and FILE:shennong</code> |
| Log entries that contain the following keyword: and | <code>and</code> |
| Log entries whose FILE field contains apsara or shennong | <code>FILE:apsara or FILE:shennong</code> |
| Log entries whose file info field contains apsara | <code>"file info":apsara</code> |
| Log entries that contain double quotation mark (") | <code>\"</code> |
| Log entries with words that are prefixed with shen | <code>shen*</code> |
| Log entries whose FILE field is prefixed with shen | <code>FILE:shen*</code> |
| Log entries whose value of the FILE field is shen* | <code>FILE: "shen*"</code> |
| Log entries with words that are prefixed shen, are suffixed with ong, and contain a single character in between | <code>shen?ong</code> |
| Log entries with words that are prefixed with shen and words that are prefixed with aps | <code>shen* and aps*</code> |
| Log entries of topic1 and topic2 | <code>__topic__:topic1 or __topic__ : topic2</code> |
| Log entries with a tag whose key is tagkey1 and value is tagvalue2 | <code>__tag__ : tagkey1 : tagvalue2</code> |
| Log entries whose value of the latency field is greater than or equal to 100 and less than 200 | <code>latency>= 100 and latency < 200 or latency in [100 200)</code> |
| Log entries whose value of the latency field is greater than 100 | <code>latency > 100</code> |
| Log entries that do not contain spider and whose http_referer field does not contain opx | <code>not spider not bot not http_referer:opx</code> |
| Log entries whose cdnIP field is not empty | <code>not cdnIP:""</code> |
| Log entries that do not contain the cdnIP field | <code>not cdnIP:*</code> |
| Log entries that contain the cdnIP field | <code>cdnIP:*</code> |
| Log entries that contain a specified URL | <code>* select * where url = 'www.xxxxx.com'</code> |

Topic-specific search

Each Logstore is divided into one or more topics. You can divide a Logstore into multiple topics if you need level-2 categories of log entries. When you query logs, you can specify topics to increase efficiency.

In a search statement, you can specify one or more topics to query. If no topic is specified, log entries are queried from all topics.

For example, you can classify log entries into multiple topics based on domain names.

Log topics

| time | ip | method | url | host | topic |
|------------|-----------|--------|------------|--------------|-------|
| 1481270421 | 127.0.0.1 | POST | /users?u=1 | a.aliyun.com | siteA |
| 1481270422 | 127.0.0.1 | POST | /users?u=1 | a.aliyun.com | siteA |
| 1481270423 | 127.0.0.1 | POST | /users?u=1 | b.aliyun.com | siteB |
| 1481270424 | 127.0.0.1 | POST | /users?u=1 | b.aliyun.com | siteB |
| 1481270425 | 127.0.0.1 | POST | /users?u=1 | c.aliyun.com | siteC |
| 1481270426 | 127.0.0.1 | POST | /users?u=1 | c.aliyun.com | siteC |
| 1481270427 | 127.0.0.1 | POST | /users?u=1 | d.aliyun.com | siteD |
| 1481270428 | 127.0.0.1 | POST | /users?u=1 | d.aliyun.com | siteD |
| 1481270429 | 127.0.0.1 | POST | /users?u=1 | e.aliyun.com | siteE |
| 1481270430 | 127.0.0.1 | POST | /users?u=1 | e.aliyun.com | siteE |

Syntax of topic-specific search:

- In a search statement, you can specify one or more topics to query. If no topic is specified, log entries are queried from all topics.
- The topic-specific search syntax is `__topic__:topicName`. You can also specify a topic in a URL.
- You can query log entries from multiple topics. For example, the `__topic__:topic1 or __topic__:topic2` statement returns the log entries in topic1 and topic2.

23.4.7.2. LiveTail

This topic describes how to use LiveTail to monitor and analyze log data. LiveTail is an interactive feature provided in the Log Service console to monitor and extract key log data in real time.

Prerequisites

- LiveTail is available only after logs are collected.
- LiveTail can only monitor and extract log data collected by Logtail.

Context

In online O&M scenarios, you often need to monitor collected log data in real time and extract key information from the latest log data to locate error causes. In traditional O&M, you must run the `tail -f` command on servers to monitor log files in real time. To easily obtain the required real-time log information, you can include the `grep` or `grep -v` command to filter log entries by keyword. To simplify online O&M, Log Service provides LiveTail in the console to monitor and analyze online log data in real time.

Benefits

- Monitors real-time log information and filters log data by keyword.
- Logs collected based on the log collection configurations are identified by index.
- Log fields are delimited. This allows you to query contextual logs that contain delimiters.
- Log files can be tracked based on a single log entry and monitored in real time without the need to connect to online servers.

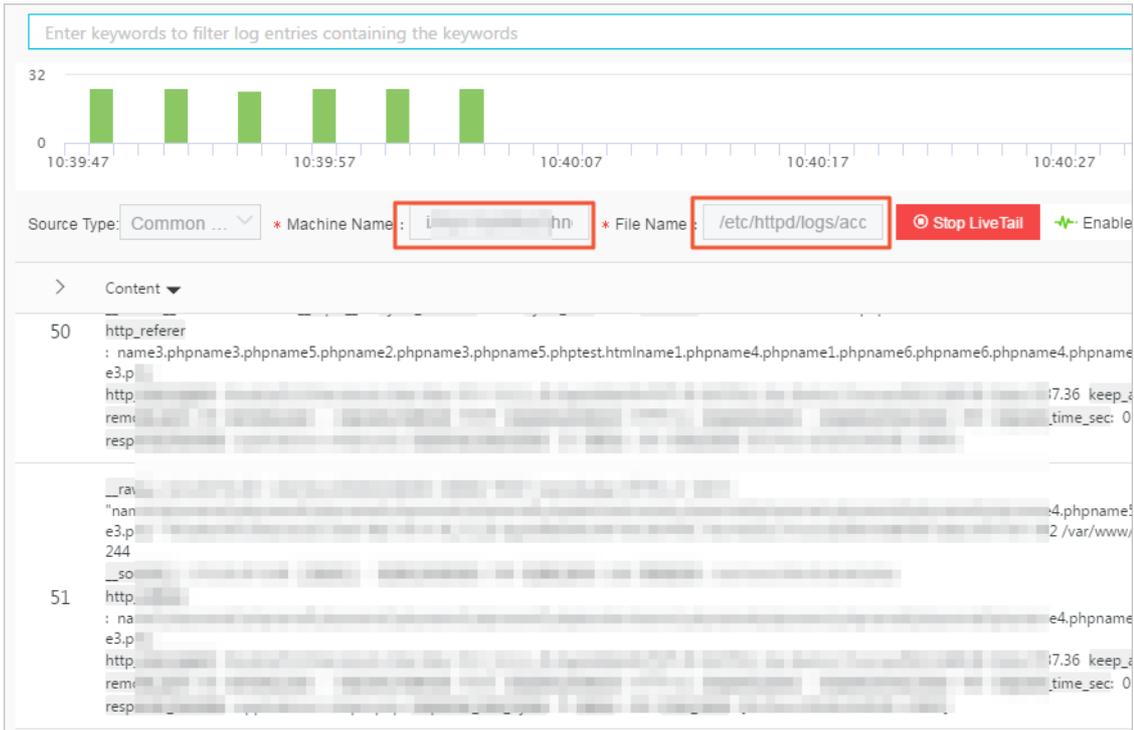
Use LiveTail to monitor logs in a Logstore in real time

1. [Log on to the Log Service console](#).
2. In the Projects section, click the target project.
3. Click the  icon next to the name of the Logstore, and then select **Search & Analysis**.
4. (Optional)Start LiveTail.

- i. On the **Raw Logs** tab, click the  icon next to the sequence number of the specified raw log entry, and then select **LiveTail**.

The system starts LiveTail and starts timing. The **Source Type**, **Machine Name**, and **File Name** fields are automatically filled in based on the specified raw log entry.

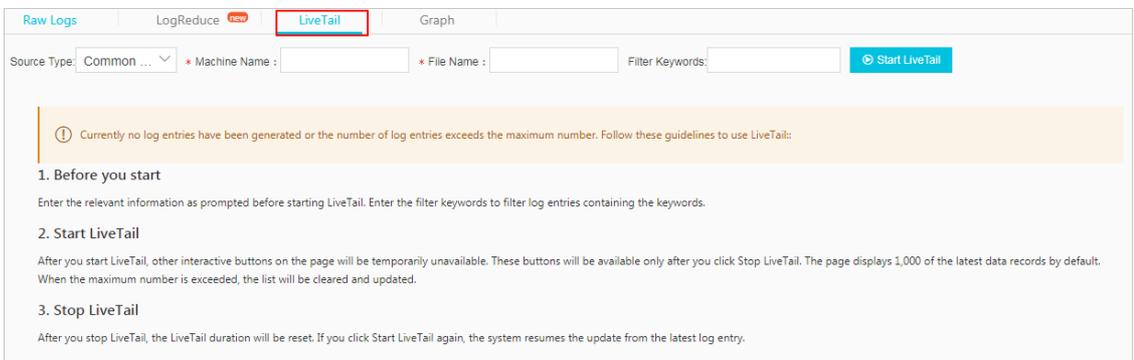
After LiveTail is started, log data collected by Logtail is displayed in order on the page in real time. The latest log data is displayed at the bottom of the page by default. You can view the latest log data without the need to drag the scroll bar. Up to 1,000 log entries can be displayed on the page. If more than 1,000 log entries are collected, the page is automatically refreshed to show the latest 1,000 log entries.



- ii. (Optional) You can also enter keywords in the search box to display log entries that contain the keywords in the monitoring list. By filtering log entries that contain the keyword, you can monitor specific log entries in real time.

5. Customize LiveTail.

- i. On the Search & Analysis page, click the **LiveTail** tab.



ii. Configure LiveTail.

| Parameter | Required | Description |
|-----------------|----------|--|
| Source Type | Yes | The source of log entries. Valid values: <ul style="list-style-type: none"> Physical servers Kubernetes containers Docker |
| Machine Name | Yes | The name of the server from which log entries are collected. |
| File Name | Yes | The full path and name of the log file. |
| Filter Keywords | No | A keyword. After you set a keyword, only log entries that contain the keyword are displayed in the log monitoring list. |

iii. Click Start LiveTail.

After LiveTail is started, log data collected by Logtail is displayed in order on the page in real time. The latest log data is displayed at the bottom of the page by default. You can view the latest log data without the need to drag the scroll bar. Up to 1,000 log entries can be displayed on the page. When more than 1,000 log entries are collected, the page is refreshed to show the latest 1,000 log entries.

6. To analyze logs during real-time log monitoring, click Stop LiveTail.

After you stop LiveTail, the LiveTail timing and the real-time log data update also stop.

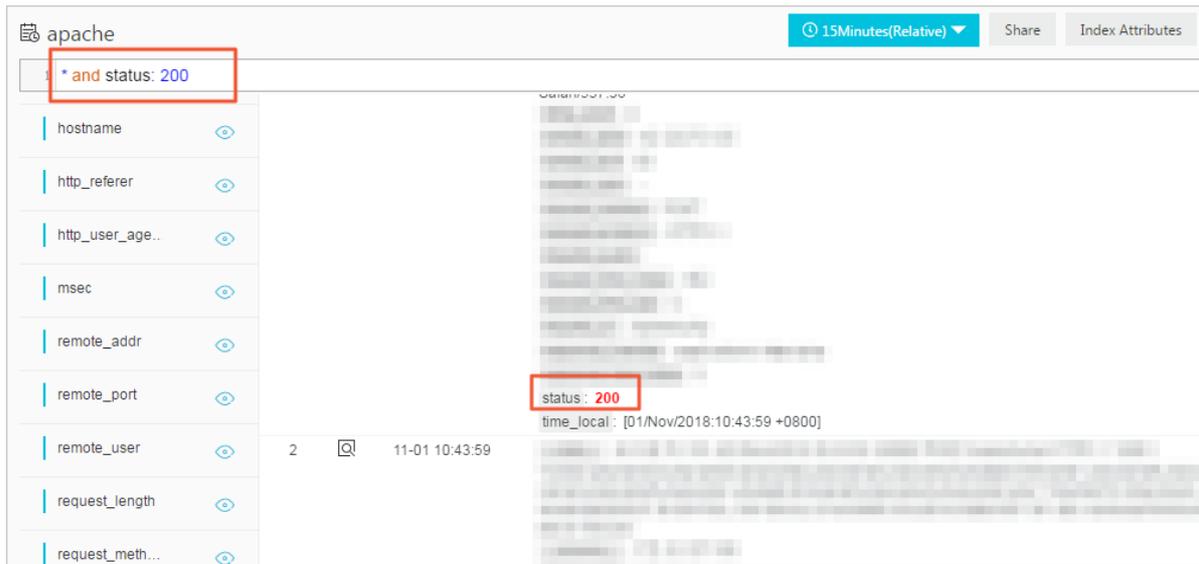
Log Service provides multiple methods to analyze exceptions that are found during log monitoring. For more information, see [Use LiveTail to analyze logs](#).

Use LiveTail to analyze logs

After you stop LiveTail, real-time log updates in the log monitoring list also stop. You can analyze and fix errors that are found during log monitoring.

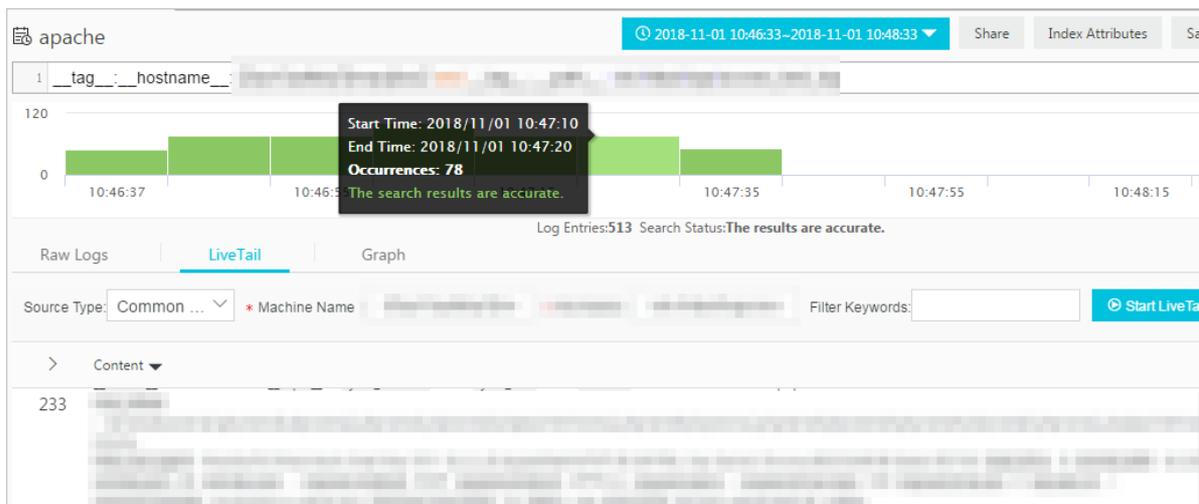
- View log entries that contain specified fields

Log fields are delimited. You can click the value of the specified exception field on the **LiveTail** tab. Then, you are automatically forwarded to the **Raw logs** tab and the value of the exception field is used as a keyword to filter all log entries that contain the field and the keyword. You can also analyze log entries that contain the keyword based on contextual queries and statistical charts.



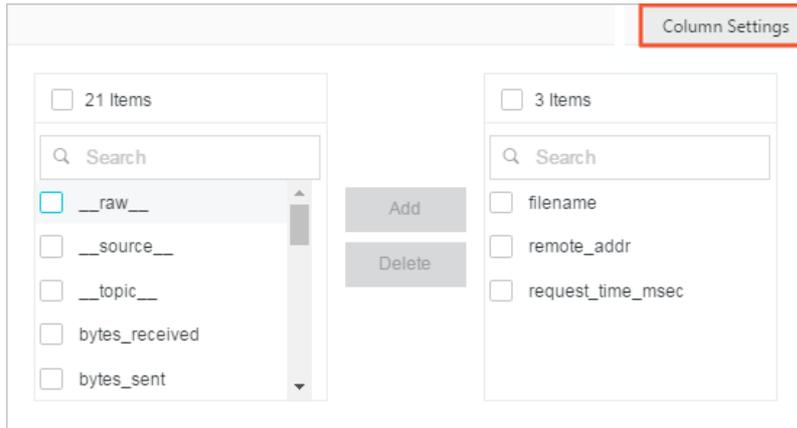
- Narrow down the time range of a log query based on the log distribution histogram

After LiveTail is started, the log distribution histogram is updated at the same time. If you find a distribution exception (for example, a big increase in the number of log entries) in a time range, you can click the corresponding green rectangle to narrow down the time range of the log query. The timeline of the raw logs is associated with the timeline that you click on the LiveTail tab. You can view all relevant raw logs and log distribution details during this time range.



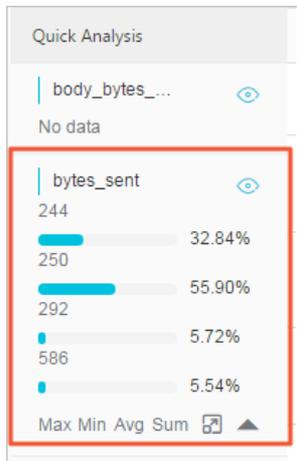
- Highlight key information based on column settings

On the LiveTail tab, click **Column Settings** in the upper-right corner of the log monitoring list. Then, you can specify a field as a separate column to highlight relevant data.



- Analyze log data

On the **LiveTail** tab, click the arrow in the upper-left corner of the log monitoring list to show the Quick Analysis pane. The time range for quick analysis starts from the time when LiveTail is started and ends at the time when LiveTail is stopped. The quick analysis provided on the LiveTail tab is the same as that provided on the Raw Logs tab. For more information, see [Quick analysis](#).



23.4.7.3. LogReduce

This topic describes how to use LogReduce to group log data in Log Service. The LogReduce feature groups similar log entries by detecting same log patterns during text log collection.

Context

The LogReduce feature allows you to group text logs of multiple formats. You can locate errors, detect anomalies, roll back versions, and perform other O&M operations in DevOps scenarios. You can also detect network intrusions to ensure data security. In addition, you can save the log grouping result to a dashboard as an analysis chart, and then view the grouped data in real time.

Benefits

- Various formats of logs such as Log4j logs, JSON-formatted logs, and syslog logs can be grouped.
- Hundreds of millions of log entries can be grouped in seconds.
- Log entries can be grouped in any pattern.
- Raw log entries can be retrieved based on the signature of log entries grouped in a pattern.
- The number of log entries grouped in a log pattern in different time ranges can be compared.
- The precision of log grouping can be adjusted based on your needs.

Billing method

After the LogReduce feature is enabled, the size of indexes increases by 10% of the raw log size. For example, if the size of raw log data is 100 GB per day, the size of log indexes increases by 10 GB.

| Raw log size | Index percentage | Size of indexes generated by LogReduce | Index size |
|--------------|------------------|--|------------|
| 100 GB | 20% (20 GB) | 100 × 10% | 30 GB |
| 100 GB | 40% (40 GB) | 100 × 10% | 50 GB |
| 100 GB | 100% (100 GB) | 100 × 10% | 110 GB |

Enable LogReduce of a Logstore

The LogReduce feature is disabled by default.

1. [Log on to the Log Service console](#).
2. Click the target project in the Projects section.
3. Click the  icon next to the name of the Logstore, and then select **Search & Analysis**.
4. Configure an index.
 - If you have enabled the index feature and configured indexes for the Logstore, choose **Index Attributes > Modify**.
 - If you have not enabled the index feature, click **Enable**.
5. Set the index attributes and turn on the **LogReduce** switch. Click **OK**.

After LogReduce is enabled, Log Service groups log data that has been collected. Then, you can perform the following operations:

- [View log grouping results and raw logs](#)
- [Adjust the precision of log grouping](#)
- [Compare the number of log entries grouped in different time ranges](#)

View log grouping results and raw logs

1. On the Search & Analysis page, enter a search and analytic statement in the search box, and then click **Search & Analytics**.

You can use keywords to filter grouped log entries.

 **Note** SQL statements are not supported. This means analysis results cannot be grouped.

2. Click the **LogReduce** tab to view the log grouping result.

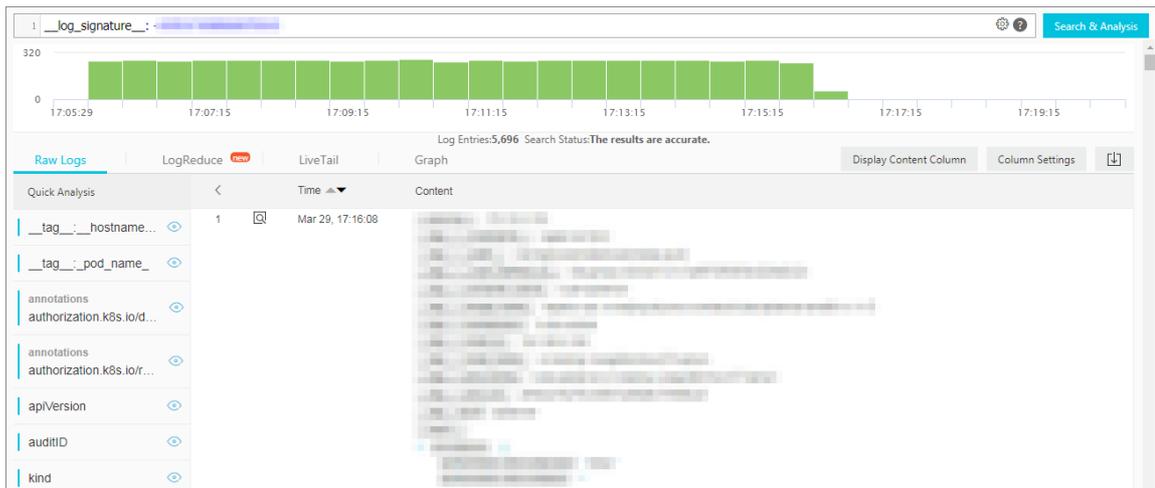
The filtered log grouping result is displayed on the **LogReduce** tab.

| Item | Description |
|----------------|---|
| Number | The sequence number of a log group. |
| Count | The number of log entries in a log group. |
| Pattern | The log pattern. Each log group has one or more sub-patterns. |

- Move the pointer over a value in the **Count** column to view the sub-patterns of the corresponding log group and the percentage of each sub-pattern in the log group. You can also click the plus sign (+) before the count value to expand the sub-pattern list.

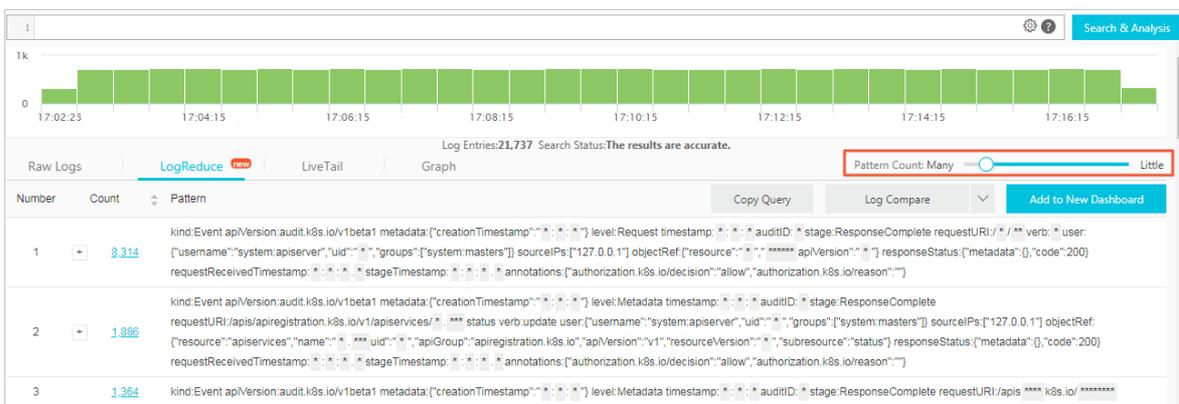


- o Click a count value to view the raw log entries of the corresponding log group.



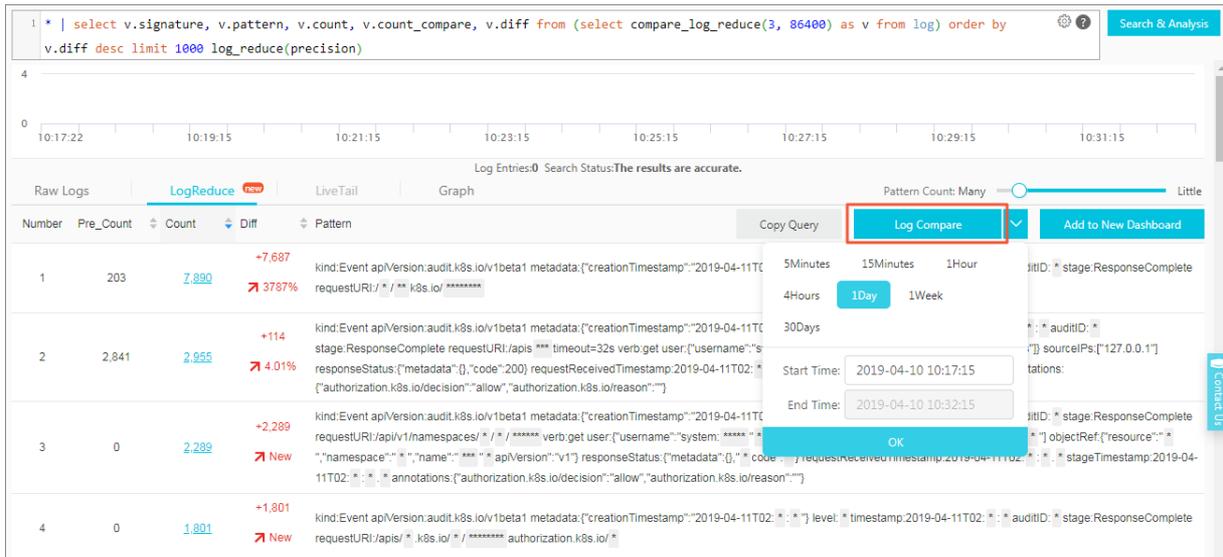
Adjust the precision of log grouping

1. On the Search & Analysis page, click the LogReduce tab.
2. In the upper-right corner of the tab, drag the Pattern Count slider to adjust the precision of log grouping.
 - o If you drag the slider towards Many, you can obtain a more precise log grouping result with more detailed patterns.
 - o If you drag the slider towards Little, you can obtain a less precise log grouping result with less detailed patterns.



Compare the number of log entries grouped in different time ranges

Click **Log Compare** on the **LogReduce** tab, select a time range, and then click **OK**.



| Item | Description |
|-----------|--|
| Number | The sequence number of a log group. |
| Pre_Count | The number of log entries grouped by the current pattern in the previous time range. |
| Count | The number of log entries grouped by the current pattern in the current time range. |
| Diff | The difference between the Pre_Count value and Count value. |
| Pattern | The pattern of a log group. |

SQL statement examples:

- Obtain a log grouping result.

- SQL statement:

```
* | select a.pattern, a.count, a.signature, a.origin_signatures from (select log_reduce(3) as a from log) limit 1000
```

Note When you view the log grouping result in the Log Service console, you can click **Copy Query** to obtain the relevant SQL statement.

- Input parameter: log_reduce (precision)

precision: an integer from 1 to 16 that can be set as the log grouping precision. A smaller value indicates a higher precision and more patterns. The default value is 3.

- Returned fields:

- pattern: the sub-patterns of log entries in a log group.
- count: the number of log entries in a log group.
- signature: the log pattern of a log group.
- origin_signatures: the original signature of a log group. You can use this field to query raw log entries of the log group.

- Compare the number of log entries grouped in different time ranges.

- SQL statement:

```
* | select v.pattern, v.signature, v.count, v.count_compare, v.diff from (select compare_log_reduce(3, 86400) as v from log) order by v.diff desc limit 1000
```

 **Note** After you click **Log Compare** in the Log Service console, you can click **Copy Query** to obtain the SQL statement.

- Input parameters: `compare_log_reduce(precision, compare_interval)`
 - `precision`: an integer from 1 to 16 that can be set as the log grouping precision. A smaller value indicates a higher precision and more patterns. The default value is 3.
 - `compare_interval`: the number of seconds between when the previous log entries and the current log entries were generated. The value of this parameter must be a positive integer.
- Returned fields:
 - `pattern`: the sub-patterns of log entries in a log group.
 - `signature`: the log pattern of a log group.
 - `count`: the number of log entries in a log group.
 - `count_compare`: the number of log entries for a same-pattern log group within the specified time range.
 - `Diff`: the difference between the values of the `count` field and the `count_compare` field.

23.4.7.4. Contextual query

This topic describes the contextual query feature provided in the Log Service console. You can use this feature to query the full context of the log file where specified log entries are obtained.

Prerequisites

The index feature is enabled.

Context

The contextual query feature identifies the server and file where a specified log entry resides. It then queries several log entries before and after the log entry in the original log file. This helps you locate errors during troubleshooting.

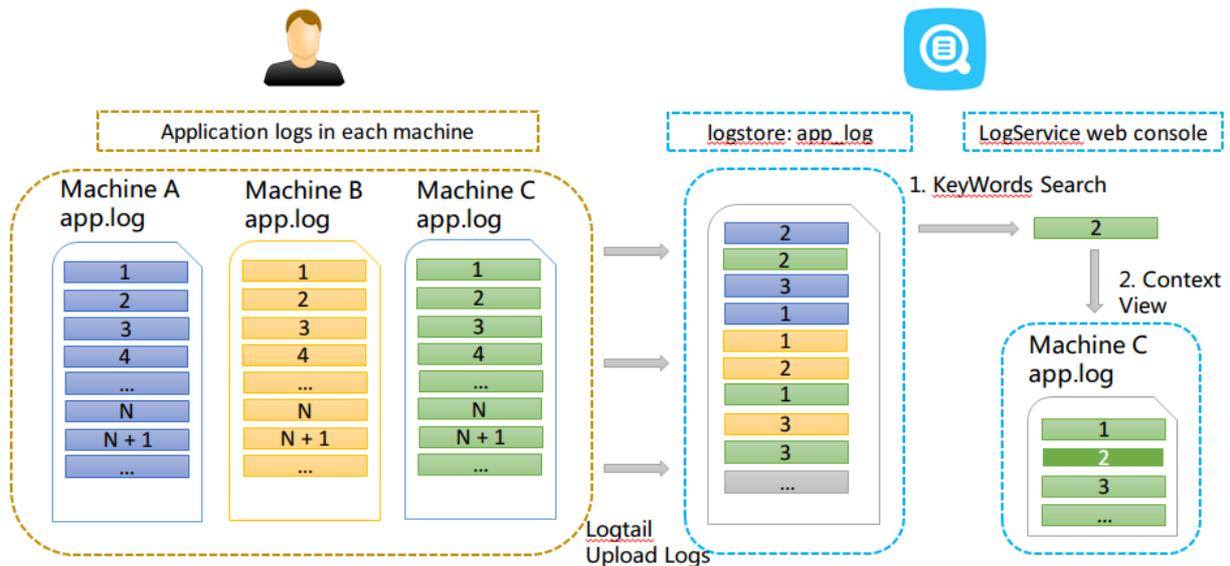
Scenarios

For example, a transaction on an O2O takeout website is logged in an application log file on a server as follows: **User login > Browse products > Select an item > Add to a shopping cart > Place an order > Payment > Deduction > Generate an order.**

If the order fails, the O&M personnel must locate the cause at the earliest opportunity. In traditional contextual queries, the O&M personnel must be authorized before logging on to each server where the application is deployed. Then, the O&M personnel must use the order ID as a keyword to search application log files to locate the cause of the failure.

In Log Service, the O&M personnel can perform the following steps to locate the cause of the failure:

1. Install Logtail on servers. Then log on to the Log Service console to add the servers to machine groups and configure log collection. After the configurations are complete, Logtail starts to upload incremental logs.
2. On the search and analysis page of the Log Service console, specify the time range and find the error log entry based on the order ID.
3. Based on the error log entry, page up until you find log entries related to the error log entry. For example, you may want to find a log entry that records a credit card payment failure.



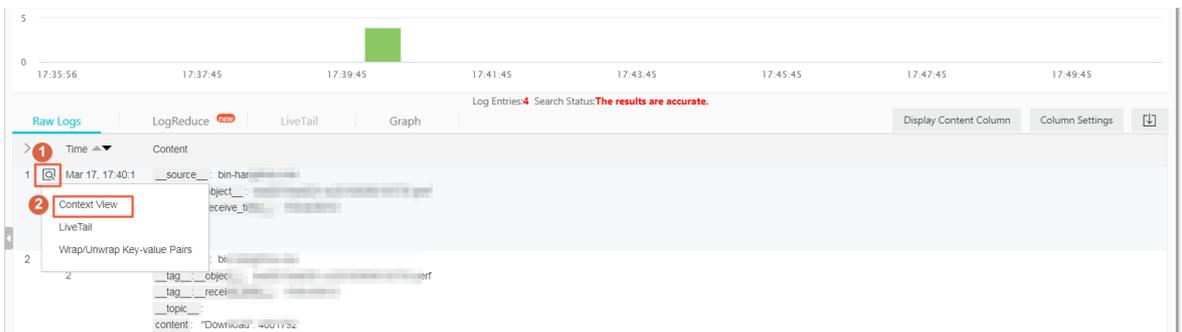
Note The contextual query feature does not support syslog.

Benefits

- Intrusions into applications or changes to log file formats are avoided.
- You can view contextual log entries of a specified log entry in a log file on a server that has been registered in the Log Service console. This helps you avoid logging on to each server to search for logs that you want.
- You can specify the time range based on the time when an event occurs to locate suspicious log entries. Then you can perform a contextual query in the Log Service console. This improves troubleshooting efficiency.
- Data loss caused by log file rotation or insufficient storage space is avoided. You can view historical log data in the Log Service console at any time.

Procedure

1. Log on to the Log Service console.
2. Click the target project in the Projects section.
3. Click the  icon next to the name of the Logstore, and then select **Search & Analysis**.
4. Enter a search and analytic statement, select a time range, and then click **Search & Analytics**.
On the query results page, if the **Context View** icon is available in the drop-down list of the icon to the left of a log entry, the log entry supports contextual query.



5. Click the icon to the left of a log entry, and select **Context View** from the drop-down list. On the page that appears, view the contextual log entries of the selected log entry.
6. Scroll up and down to view more contextual log entries. To view earlier or later contextual log entries, click

Old or New.

23.4.7.5. Saved search

This topic describes how to save a search and analytic statement as a saved search for a Logstore. The saved search feature allows you to search and analyze log data in an efficient way.

Prerequisites

The index feature is enabled and configured.

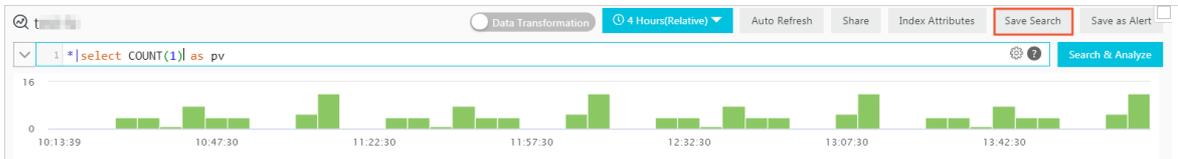
Context

If you need to frequently run a search and analytic statement, you can save the statement as a saved search. In later queries, you can click the name of the saved search on the left side of the search page to run the statement and view the result. You can also use the saved search in alert configurations. Log Service periodically runs the search and analytic statement and sends an alert if a query result meets the trigger condition.

If you want to select **Open Saved Search** in the Event Action field when you configure drill-down analysis, you must preset a saved search and set a **placeholder** in the query statement. For more information, see [Drill-down analysis](#).

Procedure

1. [Log on to the Log Service console](#).
2. Click the  icon next to the name of the Logstore, and then select **Search & Analysis**.
3. Enter a search and analytic statement in the search box, set a time range, and then click **Search & Analyze**.
4. Click **Save Search** in the upper-right corner of the page.



5. Configure the saved search.
 - i. Enter a **Saved Search Name**.
 - The name can contain only lowercase letters, digits, hyphens (-), and underscores (_).
 - The name must start and end with a lowercase letter or digit.
 - The name must be 3 to 63 characters in length.

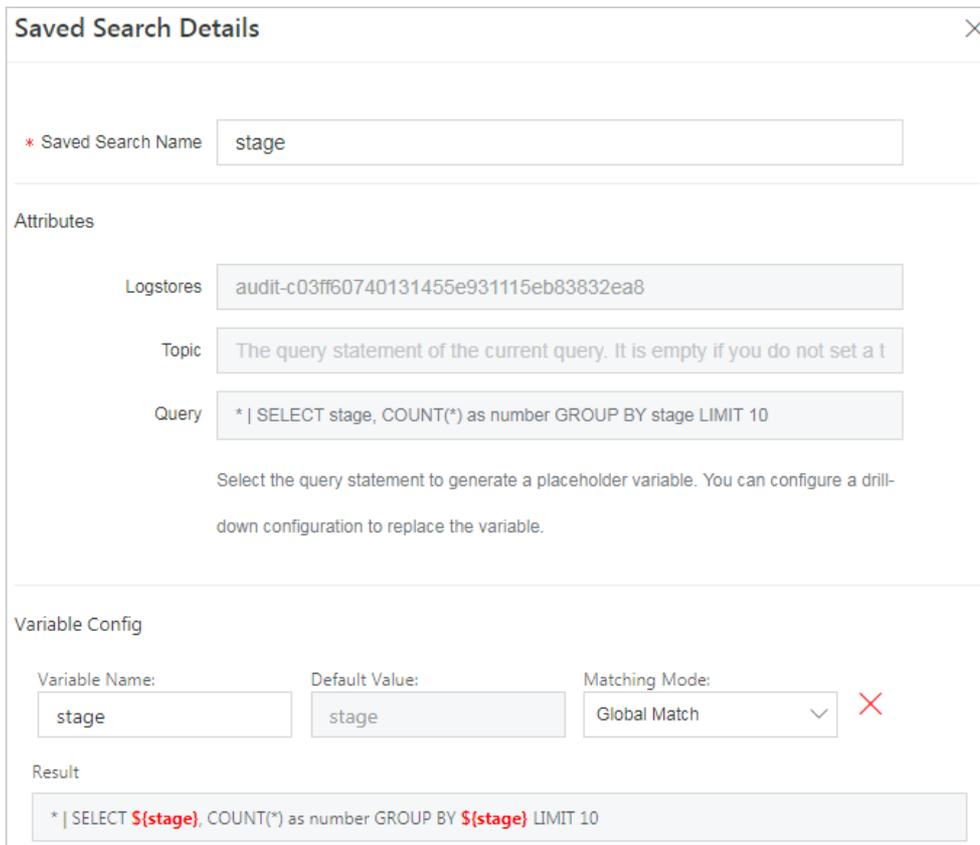
- ii. Check the values of the **Logstores**, **Topic**, and **Query** parameters.

If the values of the **Logstores** and **Topic** parameters do not meet your requirements, follow these steps: Return to the Logstores page. On this page, find and click the name of the target Logstore. On the page that appears, enter the search and analytic statement in the Search & Analyze search box, and then click **Save Search** again.

- iii. (Optional) Select a part of the query statement, and then click **Generate Variable**.

The generated variable is a placeholder variable. You can set the placeholder name in the **Variable Name** field. The selected characters are displayed in the **Default Value** field.

Note If you use this saved search for drill-down analysis in another chart where the **variable** is the same as the **Variable Name**, the **Default Value** is replaced with the chart value that you click to trigger the drill-down event. The search and analytic statement with the replaced chart value is executed. For more information, see [Drill-down analysis](#).

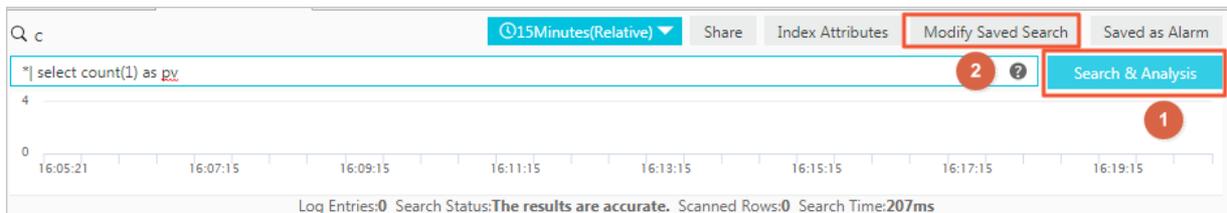


- 6. Click **OK**.

What's next

To modify a saved search, perform the following operations:

Enter a new search and analytic statement, click **Search & Analytics** to run the statement, and then click **Modify Saved Search**.



23.4.7.6. Quick analysis

This topic describes the quick analysis feature of Log Service. You can use this feature to query log data with one click. This feature allows you to analyze the distribution of a field in a specified time range and reduce the query cost of key data.

Features

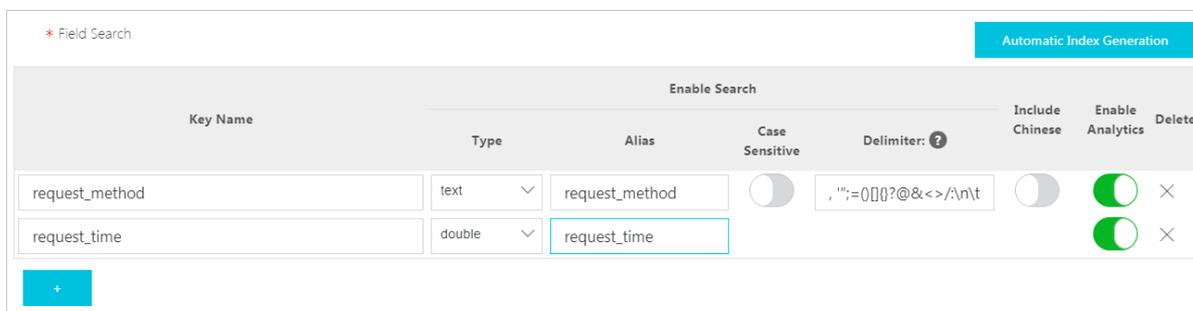
- Groups the first 100,000 values of a `text` field and provides statistics for the top 10 groups.
- Generates `approx_distinct` statements for `text` fields.
- Allows you to perform histogram-based statistics for the approximate distribution of `long` or `double` fields.
- Allows you to search for the maximum, minimum, and average of `long` or `double` fields and calculate the sum of the fields.
- Generates a query statement based on the quick analysis feature.

Prerequisites

Field indexes are configured.

- Indexes are configured for the fields that you need to search and analyze. For more information about how to enable the indexing feature, see [Enable the index feature and configure indexes for a Logstore](#).
- The name of a field is specified as the `key`. The data type, alias, and delimiter of the field are configured.

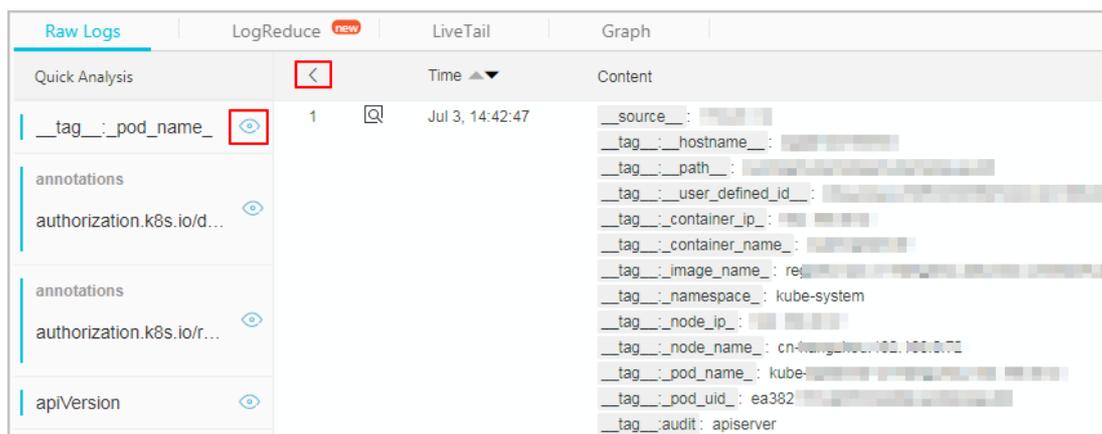
For example, if a log entry contains the `request_method` and `request_time` fields, you can configure indexes for the two fields, as shown in the following figure.



Instructions

After you configure indexes for specified fields, you can go to the Search & Analysis page and click the **Raw Logs** tab to view the specified fields. The fields are listed in the **Quick Analysis** pane on the left of the raw log entries. You can click the icon above the serial number to hide the Quick Analysis pane. You can also click the **Eye** icon to perform quick analysis based on the **Current Time Zone** and **Current Search** conditions.

Quick analysis



Data of the text type

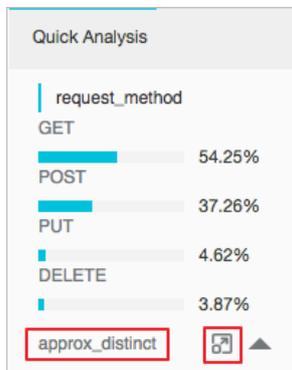
- Group and analyze log data

Click the **Eye** icon next to a `text` field to group the first 100,000 field values and return the percentages of the top 10 groups.

The following statement is used:

```
$Search | select ${keyName} , pv, pv *1.0/sum(pv) over() as percentage from( select count(1) as pv , "${keyName}" from (select "${keyName}" from log limit 100000) group by "${keyName}" order by pv desc) order by pv desc limit 10
```

The following figure shows the grouping and analytics result of the `request_method` field. `GET` requests account for the majority of requests.



- Calculate the number of unique values in a field

Under the target fields in the **Quick Analysis** pane, click **Count Distinct Values** to calculate the number of unique values in the `${keyName}` field.

- Fill the Search & Analyze search box with the grouping and analytics statement

Click the **Count Distinct Values** button on the right of the  icon. The Search & Analyze search box is filled with the grouping and analytics statement. You can edit the statement.

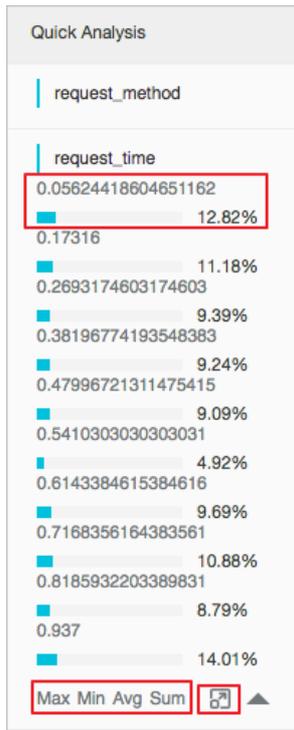
Data of the long and double types

- Display approximate distribution by using histograms

The number of `long` and `double` field values is large. The preceding grouping and analytics method is not suitable for the long and double data types. You can use the following statement to divide field values into 10 groups and display the approximate distribution of the values in a histogram:

```
$Search | select numeric_histogram(10, ${keyName})
```

The following figure shows the approximate distribution of the values in the `request_time` field. The largest percentage of request time is distributed around 0.059 seconds.



- Perform quick analysis by using the `Max`, `Min`, `Avg`, and `Sum` functions. You can click `Max` under a field to search for the maximum value, `Min` to search for the minimum value, `Avg` to calculate the average value, and `Sum` to calculate the sum of the values.
- Fill the Search & Analyze search box with the query statement of the histogram approximate distribution. Click the  icon next to `Sum`. The Search & Analyze search box is filled with the query statement of the histogram approximate distribution. You can edit the statement.

23.4.8. Analysis grammar

23.4.8.1. General aggregate functions

This topic describes the syntax and examples of general aggregate functions.

The search and analytics feature of Log Service allows you to use general aggregate functions for log analysis. The following table describes the supported general aggregate functions.

| Function | Description | Example |
|---------------------------|---|--|
| <code>arbitrary(x)</code> | Returns a random value from among the values in the x field. | <code>latency > 100 select arbitrary(method)</code> |
| <code>avg(x)</code> | Returns the arithmetic mean of all values in the x field. | <code>latency > 100 select avg(latency)</code> |
| <code>checksum(x)</code> | Returns a Base64-encoded checksum of the values of the x field. | <code>latency > 100 select checksum(method)</code> |
| <code>count(*)</code> | Calculates the number of values of a field. | - |

| Function | Description | Example |
|---------------------------------|--|---|
| <code>count(x)</code> | Counts the number of non-null values of the x field. | <code>latency > 100 select count(method)</code> |
| <code>count(digit)</code> | Counts the number of values of a field. The <code>count(digit)</code> function is equivalent to the <code>count(1)</code> and <code>count(*)</code> functions. | - |
| <code>count_if(x)</code> | Returns the number of the occurrences of the TRUE value. | <code>latency > 100 select count_if(url like '%abc')</code> |
| <code>geometric_mean(x)</code> | Returns the geometric mean of the values in the x field. | <code>latency > 100 select geometric_mean(latency)</code> |
| <code>max_by(x,y)</code> | Returns the value of the x field that corresponds to the maximum value of the y field. | <code>* select min_by(method,latency,3)</code> : queries the method that corresponds to the maximum latency. |
| <code>max_by(x,y,n)</code> | Returns n values of the x field that corresponds to the n largest value of the y field. | <code>* select min_by(method,latency,3)</code> : queries the three methods that corresponds to the three maximum latencies. |
| <code>min_by(x,y)</code> | Returns the value of the x field that corresponds to the smallest value of the y field. | <code>* select min_by(x,y)</code> : queries the method that corresponds to the minimum latency. |
| <code>min_by(x,y,n)</code> | Returns n values of the x field that corresponds to the n smallest values of the y field. | <code>* select min_by(method,latency,3)</code> : queries the three methods that corresponds to the three minimum latencies. |
| <code>max(x)</code> | Returns the maximum value among all values in the x field. | <code>latency > 100 select max(inflow)</code> |
| <code>min(x)</code> | Returns the minimum value among all values in the x field. | <code>latency > 100 select min(inflow)</code> |
| <code>sum(x)</code> | Returns the sum among all values in the x field. | <code>latency > 10 select sum(inflow)</code> |
| <code>bitwise_and_agg(x)</code> | Returns the bitwise AND of all values in the x field in two's complement representation. | - |
| <code>bitwise_or_agg(x)</code> | Returns the bitwise OR of all values in the x field in two's complement representation. | - |

23.4.8.2. Security check functions

Security check functions in Log Services are designed based on the globally shared WhiteHat Security asset library. This topic describes security check functions that you can use to check whether an IP address, domain name, or URL in logs is secure.

Scenarios

- O&M personnel of enterprises and institutions in Internet, gaming, information, and other industries that require robust O&M services can use security check functions to identify suspicious requests or attacks. They can also use the functions to implement in-depth analysis and defend against potential attacks.
- O&M personnel of enterprises and institutions in banking, securities, e-commerce, and other industries that require strong protection for internal assets can use security check functions to identify requests to suspicious websites and downloads initiated by trojans. Then the O&M personnel can take immediate actions to prevent potential losses.

Features

- Reliability: built upon the globally shared WhiteHat Security asset library that is updated in a timely manner.
- Efficiency: capable of screening millions of IP addresses, domain names, and URLs within seconds.
- Ease of use: supports the analysis of network logs by using the security_check_ip, security_check_domain, and security_check_url functions.
- Flexibility: supports interactive queries, report creation, and alert configurations and subsequent actions.

Functions

| Function | Description | Example |
|-----------------------|---|--|
| security_check_ip | <p>Checks whether an IP address is secure.</p> <ul style="list-style-type: none"> • The value 1 indicates that the specified IP address is suspicious. • The value 0 indicates that the specified IP address is secure. | <pre>select security_check_ip(real_client_ip)</pre> |
| security_check_domain | <p>Checks whether a domain name is secure.</p> <ul style="list-style-type: none"> • The value 1 indicates that the specified domain name is suspicious. • The value 0 indicates that the specified domain name is secure. | <pre>select security_check_domain(site)</pre> |
| security_check_url | <p>Checks whether a URL is secure.</p> <ul style="list-style-type: none"> • The value 1 indicates that the specified URL is suspicious. • The value 0 indicates that the specified URL is secure. | <pre>select security_check_domain(concat(host ,url))</pre> |

Examples

- Check external suspicious requests and generate reports

For example, an e-commerce enterprise collects logs from its NGINX servers and wants to scan suspicious client IP addresses. To do this, the enterprise can pass the ClientIP field in logs that are collected from the NGINX servers to the security_check_ip function and filter out IP addresses associated with the returned value 1. Then the enterprise can query the countries where the IP addresses are located and ISPs to which the IP addresses belong.

SQL statement for this scenario:

```
* | select ClientIP, ip_to_country(ClientIP) as country, ip_to_provider(ClientIP) as provider, count(1) as PV where security_check_ip(ClientIP) = 1 group by ClientIP order by PV desc
```

Display the ISPs and countries in a map.

| client_ip | country | provider | PV |
|-----------|---------|----------|----|
| 180 | CN | E | 3 |
| 103 | CN | | 3 |
| 180 | CN | E | 1 |

- Check internal suspicious requests and send alerts

For example, a securities operator collects logs of its internal devices that access the Internet through gateways. To check requests to suspicious websites, the operator can run the following statement:

```
* | select client_ip, count(1) as PV where security_check_ip(remote_addr) = 1 or security_check_site(site) = 1 or security_check_url(concat(site, url)) = 1 group by client_ip order by PV desc
```

The operator can save this statement as a saved search and configure an alert. An alert is triggered when a client frequently accesses suspicious websites. The statement in the alert can be configured to run every five minutes to check if a client has frequently (more than five times) accessed suspicious websites in the past one hour. The following figure shows the configurations of an alert.

Create Alert

Alert Configuration
Notifications

* Alert Name 5/64

* Add to New Dashboard Create 12/64

* Chart Name 5/64

Query

* | select client_ip, count(1) as PV where security_check_ip(remote_addr) = 1 or security_check_site(site) = 1 or security_check_url(concat(site, url)) = 1 group by client_ip order by PV desc

* Search Period 🕒 15 Minutes(Relative)

* Check Frequency Fixed Interval + - Minutes

* Trigger Condition 4/128

Five basic operators are supported: plus (+), minus (-), multiplication (*), division (/), and modulo (%). Eight comparison operators are supported: greater than (>), greater than or equal to (>=), less than (<), less than or equal to (<=), equal to (==), not equal to (!=), regex match (=~), and negated regex match (!~). [Documentation](#)

Advanced >

Next
Cancel

23.4.8.3. Map functions

This topic describes the syntax and examples of map functions.

The following table describes the supported map functions.

| Function | Description | Example |
|--|---|---|
| Subscript operator [] | Retrieves the value corresponding to a specified key from a map. | - |
| histogram(x) | Groups the values of the parameter x and calculates the number of occurrences of each value. The syntax is equivalent to <code>select count group by x</code> . Note The returned data is in the JSON format. | The statement <code>latency > 10 select histogram(status)</code> is equivalent to the statement <code>latency > 10 select count(1) group by status</code> . |
| histogram_u(x) | Groups the values of the parameter x and calculates the number of occurrences of each value. Note The returned data is in the format of multiple rows and columns. | The statement <code>latency > 10 select histogram_u(status)</code> is equivalent to the statement <code>latency > 10 select count(1) group by status</code> . |
| map_agg(Key,Value) | Returns a random value of the key in the format of a map that consists of key-value pairs. | <code>latency > 100 select map_agg(method,latency)</code> |
| multimap_agg(Key,Value) | Returns all values of the key in the format of a map that consists of key-value pairs. | <code>latency > 100 select multimap_agg(method,latency)</code> |
| cardinality(x) → bigint | Returns the cardinality of the map x. | - |
| element_at(map <K, V> , key) → V | Returns the value for the specified key. | - |
| map() → map <unknown, unknown> | Returns an empty map. | - |
| map(array <K> , array <V>) → map <K,V> | Returns a map where each key-value pair consists of two elements from two separate arrays. | <code>SFI FCT map(ARRAY[1,3], ARRAY[2,4]); - {1 -> 2, 3 -> 4}</code> |
| map_from_entries(array <row<K, V>>) → map <K,V> | Converts a multi-dimensional array to a map. | <code>SFI FCT map from entries(ARRAY[(1, 'x'), (2, 'y')]); - {1 -> 'x', 2 -> 'y'}</code> |

| Function | Description | Example |
|---|--|---------|
| <code>map_concat(map1 <K, V> , map2 <K, V> , ..., mapN <K, V>) → map <K, V></code> | Returns a map that is the union of all specified maps. If a key is found in multiple specified maps, the value of the key in the returned map is the value of the key that occurs in the last specified map. | - |
| <code>map_filter(map <K, V> , function) → map <K, V></code> | For more information, see the <code>map_filter</code> function in Lambda functions . | - |
| <code>map_keys(x <K, V>) → array <K></code> | Returns an array of keys in the specified map. | - |
| <code>map_values(x <K, V>) → array <V></code> | Returns an array of values in the specified map. | - |

23.4.8.4. Approximate functions

This topic describes the syntax and examples of approximate functions that Log Service supports for log analysis.

The following table describes the supported approximate functions.

| Function | Description | Example |
|--|--|---|
| <code>approx_distinct(x)</code> | Returns the approximate number of distinct values of the x field. | - |
| <code>approx_percentile(x,percentage)</code> | Returns the value located at the specified approximate percentage among the sorted values of the x field. | <code>approx_percentile(x,0.5)</code> : returns the median among the sorted values of the x field. |
| <code>approx_percentile(x, percentages)</code> | Returns values located at multiple specified approximate percentages among the sorted values of the x field. This function works in a similar manner to the <code>approx_percentile(x,percentage)</code> function. | <code>approx_percentile(x,array[0.1,0.2])</code> |
| <code>numeric_histogram(buckets, Value)</code> | Distributes all values of the <i>Value</i> field into multiple buckets. The <i>buckets</i> parameter specifies the number of buckets. The key of every bucket and the number of values in a bucket are returned. This function is equivalent to <code>select count group by</code> . Note The response is in the JSON format. | <code>method:POST select numeric_histogram(10,latency)</code> : distributes the latencies of POST requests into 10 buckets and calculates the number of latencies in each bucket. |

| Function | Description | Example |
|--|--|--|
| <code>numeric_histogram_u(buckets, Value)</code> | <p>Distribute all values of the <i>Value</i> field into multiple buckets. The <i>buckets</i> parameter specifies the number of buckets.</p> <p>The key of every bucket and the number of values in a bucket are returned. This function is equivalent to <code>select count group by</code>.</p> <p>Note The returned data is in the format of multiple rows and columns.</p> | <pre>method:POST select numeric_histogram(10,latency) :</pre> <p>distributes the latency data of POST requests into 10 buckets and calculates the number of latency data in each bucket.</p> |

Note The number of values in every bucket is evenly distributed. It is returned along with the average value of all values in a bucket.

23.4.8.5. Mathematical statistics functions

This topic describes the syntax and examples of mathematical statistics functions.

The search and analytics feature of Log Service allows you to use mathematical statistics functions for log analysis. The following table describes the supported mathematical statistics functions.

| Function | Description | Example |
|---|--|---|
| <code>corr(y, x)</code> | Returns the correlation coefficient of input values. The result ranges from 0 to 1. | <pre>latency>100 select corr(latency,request_size)</pre> |
| <code>covar_pop(y, x)</code> | Returns the population covariance of input values. | <pre>latency>100 select covar_pop(request_size,latency)</pre> |
| <code>covar_samp(y, x)</code> | Returns the sample covariance of input values. | <pre>latency>100 select covar_samp(request_size,latency)</pre> |
| <code>regr_intercept(y, x)</code> | Returns the linear regression intercept of input values. <i>y</i> is the dependent value. <i>x</i> is the independent value. | <pre>latency>100 select regr_intercept(request_size,latency)</pre> |
| <code>regr_slope(y, x)</code> | Returns the linear regression slope of input values. <i>y</i> is the dependent value. <i>x</i> is the independent value. | <pre>latency>100 select regr_slope(request_size,latency)</pre> |
| <code>stddev(x)</code> or <code>stddev_samp(x)</code> | Returns the sample standard deviation of the values in the <i>x</i> field. | <pre>latency>100 select stddev(latency)</pre> |
| <code>stddev_pop(x)</code> | Returns the population standard deviation of the values in the <i>x</i> field. | <pre>latency>100 select stddev_pop(latency)</pre> |
| <code>variance(x)</code> or <code>var_samp(x)</code> | Returns the sample variance of the values in the <i>x</i> field. | <pre>latency>100 select variance(latency)</pre> |

| Function | Description | Example |
|-------------------------|---|--|
| <code>var_pop(x)</code> | Returns the population variance of the values in the x field. | <code>latency>100 select variance(latency)</code> |

23.4.8.6. Mathematical calculation functions

This topic describes the syntax and examples of mathematical calculation functions.

By including mathematical calculation functions in SQL statements, you can perform mathematical calculation on log query results.

Mathematical operators

Mathematical operators include the plus sign (+), minus sign (-), multiplication sign (*), division sign (/), and percent sign (%). These operators can be used in SELECT statements.

Example:

```
*|select avg(latency)/100 , sum(latency)/count(1)
```

Mathematical calculation functions

Log Service supports the following mathematical calculation functions.

| Function | Description |
|--------------------------------------|--|
| <code>abs(x)</code> | Returns the absolute values of the values in the x field. |
| <code>cbrt(x)</code> | Returns the cube roots of the values in the x field. |
| <code>ceiling(x)</code> | Returns the rounded-up nearest integers of the values in the x field. |
| <code>cosine_similarity(x,y)</code> | Returns the cosine similarity between the sparse vectors x and y. |
| <code>degrees</code> | Converts angles in radians to degrees. |
| <code>e()</code> | Returns the Euler's number. |
| <code>exp(x)</code> | Returns Euler's number raised to the power of the values in the x field. |
| <code>floor(x)</code> | Returns the rounded-down nearest integers of the values in the x field. |
| <code>from_base(string,radix)</code> | Returns the radix number representation of a string. |
| <code>ln(x)</code> | Returns the natural logarithm of x. |
| <code>log2(x)</code> | Returns the base 2 logarithm of x. |
| <code>log10(x)</code> | Returns the base 10 logarithm of x. |
| <code>log(x,b)</code> | Returns the base b logarithm of x. |
| <code>pi()</code> | Returns the constant Pi. |

| Function | Description |
|-------------------------------|---|
| <code>pow(x,b)</code> | Returns x raised to the power of b. |
| <code>radians(x)</code> | Converts angle x in degrees to radians. |
| <code>rand()</code> | Returns a random number. |
| <code>random(0,n)</code> | Returns a random number from 0 to n (exclusive). |
| <code>round(x)</code> | Returns x rounded to the nearest integer. |
| <code>round(x,y)</code> | Returns x rounded to y decimal places. For example, <code>round(1.012345,2) = 1.01</code> . |
| <code>sqrt(x)</code> | Returns the square root of x. |
| <code>to_base(x,radix)</code> | Returns the radix number representation of x. |
| <code>truncate(x)</code> | Returns x rounded to integer by dropping digits after the decimal point. |
| <code>acos(x)</code> | Returns the arc cosine of x. |
| <code>asin(x)</code> | Returns the arc sine of x. |
| <code>atan(x)</code> | Returns the arc tangent of x. |
| <code>atan2(y,x)</code> | Returns the arc tangent of y/x. |
| <code>cos(x)</code> | Returns the cosine of x. |
| <code>sin(x)</code> | Returns the sine of x. |
| <code>cosh(x)</code> | Returns the hyperbolic cosine of x. |
| <code>tan(x)</code> | Returns the tangent of x. |
| <code>tanh(x)</code> | Returns the hyperbolic tangent of x. |
| <code>infinity()</code> | Returns the value representing positive infinity. |
| <code>is_infinity(x)</code> | Determines whether x is infinite. |
| <code>is_finity(x)</code> | Determines whether x is finite. |
| <code>is_nan(x)</code> | Determines whether x is not-a-number. |

23.4.8.7. String functions

This topic describes string functions that Log Service supports for log data search and analytics.

This following table lists the functions and their descriptions.

| Function | Description |
|---------------------|--|
| <code>chr(x)</code> | Converts an integer to an ASCII character. For example, the result of <code>chr(65)</code> is <code>A</code> . |

| Function | Description |
|--|--|
| <code>codepoint(x)</code> | Converts an ASCII character to a code point of the integer type. For example, the result of <code>codepoint('A')</code> is <code>65</code> . |
| <code>length(x)</code> | Returns the length of a string. |
| <code>levenshtein_distance(string1, string2)</code> | Returns the Levenshtein distance of string1 and string2. |
| <code>lower(string)</code> | Converts all uppercase characters in a string into lowercase characters. |
| <code>lpad(string, size, padstring)</code> | Pads a string to the specified size. If the length of the string is shorter than the specified size, padstring is used to left pad the string. If the length of the string is longer than the specified size, the string is truncated with the specified size. |
| <code>rpadd(string, size, padstring)</code> | Right pads a <code>string</code> with the specified padding. The implementation is similar to that of the <code>lpad</code> function. |
| <code>ltrim(string)</code> | Removes whitespace characters from the left side of a string. |
| <code>replace(string, search)</code> | Removes all occurrences of a substring search from a string. |
| <code>replace(string, search, rep)</code> | Replaces all occurrences of a substring search in a string with another substring rep. |
| <code>reverse(string)</code> | Reverses a string. |
| <code>rtrim(string)</code> | Removes whitespace characters from the right side of a string. |
| <code>split(string, delimiter, limit)</code> | Splits a string based on the specified delimiter and returns an array with the maximum number of elements at limit. The index of the first element in the array is 1. |
| <code>split_part(string, delimiter, offset)</code> | Splits a string based on a delimiter into an array of substrings and returns the element with the index specified by the offset parameter. |
| <code>split_to_map(string, entryDelimiter, keyValueDelimiter)</code> → <code>map<varchar, varchar></code> | Splits a string based on the entryDelimiter into multiple entries, each of which is then split based on the keyValueDelimiter into a key and value. This function returns a map. |
| <code>position(substring IN string)</code> | Returns the position of the first occurrence of the specified substring in a string. |
| <code>strpos(string, substring)</code> | Returns the position of the first occurrence of the specified substring in a string. Positions start with 1. If the substring is found, 0 is returned. |
| <code>substr(string, start)</code> | Returns a substring from the start position. Positions start with 1. |

| Function | Description |
|---|---|
| <code>substr(string, start, length)</code> | Returns a substring of a specified length from start position. Positions start with 1. The length parameter specifies the length of the substring returned. |
| <code>trim(string)</code> | Removes leading and trailing whitespace characters from a string. |
| <code>upper(string)</code> | Converts all lowercase characters in a string into uppercase characters. |
| <code>concat(string,string...)</code> | Concatenates multiple strings into a single string. |
| <code>hamming_distance (string1,string2)</code> | Returns the Hamming distance of string1 and string2. |

Note Strings must be enclosed in single quotation marks ('). Double quotation marks (") are used to enclose field names. For example, `a='abc'` indicates `field a = string 'abc'`, and `"a"="abc"` indicates `field a = field abc`.

23.4.8.8. Date and time functions

This topic describes the available time functions, date functions, interval functions, and a time series padding function in Log Service. You can use these functions when you analyze data.

Date and time data types

- Unix timestamp: specifies the number of seconds that have elapsed since 1970-01-01 00:00:00 UTC. The value is in the format of an integer. For example, `1512374067` indicates `Mon Dec 4 15:54:27 CST 2017`. The `__time__` field in every log entry is of this type.
- Timestamp: specifies the date and time in the format of a string. For example, `2017-11-01 13:30:00`.

Date functions

The following table lists the commonly used date functions that are supported in Log Service.

| Function | Description | Example |
|---|---|---|
| <code>current_date</code> | Returns the current date. | <code>latency>100 select current_date</code> |
| <code>current_time</code> | Returns the current time. | <code>latency>100 select current_time</code> |
| <code>current_timestamp</code> | Returns the current timestamp. This function is equivalent to the combination of the <code>current_date</code> and <code>current_time</code> functions. | <code>latency>100 select current_timestamp</code> |
| <code>current_timezone()</code> | Returns the current time zone. | <code>latency>100 select current_timezone()</code> |
| <code>from_iso8601_timestamp(string)</code> | Parses an ISO 8601 formatted string into a timestamp that specifies the time zone. | <code>latency>100 select from_iso8601_timestamp(iso8601)</code> |
| <code>from_iso8601_date(string)</code> | Parses an ISO 8601 formatted string into a date. | <code>latency>100 select from_iso8601_date(iso8601)</code> |

| Function | Description | Example |
|---|--|---|
| <code>from_unixtime(unixtime)</code> | Parses a Unix timestamp into a timestamp. | <code>latency>100 select from_unixtime(1494985275)</code> |
| <code>from_unixtime(unixtime,string)</code> | Parses a Unix timestamp into a timestamp based on the time zone that is specified by the string parameter. | <code>latency>100 select from_unixtime(1494985275,'Asia/Shanghai')</code> |
| <code>localtime</code> | Returns the local time. | <code>latency>100 select localtime</code> |
| <code>localtimestamp</code> | Returns the local timestamp. | <code>latency>100 select localtimestamp</code> |
| <code>now()</code> | Returns the current date and time. This function is equivalent to the <code>current_timestamp</code> function. | N/A |
| <code>to_unixtime(timestamp)</code> | Parses a timestamp into a Unix timestamp. | <code>* select to_unixtime('2017-05-17 09:45:00.848 Asia/Shanghai')</code> |

Time functions

The following table lists the time functions that you can use in Log Service to parse time in the formats supported in MySQL, such as %a, %b, and %y.

| Function | Description | Example |
|---|--|---|
| <code>date_format(timestamp, format)</code> | Formats a timestamp in the specified format. | <code>latency>100 select date_format(date_parse('2017-05-17 09:45:00','%Y-%m-%d %H:%i:%S'), '%Y-%m-%d')</code> |
| <code>date_parse(string, format)</code> | Parses a formatted string into a timestamp. | <code>latency>100 select date_format(date_parse(time,'%Y-%m-%d %H:%i:%S'), '%Y-%m-%d')</code> |

Formats

| Format | Description |
|--------|---|
| %a | The week day in abbreviation, such as Sun and Sat. |
| %b | The month in abbreviation, such as Jan and Dec. |
| %c | The month in the numeric format. Valid values: 1 to 12. |
| %D | The day of the month with a suffix, such as 0th, 1st, 2nd, and 3rd. |
| %d | The day of the month in the numeric format. Valid values: 01 to 31. |
| %e | The day of the month in the numeric format. Valid values: 1 to 31. |
| %H | The hour that applies the 24-hour clock convention. Valid values: 00 to 23. |
| %h | The hour that applies the 12-hour clock convention. Valid values: 01 to 12. |
| %l | The hour that applies the 12-hour clock convention. Valid values: 1 to 12. |

| Format | Description |
|--------|--|
| %i | The minute in the numeric format. Valid values: 00 to 59. |
| %j | The day of the year. Valid values: 001 to 366. |
| %k | The hour. Valid values: 0 to 23. |
| %l | The hour. Valid values: 1 to 12. |
| %M | The month. Valid values: January to December. |
| %m | The month in the numeric format. Valid values: 01 to 12. |
| %p | The abbreviation that indicates the morning or afternoon. Valid values: a.m. and p.m.. |
| %r | The time that applies the 12-hour clock convention: <code>hh:mm:ss AM/PM</code> . |
| %S | The second. Valid values: 00 to 59. |
| %s | The second. Valid values: 00 to 59. |
| %T | The time that applies the 24-hour clock convention, formatted in <code>hh:mm:ss</code> . |
| %U | The week of the year. Sunday is the first day of a week. Valid values: 00 to 53. |
| %u | The week of the year. Monday is the first day of a week. Valid values: 00 to 53. |
| %V | The week of the year. Sunday is the first day of a week. This format is used together with %X. Valid values: 01 to 53. |
| %v | The week of the year. Monday is the first day of a week. This format is used in combination with %x. Valid values: 01 to 53. |
| %W | The day of the week. Valid values: Sunday to Saturday. |
| %w | The day of the week. Valid values: 0 to 6. The value 0 indicates Sunday. |
| %Y | The year in the 4-digit format. |
| %y | The year in the 2-digit format. |
| %% | Escapes the second percent sign (%). |

Truncation functions

Log Service supports a truncation function, which can truncate a time by the second, minute, hour, day, month, or year. Typically, this function is used for time-based analytics.

- Syntax

```
date_trunc(unit, x)
```

- Parameters

The value of the x parameter can be a timestamp or Unix timestamp.

The following table lists the values of the unit parameter and the responses when the x parameter is set to `2001-08-22 03:04:05.000` .

| Unit | Response |
|---------|-------------------------|
| second | 2001-08-22 03:04:05.000 |
| minute | 2001-08-22 03:04:00.000 |
| hour | 2001-08-22 03:00:00.000 |
| day | 2001-08-22 00:00:00.000 |
| week | 2001-08-20 00:00:00.000 |
| month | 2001-08-01 00:00:00.000 |
| quarter | 2001-07-01 00:00:00.000 |
| year | 2001-01-01 00:00:00.000 |

● Example

The `date_trunc` function is applicable only to analytics at a fixed time interval. To implement analytics at a flexible interval, for example, every 5 minutes, you need to use a GROUP BY clause according to the mathematical modulus method.

```
* | SELECT count(1) as pv, __time__ - __time__% 300 as minute5 group by minute5 limit 100
```

In the preceding statement, `%300` indicates that modulo and truncation are performed every 5 minutes.

The following example shows how to use the `date_trunc` function:

```
*|select date_trunc('minute' , __time__) as t,
truncate (avg(latency) ) ,
current_date
group by t
order by t desc
limit 60
```

Interval functions

Interval functions perform interval-related calculation. For example, you can use interval functions to add or subtract an interval based on a date, or calculate the interval between two dates.

| Function | Description | Example |
|--|--|---|
| <code>date add (unit, value, timestamp)</code> | Adds an interval <code>value</code> of the <code>unit</code> type to a <code>timestamp</code> . To subtract an interval, use a negative <code>value</code> . | <code>date add('day',-7,'2018-08-09 00:00:00')</code> : indicates seven days before August 9. |
| <code>date diff(unit, timestamp1, timestamp2)</code> | Returns the time difference between <code>timestamp1</code> and <code>timestamp2</code> expressed in terms of <code>unit</code> . | <code>date diff('day', '2018-08-02 00:00:00', '2018-08-09 00:00:00') = 7</code> |

The following table lists the values of the `unit` parameter that are supported by the interval functions.

| Value | Description |
|-------------|------------------|
| millisecond | The millisecond. |

| Value | Description |
|---------|------------------------------------|
| second | The second. |
| minute | The minute. |
| hour | The hour. |
| day | The day. |
| week | The week. |
| month | The month. |
| quarter | The quarter, namely, three months. |
| year | The year. |

Time series padding function

The time series padding function is used to pad time series and corresponding data.

Note This function must be used in combination with the `group by time order by time` clause. When used together, the `order by` clause does not support the `desc` sorting method.

- Syntax

```
time_series(time_column, window, format, padding_data)
```

- Parameters

| Parameter | Description |
|---------------------|--|
| <i>time_column</i> | The name of the time field in a log entry. The default field name is <code>__time__</code> . The field value is of the LONG or TIMESTAMP type. |
| <i>window</i> | The time window for a data query. It is composed of a number and a unit. Unit: s (seconds), m (minutes), H (hours), and d (days). For example, 2h, 5m, or 3d. |
| <i>format</i> | The MySQL time format displayed. |
| <i>padding_data</i> | The content to be added for a time point. Valid values: <ul style="list-style-type: none"> ◦ 0: adds 0. ◦ null: adds null. ◦ last: adds the value corresponding to the last time point. ◦ next: adds the value corresponding to the next time point. ◦ avg: adds the average value of the last and next values. |

- Example

The following statement is used to format data every 2 hours:

```
* |select time_series(__time__, '2h', '%Y-%m-%d %H:%i:%s', '0') as stamp, count(*) as num from log group by stamp order by stamp
```

23.4.8.9. URL functions

This topic describes the syntax of URL functions and provides examples.

URL functions extract fields from standard URLs. A standard URL is described as follows:

```
[protocol:][//host[:port]][path][? query][#fragment]
```

Common URL functions

| Function | Description | Example | |
|---|---|--|---|
| | | Statement | Response |
| <code>url_extract_fragment(url)</code> | Extracts the fragment identifier from a URL and returns the fragment identifier of the VARCHAR type. | <pre>* select url_extract_fragment('https://sls.console.aliyun.com/#/project/dashboard-demo/categoryList')</pre> | <code>/project/dashboard-demo/categoryList</code> |
| <code>url_extract_host(url)</code> | Extracts the host information from a URL and returns the host information of the VARCHAR type. | <pre>* select url_extract_host('http://www.aliyun.com/product/sls')</pre> | <code>www.aliyun.com</code> |
| <code>url_extract_parameter(url, name)</code> | Extracts the value of the name parameter in the query from a URL and returns the value of the VARCHAR type. | <pre>* select url_extract_parameter('http://www.aliyun.com/product/sls?userid=testuser','userid')</pre> | <code>testuser</code> |
| <code>url_extract_path(url)</code> | Extracts the path from a URL and returns the path of the VARCHAR type. | <pre>* select url_extract_path('http://www.aliyun.com/product/sls?userid=testuser')</pre> | <code>/product/sls</code> |
| <code>url_extract_port(url)</code> | Extracts the port number from a URL and returns the port number of the BIGINT type. | <pre>* select url_extract_port('http://www.aliyun.com:80/product/sls?userid=testuser')</pre> | <code>80</code> |
| <code>url_extract_protocol(url)</code> | Extracts the protocol from a URL and returns the protocol of the VARCHAR type. | <pre>* select url_extract_protocol('http://www.aliyun.com:80/product/sls?userid=testuser')</pre> | <code>http</code> |
| <code>url_extract_query(url)</code> | Extracts the query string from a URL and returns the query string of the VARCHAR type. | <pre>* select url_extract_query('http://www.aliyun.com:80/product/sls?userid=testuser')</pre> | <code>userid=testuser</code> |
| <code>url_encode(value)</code> | Encodes a URL. | <pre>* select url_encode('http://www.aliyun.com:80/product/sls?userid=testuser')</pre> | <code>http%3a%2f%2fwww.aliyun.com%3a80%2fproduct%2fsls%3fuserid%3dtestuser</code> |

| Function | Description | Example | |
|--------------------------------|----------------|--|---|
| | | Statement | Response |
| <code>url decode(value)</code> | Decodes a URL. | <code>* select url decode('http%3a%2f%2fwww.alivun.com%3a80%2fproduct%2fsls%3fuserid%3dtestuser')</code> | <code>http://www.alivun.com:80/product/sls?userid=testuser</code> |

23.4.8.10. Regular expression functions

This topic describes the available regular expression functions. You can use these functions when you query and analyze data in Log Service.

A regular expression function parses a string and returns the required substrings.

The following table lists common regular expression functions.

| Function | Description | Example |
|---|---|--|
| <code>regexp_extract_all(string, pattern)</code> | Returns an array where each element is a substring that matches the regular expression. These substrings derive from the specified string. | The returned result of <code>* SELECT regexp_extract_all('5a 67b 890m', '(d+)')</code> is <code>['5','67','890']</code> . The returned result of <code>* SELECT regexp_extract_all('5a 67a 890m', '(d+)a')</code> is <code>['5a','67a']</code> . |
| <code>regexp_extract_all(string, pattern, group)</code> | Returns an array where each element is a part of a substring that matches the regular expression. This part is the content in the group parameter value of the <code>()</code> of a substring that derives from the specified string. | The returned result of <code>* SELECT regexp_extract_all('5a 67a 890m', '(d+)a', 1)</code> is <code>['5','67']</code> . |
| <code>regexp_extract(string, pattern)</code> | Returns the first substring of the specified string that matches the regular expression. | The returned result of <code>* SELECT regexp_extract('5a 67b 890m', '(d+)')</code> is <code>'5'</code> . |
| <code>regexp_extract(string, pattern, group)</code> | Returns a part of the first substring that matches the regular expression. This part is the content in the group parameter value of the <code>()</code> of the substring that derives from the specified string. | The returned result of <code>* SELECT regexp_extract('5a 67b 890m', '(d+)([a-z]+)', 2)</code> is <code>'a'</code> . |
| <code>regexp_like(string, pattern)</code> | Returns a Boolean value. If the string and its substrings cannot match the regular expression, the value <code>False</code> is returned. | The returned result of <code>* SELECT regexp_like('5a 67b 890m', '(d+m)')</code> is <code>True</code> . |
| <code>regexp_replace(string, pattern, replacement)</code> | Replaces the substrings of the specified string that match the regular expression with the value of the replacement parameter. | The returned result of <code>* SELECT regexp_replace('5a 67b 890m', '(d+)', 'a')</code> is <code>'aa ab am'</code> . |

| Function | Description | Example |
|--|---|---|
| <code>regexp_replace(string, pattern)</code> | Removes the substrings of the specified string that match the regular expression. This function is equivalent to <code>regexp_replace(string, pattern, '')</code> . | The returned result of <code>* SELECT regexp_replace('5a 67b 890m', '\d+')</code> is <code>'a b m'</code> . |
| <code>regexp_split(string, pattern)</code> | Returns an array where each element is a substring of the specified string that is split based on the regular expression. | The returned result of <code>* SELECT regexp_split('5a 67b 890m', '\d+')</code> is <code>['a','b','m']</code> . |

23.4.8.11. JSON functions

JSON functions can convert a string into a JSON type and extract JSON fields. The two major JSON data types are map and array. If a string cannot be converted to a value of the JSON type, null value is returned.

For information about how to expand JSON data into multiple rows, see [UNNEST function](#).

The following table lists the JSON functions that Log Service supports:

| Function | Description | Example |
|---|---|--|
| <code>json_parse(string)</code> | Converts a string to a JSON-formatted data. | <code>SELECT json_parse('[1, 2, 3]')</code> : returns a JSON array. |
| <code>json_format(json)</code> | Converts JSON-formatted data to a string. | <code>SELECT json_format(json_parse('[1, 2, 3]'))</code> : returns a string. |
| <code>json_array_contains(json, value)</code> | Determines whether a value exists in a JSON array or in a string that contains a JSON array. | <code>SELECT json_array_contains(json_parse('[1, 2, 3]'), 2)</code> or <code>SELECT json_array_contains('[1, 2, 3]', 2)</code> |
| <code>json_array_get(json_array, index)</code> | Retrieves the element at the specified index in the JSON array. This function is equivalent to <code>json_array_contains</code> . | <code>SELECT json_array_get(['a', 'b', 'c'], 0)</code> : returns 'a' |
| <code>json_array_length(json)</code> | Returns the length of the JSON array. | <code>SELECT json_array_length('[1, 2, 3]')</code> : returns 3 |
| <code>json_extract(json, json_path)</code> | Extracts a value from a JSON object and returns a JSON object. The <code>json_path</code> expression works in a similar manner to <code>\$.store.book[0].title</code> . | <code>SELECT json_extract(json, '\$.store.book')</code> |
| <code>json_extract_scalar(json, json_path)</code> | Returns a string. This function works in a similar manner to <code>json_extract</code> . | N/A |
| <code>json_size(json, json_path)</code> | Retrieves the length of a JSON object or array. | <code>SELECT json_size('[1, 2, 3]')</code> : returns 3. |

23.4.8.12. Type conversion functions

Type conversion functions convert the data type of a specified value or column in a query.

You can use the index attribute feature of Log Service to set the data type of a field to LONG, DOUBLE, TEXT, or JSON. You can also query fields of various data types, including BIGINT, DOUBLE, VARCHAR, and TIMESTAMP. To query fields of a specific data type, you can use type conversion functions to convert the data type configured in an index into the data type that you use in a query.

Syntax

Note We recommend that you use the `try_cast()` function if a log contains dirty data. Otherwise, a query may fail due to the dirty data.

- Convert the data type of a column of values or a specific value into the specified type in a query. If the data type of a value fails to be converted, the query is terminated.

```
cast([key|value] AS type)
```

- Convert the data type of a column of values or a specific value into the specified type in a query. If the data type of a value fails to be converted, NULL is returned for the value, and the query continues.

```
try_cast([key|value] AS type)
```

| Parameter | Description |
|-----------|---|
| key | The key of a field whose value data type is to be converted. |
| value | The field value whose data type is to be converted into the specified type. |

Example

- To convert the numeric value 123 to a value of the VARCHAR type, use the following statement:

```
cast(123 AS varchar)
```

- To convert the data type of the uid field values to the VARCHAR type, use the following statement:

```
cast(uid AS varchar)
```

23.4.8.13. IP functions

This topic describes the syntax of IP functions and provides examples.

IP functions can identify whether an IP address is an intranet or Internet IP address. IP functions can also identify the country, province, and city where an IP address resides. For information about geohash functions, see [Geography functions](#).

| Function | Description | Example |
|--------------------------------|---|---------------------------------------|
| <code>ip_to_domain(ip)</code> | Identifies whether an IP address is an intranet or Internet IP address. This function returns intranet or internet. | <code>SELECT ip_to_domain(ip)</code> |
| <code>ip_to_country(ip)</code> | Identifies the country where an IP address resides. | <code>SELECT ip_to_country(ip)</code> |

| Function | Description | Example |
|--------------------------------------|--|---|
| <code>ip_to_province(ip)</code> | Identifies the province where an IP address resides. | <code>SELECT ip_to_province(ip)</code> |
| <code>ip_to_city(ip)</code> | Identifies the city where an IP address resides. | <code>SELECT ip_to_city(ip)</code> |
| <code>ip_to_geo(ip)</code> | Identifies the longitude and latitude of the city where an IP address resides. The result is returned in the format of <code>latitude, longitude</code> . | <code>SELECT ip_to_geo(ip)</code> |
| <code>ip_to_city_geo(ip)</code> | Identifies the longitude and latitude of the city where an IP address resides. Each city has only one longitude and latitude. The result is returned in the format of <code>latitude, longitude</code> . | <code>SELECT ip_to_city_geo(ip)</code> |
| <code>ip_to_provider(ip)</code> | Identifies the network service provider that assigns an IP address. | <code>SELECT ip_to_provider(ip)</code> |
| <code>ip_to_country(ip,'en')</code> | Identifies the country where an IP address resides. The function returns a country code. | <code>SELECT ip_to_country(ip,'en')</code> |
| <code>ip_to_country_code(ip)</code> | Identifies the country where an IP address resides. The function returns a country code. | <code>SELECT ip_to_country_code(ip)</code> |
| <code>ip_to_province(ip,'en')</code> | Identifies the province where an IP address resides. | <code>SELECT ip_to_province(ip,'en')</code> |
| <code>ip_to_city(ip,'en')</code> | Identifies the city where an IP address resides. | <code>SELECT ip_to_city(ip,'en')</code> |

Example

- To query the number of requests that are not sent from an intranet, run the following statement:

```
* | SELECT count(1) where ip_to_domain(ip)!='intranet'
```

- To query the top 10 provinces from which requests are sent, run the following statement:

```
* | SELECT count(1) as pv, ip_to_province(ip) as province GROUP BY province order by pv desc limit 10
```

Sample response

```
[
  {
    "__source__": "",
    "__time__": "1512353137",
    "province": "Zhejiang",
    "pv": "4045"
  }, {
    "__source__": "",
    "__time__": "1512353137",
    "province": "Shanghai",
    "pv": "3727"
  }, {
    "__source__": "",
    "__time__": "1512353137",
    "province": "Beijing",
    "pv": "954"
  }, {
    "__source__": "",
    "__time__": "1512353137",
    "province": "intranet IP address",
    "pv": "698"
  }, {
    "__source__": "",
    "__time__": "1512353137",
    "province": "Guangdong",
    "pv": "472"
  }, {
    "__source__": "",
    "__time__": "1512353137",
    "province": "Fujian",
    "pv": "71"
  }
]
```

The response contains an intranet IP address. You can use the SELECT statement to filter requests that are sent from the IP address.

- To query the top 10 provinces from which intranet requests are sent, run the following statement:

```
* | SELECT count(1) as pv, ip_to_province(ip) as province WHERE ip_to_domain(ip) != 'intranet' GROUP BY province ORDER BY pv desc limit 10
```

- To query the average request latency, the maximum request latency, and the request with the maximum latency from each country, run the following statement:

```
* | SELECT AVG(latency),MAX(latency),MAX_BY(requestId, latency) ,ip_to_country(ip) as country group by country limit 100
```

- To query the average latency of requests supported by each network service provider, run the following statement:

```
* | SELECT AVG(latency) , ip_to_provider(ip) as provider group by provider limit 100
```

- To query the longitude and latitude of the city to which an IP address belongs and show the city on a map, run the following statement:

```
* | SELECT count(1) as pv , ip_to_geo(ip) as geo group by geo order by pv desc
```

The following table shows the format of the result.

| | |
|-----|------------------|
| pv | geo |
| 100 | 35.3284,-80.7459 |

23.4.8.14. GROUP BY syntax

This topic describes the GROUP BY syntax.

GROUP BY statements support multiple columns. A GROUP BY statement allows you to specify a column alias in a SELECT statement to act as the corresponding KEY.

Example:

```
method:PostLogstoreLogs |select avg(latency),projectName,date_trunc('hour',__time__) as hour group by projectName,hour
```

The hour alias indicates the third SELECT column `date_trunc('hour',__time__)`. This improves the performance of complicated queries.

The GROUP BY statement supports the GROUPING SETS, CUBE, and ROLLUP clauses.

Example:

```
method:PostLogstoreLogs |select avg(latency) group by cube(projectName,logstore)
method:PostLogstoreLogs |select avg(latency) group by GROUPING SETS (( projectName,logstore), (projectName,method))
method:PostLogstoreLogs |select avg(latency) group by rollup(projectName,logstore)
```

Examples

- Use GROUP BY based on time

Each log has a built-in time column named `__time__`. When the analytics feature is enabled on one of the columns, the statistics of the time column are included.

The `date_trunc` function can truncate the time column to minute, hour, day, month, and year. The `date_trunc` function accepts an aligned unit and a column of the Unix timestamp type, such as `__time__`.

- PV statistics per hour and per minute

```
* | SELECT count(1) as pv , date_trunc('hour',__time__) as hour group by hour order by hour limit 100
* | SELECT count(1) as pv , date_trunc('minute',__time__) as minute group by minute order by minute limit 100
```

Note 100 specifies that up to 100 rows can be retrieved. If a LIMIT clause is not specified, up to 10 rows of data are retrieved by default.

- `date_trunc` functions are available only for statistics at a fixed time interval. For statistics based on flexible time dimensions, for example, every 5 minutes, run a GROUP BY statement based on the mathematical modulus method.

```
* | SELECT count(1) as pv, __time__ - __time__% 300 as minute5 group by minute5 limit 100
```

In the preceding statement, `%300` indicates that the time is truncated in mod every 5 minutes.

- Retrieve non-aggregation columns from a GROUP BY statement

In standard SQL, if a GROUP BY statement is used during the SELECT operation, Log Service selects only the raw data of the column specified in the GROUP BY statement or performs aggregation on any columns.

For example, the following statement is invalid. Log Service cannot determine which row of b to return during the GROUP BY operation based on a, because b is not a GROUP BY column.

```
*|select a, b, count(c) group by a
```

Instead, you can use the `arbitrary()` function to return b.

```
*|select a, arbitrary(b), count(c) group by a
```

23.4.8.15. Window functions

This topic describes the syntax for window functions.

Window functions are used to perform calculations across rows of a log. Common SQL aggregate functions calculate the results of only one row or aggregate all rows into one row for calculation. Window functions support cross-row calculation and fill the calculation results in each row.

Syntax for window functions is

```
SELECT key1, key2, value,
       rank() OVER (PARTITION BY key2
                   ORDER BY value DESC) AS rnk
FROM orders
ORDER BY key1, rnk
```

The important part is

```
rank() OVER (PARTITION BY KEY1 ORDER BY KEY2 DESC)
```

`rank()` is an aggregate function. You can use any functions in analysis syntax or the function listed in this topic. `PARTITION BY` indicates the buckets based on which values are calculated.

Special aggregate functions used in windows

| Function | Description |
|---|---|
| <code>rank()</code> | Returns the rank of a value within a group of values. The rank is one plus the number of preceding rows that are not peers of the current row. |
| <code>row_number()</code> | Returns a unique, sequential number for each row. |
| <code>first_value(x)</code> | Returns the first value of the window. In most cases, the function is used to obtain the maximum value after the values of the window are sorted. |
| <code>last_value(x)</code> | Returns the last value of the window. In most cases, the function is used to obtain the minimum value after the values of the window are sorted. |
| <code>nth_value(x, offset)</code> | Returns the value at the specified offset from the beginning of the window. |
| <code>lead(x, offset, default_value)</code> | Returns the value at offset rows after the current row in the window. If the target row does not exist, the <code>default_value</code> is returned. |
| <code>lag(x, offset, default_value)</code> | Returns the value at offset rows after the current row in the window. If the target row does not exist, the <code>default_value</code> is returned. |

Example

- Rank the salaries of employees in their respective departments

```
* |select department, personId, salary , rank() over(PARTITION BY department order by salary desc) as salary_rank
order by department,salary_rank
```

Results

| department | personId | salary | salary_rank |
|------------|-------------|--------|-------------|
| dev | john | 9000 | 1 |
| dev | Smith | 8000 | 2 |
| dev | Snow | 7000 | 3 |
| dev | Achilles | 6000 | 4 |
| Marketing | Blan Stark | 9000 | 1 |
| Marketing | Rob Stark | 8000 | 2 |
| Marketing | Sansa Stark | 7000 | 3 |

- Calculate the salaries of employees as percentages in their respective departments

```
* |select department, personId, salary *1.0 / sum(salary) over(PARTITION BY department ) as salary_percentage
```

Results

| department | personId | salary | salary_percentage |
|------------|-------------|--------|-------------------|
| dev | john | 9,000 | 0.3 |
| dev | Smith | 8,000 | 0.26 |
| dev | Snow | 7000 | 0.23 |
| dev | Achilles | 6000 | 0.2 |
| Marketing | Blan Stark | 9000 | 0.375 |
| Marketing | Rob Stark | 8000 | 0.333 |
| Marketing | Sansa Stark | 7000 | 0.29 |

- Calculate the daily UV increase over the previous day

```
* |select day ,uv, uv *1.0 /(lag(uv,1,0) over() ) as diff_percentage from
(
select approx_distinct(ip) as uv, date_trunc('day',__time__) as day from log group by day order by day asc
)
```

Results

| day | uv | diff_percentage |
|---------------------|-----|-----------------|
| 2017-12-01 00:00:00 | 100 | null |

| day | uv | diff_percentage |
|---------------------|-------|-----------------|
| 2017-12-02 00:00:00 | 125 | 1.25 |
| 2017-12-03 00:00:00 | 1,500 | 1.2 |
| 2017-12-04 00:00:00 | 175 | 1.16 |
| 2017-12-05 00:00:00 | 2,000 | 1.14 |
| 2017-12-06 00:00:00 | 225 | 1.125 |
| 2017-12-07 00:00:00 | 250 | 1.11 |

23.4.8.16. HAVING syntax

This topic describes the HAVING syntax.

The LogSearch/Analytics feature of Log Service supports the standard SQL HAVING clause. The HAVING clause is used with the GROUP BY statement to filter GROUP BY results.

Example:

```
method :PostLogstoreLogs |select avg(latency),projectName group by projectName HAVING avg(latency) > 100
```

Difference between HAVING and WHERE clauses

The HAVING clause is used to filter the aggregation and calculation results after you run the GROUP BY statement. The WHERE clause is used to filter the raw data during the aggregation calculation.

Example

Calculate the average rainfall of each province where the temperature is higher than 10°C, and only show the provinces with average rainfall greater than 100 mL in the final results:

```
* |select avg(rain) ,province where temperature > 10 group by province having avg(rain) > 100
```

23.4.8.17. ORDER BY syntax

This topic describes the ORDER BY syntax.

The ORDER BY keyword is used to sort output results. You can sort output results based on only one column.

- Syntax format

```
order by column name [desc|asc]
```

- Example

```
method :PostLogstoreLogs |select avg(latency) as avg_latency,projectName group by projectName
HAVING avg(latency) > 5700000
order by avg_latency desc
```

23.4.8.18. LIMIT syntax

The LIMIT clause is used to limit the number of returned rows.

Syntax formats

Log Service supports the following LIMIT syntax formats:

- Reads only the first N rows:

```
limit N
```

- Reads N rows starting from the S-th row:

```
limit S, N
```

Note

- If you use the LIMIT clause to paginate results, the final results rather than the intermediate results of the SQL query are obtained.
- You cannot apply the LIMIT clause to subqueries. For example, the following statement is not supported:

```
* | select count(1) from ( select distinct(url) from limit 0,1000)
```

- If you use the LIMIT clause for pagination, the offset value cannot exceed 1,000,000. For example, in the limit S, N clause, the sum of S and N cannot exceed 1,000,000, and the value of N cannot exceed 10,000.

Example

- To obtain the first 100 rows of results, run the following statement.

```
* | select distinct(url) from log limit 100
```

- To obtain a total of 1,000 results from row 0 to row 999, run the following statement.

```
* | select distinct(url) from log limit 0,1000
```

- To obtain a total of 1,000 results from row 1,000 to row 1,999, run the following statement:

```
* | select distinct(url) from log limit 1000,1000
```

23.4.8.19. Syntax for CASE statements and if() functions

This topic describes the syntax for CASE statements and if() functions.

CASE statements are used to classify continuous data. For example, you can use the following CASE statement to extract information from http_user_agent and classify the information into two types: Android and iOS.

```
SELECT
CASE
WHEN http_user_agent like '%android%' then 'android'
WHEN http_user_agent like '%ios%' then 'ios'
ELSE 'unknown' END
as http_user_agent,
count(1) as pv
group by http_user_agent
```

Examples

- Calculate the ratio of requests whose status code is 200 to all requests

```
* | SELECT
sum(
CASE
WHEN status =200 then 1
ELSE 0 end
)*1.0 / count(1) as status_200_percentage
```

- Calculate the distribution of latencies

```
* | SELECT `
CASE
WHEN latency < 10 then 's10'
WHEN latency < 100 then 's100'
WHEN latency < 1000 then 's1000'
WHEN latency < 10000 then 's10000'
else 's_large' end
as latency_slot,
count(1) as pv
group by latency_slot
```

Syntax for if() functions

An if() function works in the same way as a CASE statement does.

```
CASE
  WHEN condition THEN true_value
  [ ELSE false_value ]
END
```

- if(condition, true_value)
If the condition is true, the true_value column is returned. Otherwise, null is returned.
- if(condition, true_value, false_value)
If the condition is true, the true_value column is returned. Otherwise, the false_value column is returned.

Syntax for coalesce() functions

A coalesce() function returns the first non-null value from multiple columns.

```
COALESCE (value1, value2 [,...])
```

Syntax for the nullif() function

If value1 is equal to value2, null is returned. Otherwise, value1 is returned.

```
nullif(value1, value2)
```

Syntax for the try() function

The try() function catches underlying exceptions, such as division by zero errors, and returns null.

```
try(expression)
```

23.4.8.20. Nested subqueries

This describes how to use nested subqueries when you query logs.

You can use nested queries to perform more complicated queries.

Nested queries differ from non-nested queries in the need for specifying the FROM clause in the SQL statement. You must specify the `from log` keyword in each SQL statement to read raw data from logs.

Example:

```
* | select sum(pv) from
(
  select count(1) as pv from log group by method
)
```

23.4.8.21. Array functions

This topic describes the syntax of array functions. It also provides examples that show how to use these functions.

| Function | Description | Example |
|--|---|---------|
| Subscript operator [] | The subscript operator [] is used to obtain an element in the array. | N/A |
| array_distinct | Returns the distinct elements in the array. | N/A |
| array_intersect(x, y) | Returns the intersection of the x and y arrays. | N/A |
| array_union(x, y) → array | Returns the union of the x and y arrays. | N/A |
| array_except(x, y) → array | Returns the subtraction of the x and y arrays. | N/A |
| array_join(x, delimiter, null_replacement) → varchar | <p>Joins string arrays with the delimiter into a string and replaces null values with null_replacement.</p> <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e0f0ff;"> <p>Note The maximum size of the returned result of this array_join function is 1 KB. If the returned result exceeds 1 KB, the extra data will be truncated.</p> </div> | N/A |
| array_max(x) → x | Returns the maximum value in the x array. | N/A |
| array_min(x) → x | Returns the minimum value in the x array. | N/A |
| array_position(x, element) → bigint | Returns the subscript of the element in the x array. The subscript starts from 1. The value 0 is returned if no subscript is found. | N/A |
| array_remove(x, element) → array | Removes the element from the array. | N/A |
| array_sort(x) → array | Sorts the array and moves null values to the end. | N/A |

| Function | Description | Example |
|--|--|---|
| cardinality(x) → bigint | Returns the array size. | N/A |
| concat(array1, array2, ..., arrayN) → array | Concatenates arrays. | N/A |
| contains(x, element) → boolean | Returns true if the x array contains the specified element. | N/A |
| filter(array, function) → array | For more information about this Lambda function, see the filter() function in Lambda functions . | N/A |
| flatten(x) → array | Flattens an array(array(T)) to an array(T) by concatenating the contained arrays. | N/A |
| reduce(array, initialState, inputFunction, outputFunction) → x | For more information, see the reduce() function in Lambda functions . | N/A |
| reverse(x) → array | Returns an array that has the reversed order of the x array. | N/A |
| sequence(start, stop) → array | Generates a sequence of elements from start to stop. The difference between elements is 1. | N/A |
| sequence(start, stop, step) → array | Generates a sequence of elements from start to stop. The difference between elements is step. | N/A |
| sequence(start, stop, step) → array | Generates a sequence of timestamps from start to stop. The difference between timestamps is step. The type of step can be either INTERVAL DAY TO SECOND or INTERVAL YEAR TO MONTH. | N/A |
| shuffle(x) → array | Shuffles the array. | N/A |
| slice(x, start, length) → array | Returns a subset of the x array starting from the start value with the specified length. | N/A |
| transform(array, function) → array | For more information, see the transform() function in Lambda functions . | N/A |
| zip(array1, array2[, ...]) → array | Merges the specified arrays. The M-th element of the N-th argument will be the N-th field of the M-th output element. | SELECT zip(ARRAY[1, 2], ARRAY['1h' , null, '3h']) — ROW(1, '1h'), ROW(2, null), ROW(null, '3h')] |
| zip_with(array1, array2, function) → array | For more information, see the zip_with() function in Lambda functions . | N/A |

| Function | Description | Example |
|---|---|--|
| <code>array_agg (key)</code> | An aggregate function that returns an array from values in the key column. | <code>* select array_agg(key)</code> |
| <code>array_transpose(array[array[x,y,z], array[a,b,c]])</code> | Returns a new matrix by transposing the values of the previous matrix from rows to columns. | N/A |

23.4.8.22. Binary string functions

This topic describes the syntax of binary string functions. It also provides examples that show how to use these functions.

Data of the VARBINARY type is different from data of the VARCHAR type.

| Function | Description |
|--|--|
| Concatenation operator () | The result of <code>a b</code> is <code>ab</code> . |
| <code>length(binary) → bigint</code> | Returns the length in binary. |
| <code>concat(binary1, ..., binaryN) → varbinary</code> | Connects the binary strings, which is equivalent to . |
| <code>to_base64(binary) → varchar</code> | Converts a binary string to a Base64 string. |
| <code>from_base64(string) → varbinary</code> | Converts a Base64 string to a binary string. |
| <code>to_base64url(binary) → varchar</code> | Converts a string to a URL-safe Base64 string. |
| <code>from_base64url(string) → varbinary</code> | Converts a URL-safe Base64 string to a binary string. |
| <code>to_hex(binary) → varchar</code> | Converts a binary string to a hexadecimal string. |
| <code>from_hex(string) → varbinary</code> | Converts a hexadecimal string to a binary string. |
| <code>to_big_endian_64(bigint) → varbinary</code> | Convert a number to a binary string in big endian mode. |
| <code>from_big_endian_64(binary) → bigint</code> | Converts a binary string in big endian mode to a number. |
| <code>md5(binary) → varbinary</code> | Calculates the MD5 value of a binary string. |
| <code>sha1(binary) → varbinary</code> | Calculates the SHA1 value of a binary string. |
| <code>sha256(binary) → varbinary</code> | Calculates the SHA256 hash value of a binary string. |
| <code>sha512(binary) → varbinary</code> | Calculate the SHA512 value of a binary string. |
| <code>xxhash64(binary) → varbinary</code> | Calculates the xxhash64 value of a binary string. |

23.4.8.23. Bitwise operations

This topic describes the syntax for bitwise operations. It also provides examples that show how to use these operations.

| Function | Description | Example |
|--|--|---|
| <code>bit_count(x, bits) → bigint</code> | Count the number of 1 in the binary expression of x. | <pre>SELECT bit_count(9, 64); -- 2 SELECT bit_count(9, 8); -- 2 SELECT bit_count(-7, 64); -- 62 SELECT bit_count(-7, 8); -- 6</pre> |
| <code>bitwise_and(x, y) → bigint</code> | Perform the AND operation on x and y in the binary form. | N/A |
| <code>bitwise_not(x) → bigint</code> | Calculate the opposite values of all bits of x in the binary form. | N/A |
| <code>bitwise_or(x, y) → bigint</code> | Perform the OR operation on x and y in the binary form. | N/A |
| <code>bitwise_xor(x, y) → bigint</code> | Perform the XOR operation on x and y in the binary form. | N/A |

23.4.8.24. Interval-valued comparison and periodicity-valued comparison functions

Log Service allows you to use interval-valued comparison and periodicity-valued comparison functions to query and analyze log data.

You can use the functions to compare the value for one period with that for a previous period.

| Function | Description | Example |
|--|---|--|
| <code>compare(value, time_window)</code> | <p>This function compares the value calculated for the current period with that calculated for the period before <code>time_window</code>.</p> <p>The data type of the values to be compared is Double or Long. The <code>time_window</code> parameter is measured in seconds. This function returns an array.</p> <p>Possible return values include the value for the current period, the value for the period before <code>time_window</code>, and the ratio of the current value to the value before <code>time_window</code>.</p> | <pre>* select compare(nv . 86400) from (select count(1) as pv from log)</pre> |

| Function | Description | Example |
|---|--|--|
| <code>compare(value, time_window1, time_window2)</code> | This function compares the current value with the values for periods before <code>time_window1</code> and <code>time_window2</code> . The comparison results are in the JSON array format, where the values must be returned in the following order: [current value, value before <code>time_window1</code> , value before <code>time_window2</code> , current value/value before <code>time_window1</code> , current value/value before <code>time_window2</code>]. | <pre>* select compare(pv, 86400, 172800) from (select count(1) as pv from log)</pre> |
| <code>compare(value, time_window1, time_window2, time_window3)</code> | This function compares the value for the current period with the values for periods before <code>time_window1</code> , <code>time_window2</code> , and <code>time_window3</code> . The comparison results are in the JSON array format, where the values must be returned in the following order: [current value, value before <code>time_window1</code> , value before <code>time_window2</code> , value before <code>time_window3</code> , current value/value before <code>time_window1</code> , current value/value before <code>time_window2</code> , current value/value before <code>time_window3</code>]. | <pre>* select compare(pv, 86400, 172800, 604800) from (select count(1) as pv from log)</pre> |
| <code>ts_compare(value, time_window)</code> | This function compares the value for the current period with the values for periods before <code>time_window1</code> and <code>time_window2</code> and returns a JSON array. The comparison results are in the JSON array format, where the values must be returned in the following order: [current value, value before <code>time_window1</code> , current value/value before <code>time_window1</code> , Unix timestamp that indicates the start time of the previous period]. This function is used to compare time series values. You must specify the GROUP BY keyword for the time column in SQL statements. | For example, <pre>* select ts_compare(pv, 86400) as d from (select date_trunc('minute', time) as t, count(1) as pv from log group by t order by t) group by t</pre> specifies that the function compares the calculation result of every minute in the current period with that of every minute in the last period. The comparison result is <code>d: [1251.0.1264.0.0.9897151898734177, 1539843780.0.1539757380.0]:2018-10-19 14:23:00.000</code> . |

Examples

- Calculate the ratio of the PV for an hour on a day to that for the same time period on a previous day.

The start time is 2018-07-25 14:00:00, and the end time is 2018-07-25 15:00:00.

Statement for query and analysis:

```
* | select compare( pv , 86400) from (select count(1) as pv from log)
```

In the preceding statement, 86400 indicates 86,400 seconds before the current period.

Return results:

```
[9.0,19.0,0.47368421052631579]
```

In the preceding results,

- o 9.0 is the PV for the period from 2018-07-25 14:00:00 to 2018-07-25 15:00:00.
- o 19.0 is the PV for the period from 2018-07-24 14:00:00 to 2018-07-24 15:00:00.
- o 0.47368421052631579 is the ratio of the PV for the current period to that for a previous period.

If you want to expand the array into three columns of numbers, the analysis statement is

```
* | select diff[1],diff[2],diff[3] from(select compare( pv , 86400) as diff from (select count(1) as pv from log))
```

- Calculate the PV ratio for every minute of the current hour to that in the same time period as the day before, and output the results in a line chart.
 - i. Calculate the PV ratio for every minute of the current hour to that in the same time period as the day before. The start time is 2018-07-25 14:00:00, and the end time is 2018-07-25 15:00:00.

Statement for query and analysis:

```
*| select t, compare( pv , 86400) as diff from (select count(1) as pv, date_format(from_unixtime(__time__), '%H:%i') as t from log group by t) group by t order by t
```

Return results:

| t | diff |
|-------|------------------------------------|
| 14:00 | [9520.0,7606.0,1.2516434393899554] |
| 14:01 | [8596.0,8553.0,1.0050274757395066] |
| 14:02 | [8722.0,8435.0,1.0340248962655603] |
| 14:03 | [7499.0,5912.0,1.2684370771312586] |

In the preceding table, t indicates the time in the format of **Hour:Minute** . The contents of the diff column are included in an array that contains the following values:

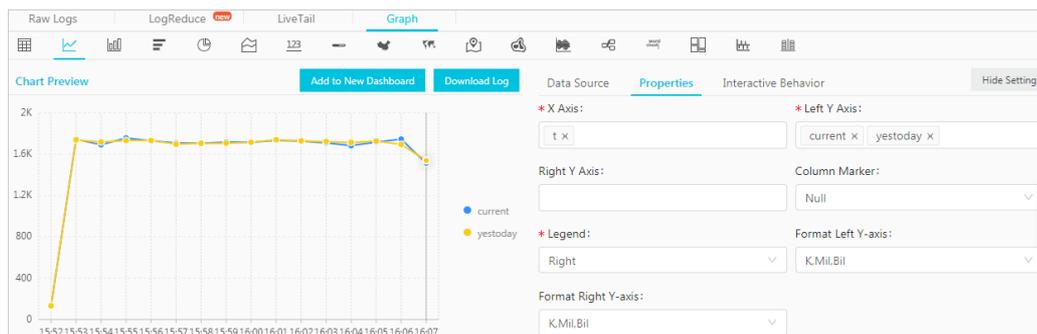
- The PV value of the current period.
- The PV value of the previous period.
- The ratio of the PV value for the current period to that for the previous period.

- ii. To show the query results in a line chart, use the following statement:

```
*|select t, diff[1] as current, diff[2] as yesterday, diff[3] as percentage from(select t, compare( pv , 86400) as diff from (select count(1) as pv, date_format(from_unixtime(__time__), '%H:%i') as t from log group by t) group by t order by t)
```

The two lines indicate the PV values of a day and the day before.

Line chart



23.4.8.25. Comparison functions and operators

This topic describes the comparison functions and operators in Log Service. You can use these functions and operators to query and analyze log data.

Comparison functions and operators

A comparison function compares two parameter values of any comparable data types, such as INTEGER, BIGINT, DOUBLE, and TEXT.

Comparison operators

A comparison operator is used to compare two values. If the statement is true, TRUE is returned. Otherwise, FALSE is returned.

| Operator | Description |
|----------|--------------------------|
| < | Less than |
| > | Greater than |
| <= | Less than or equal to |
| >= | Greater than or equal to |
| = | Equal to |
| <> | Not equal to |
| != | Not equal to |

Range operator BETWEEN

The BETWEEN operator determines whether a value falls in a specified closed interval.

- If the value falls in the specified closed interval, TRUE is returned. Otherwise, FALSE is returned.

Example: `SELECT 3 BETWEEN 2 AND 6;` . The statement is true, and TRUE is returned.

The preceding statement is equivalent to `SELECT 3 >= 2 AND 3 <= 6;` .

- The BETWEEN operator can be put behind the NOT operator to test whether a value falls out of a specified closed interval.

Example: `SELECT 3 NOT BETWEEN 2 AND 6;` . The statement is false, and FALSE is returned.

The preceding statement is equivalent to `SELECT 3 < 2 OR 3 > 6;` .

- If any of the three values is NULL, NULL is returned.

IS NULL and IS NOT NULL

The IS NULL and IS NOT NULL operators test whether a value is NULL.

IS DISTINCT FROM and IS NOT DISTINCT FROM

These operators are similar to the EQUAL TO and NOT EQUAL TO operators. However, IS DISTINCT FROM and IS NOT DISTINCT FROM can determine whether a NULL value exists.

Examples:

```
SELECT NULL IS DISTINCT FROM NULL; -- false
SELECT NULL IS NOT DISTINCT FROM NULL; -- true
```

The DISTINCT operator compares parameter values under multiple conditions, as described in the following table.

| a | b | a = b | a <> b | a DISTINCT b | a NOT DISTINCT b |
|------|------|-------|--------|--------------|------------------|
| 1 | 1 | TRUE | FALSE | FALSE | TRUE |
| 1 | 2 | FALSE | TRUE | TRUE | FALSE |
| 1 | NULL | NULL | NULL | TRUE | FALSE |
| NULL | NULL | NULL | NULL | FALSE | TRUE |

GREATEST and LEAST

These operators are used to obtain the maximum or minimum value from a row of field values.

Example:

```
select greatest(1,2,3) -- Returns 3.
```

Quantified comparison predicates: ALL, ANY, and SOME

The ALL, ANY, and SOME quantifiers can be used to determine whether a parameter value meets specified conditions.

- ALL is used to determine whether a parameter value meets all conditions. If the statement is true, TRUE is returned. Otherwise, FALSE is returned.
- ANY is used to determine whether a parameter value meets a condition. If the statement is true, TRUE is returned. Otherwise, FALSE is returned.
- SOME is used to determine whether a parameter value meets a condition. SOME is equivalent to ANY.
- ALL, ANY, and SOME must immediately follow comparison operators.

ALL and ANY support comparison under multiple conditions, as described in the following table.

| Expression | Description |
|----------------|--|
| A = ALL (...) | Returns TRUE if A matches all values. |
| A <> ALL (...) | Returns TRUE if A does not match a value. |
| A < ALL (...) | Returns TRUE if A is smaller than the smallest value. |
| A = ANY (...) | Returns TRUE if A is equal to a value. This statement is equivalent to A IN (...). |
| A <> ANY (...) | Returns TRUE if A does not match a value. |
| A < ANY (...) | Returns TRUE if A is smaller than the largest value. |

Examples:

```
SELECT 'hello' = ANY (VALUES 'hello', 'world'); -- true
SELECT 21 < ALL (VALUES 19, 20, 21); -- false
SELECT 42 >= SOME (SELECT 41 UNION ALL SELECT 42 UNION ALL SELECT 43); -- true
```

23.4.8.26. Lambda functions

This topic describes Lambda functions and provides some examples. You can use Lambda functions to analyze log data in Log Service.

Lambda expressions

Lambda expressions use the arrow operator `->`.

Examples:

```
x -> x + 1
(x, y) -> x + y
x -> regexp_like(x, 'a+')
x -> x[1] / x[2]
x -> IF(x > 0, x, -x)
x -> COALESCE(x, 0)
x -> CAST(x AS JSON)
x -> x + TRY(1 / 0)
```

Most MySQL expressions can be used in Lambda functions.

`filter(array<T>, function<T, boolean>) → ARRAY<T>`

Returns an array whose elements are filtered from the specified array based on the Lambda expression.

Examples:

```
SELECT filter(ARRAY [], x -> true); -- []
SELECT filter(ARRAY [5, -6, NULL, 7], x -> x > 0); -- [5, 7]
SELECT filter(ARRAY [5, NULL, 7, NULL], x -> x IS NOT NULL); -- [5, 7]
```

`map_filter(map<K, V>, function<K, V, boolean>) → MAP<K, V>`

Returns a map whose elements are filtered based on the Lambda expression. The map is generated from the map function.

Examples:

```
SELECT map_filter(MAP(ARRAY[], ARRAY[]), (k, v) -> true); -- {}
SELECT map_filter(MAP(ARRAY[10, 20, 30], ARRAY['a', NULL, 'c']), (k, v) -> v IS NOT NULL); -- {10 -> a, 30 -> c}
SELECT map_filter(MAP(ARRAY['k1', 'k2', 'k3'], ARRAY[20, 3, 15]), (k, v) -> v > 10); -- {k1 -> 20, k3 -> 15}
```

`reduce(array<T>, initialState S, inputFunction<S, T, S>, outputFunction<S, R>) → R`

The reduce function starts from the initial state, traverses each element in the array, and then calls `inputFunction(S,T)` to generate a new state. After all the elements in the array are traversed and the final state is generated, the reduce function calls `outputFunction` to assign the final state value to the result `R` and output the result. The procedure is described as follows:

1. Start from the initial state `S`.
2. Traverse each element `T`.
3. Calculate `inputFunction(S,T)` to generate a new state `S`.
4. Repeat steps 2 and 3 until the last element is traversed and has a new state.
5. Turn the final state `S` into the final result `R`.

Examples:

```
SELECT reduce(ARRAY [], 0, (s, x) -> s + x, s -> s); -- 0
SELECT reduce(ARRAY [5, 20, 50], 0, (s, x) -> s + x, s -> s); -- 75
SELECT reduce(ARRAY [5, 20, NULL, 50], 0, (s, x) -> s + x, s -> s); -- NULL
SELECT reduce(ARRAY [5, 20, NULL, 50], 0, (s, x) -> s + COALESCE(x, 0), s -> s); -- 75
SELECT reduce(ARRAY [5, 20, NULL, 50], 0, (s, x) -> IF(x IS NULL, s, s + x), s -> s); -- 75
SELECT reduce(ARRAY [2147483647, 1], CAST (0 AS BIGINT), (s, x) -> s + x, s -> s); -- 2147483648
SELECT reduce(ARRAY [5, 6, 10, 20], -- calculates arithmetic average: 10.25
    CAST(ROW(0.0, 0) AS ROW(sum DOUBLE, count INTEGER)),
    (s, x) -> CAST(ROW(x + s.sum, s.count + 1) AS ROW(sum DOUBLE, count INTEGER)),
    s -> IF(s.count = 0, NULL, s.sum / s.count));
```

transform(array<T>, function<T, U>) → ARRAY<U>

This Lambda function traverses each element in an array to generate a new result U.

Examples:

```
SELECT transform(ARRAY [], x -> x + 1); -- []
SELECT transform(ARRAY [5, 6], x -> x + 1); -- [6, 7] -- Increments each element by 1.
SELECT transform(ARRAY [5, NULL, 6], x -> COALESCE(x, 0) + 1); -- [6, 1, 7]
SELECT transform(ARRAY ['x', 'abc', 'z'], x -> x || '0'); -- ['x0', 'abc0', 'z0']
SELECT transform(ARRAY [ARRAY [1, NULL, 2], ARRAY[3, NULL]], a -> filter(a, x -> x IS NOT NULL)); -- [[1, 2], [3]]
```

zip_with(array<T>, array<U>, function<T, U, R>) → array<R>

This Lambda function merges two arrays and generates the element R in the new array based on element T and element U.

Examples:

```
SELECT zip_with(ARRAY[1, 3, 5], ARRAY['a', 'b', 'c'], (x, y) -> (y, x)) --Transposes the elements of the two arrays to generate a new array. Result: [['a', 1], ['b', 3], ['c', 5]]
SELECT zip_with(ARRAY[1, 2], ARRAY[3, 4], (x, y) -> x + y); -- Result: [4, 6]
SELECT zip_with(ARRAY['a', 'b', 'c'], ARRAY['d', 'e', 'f'], (x, y) -> concat(x, y)) -- Concatenates the elements of the two arrays to generate a new string. Result: ['ad', 'be', 'cf']
```

23.4.8.27. Logical functions

This topic describes the available logical functions in Log Service. You can use these functions to query and analyze log data.

Logical operators

| Operator | Description | Example |
|----------|---|---------|
| AND | The result is TRUE if both values are TRUE. | a AND b |
| OR | The result is TRUE if either value is TRUE. | a OR b |
| NOT | The result is TRUE if the value is FALSE. | NOT a |

Effect of NULL on logical operators

The following tables list the truth values when the values of a and b are TRUE, FALSE, and NULL, respectively.

Truth table 1

| a | b | a AND b | a OR b |
|-------|-------|---------|--------|
| TRUE | TRUE | TRUE | TRUE |
| TRUE | FALSE | FALSE | TRUE |
| TRUE | NULL | NULL | TRUE |
| FALSE | TRUE | FALSE | TRUE |
| FALSE | FALSE | FALSE | FALSE |
| FALSE | NULL | FALSE | NULL |
| NULL | TRUE | NULL | TRUE |
| NULL | FALSE | FALSE | NULL |
| NULL | NULL | NULL | NULL |

Truth table 2

| a | NOT a |
|-------|-------|
| TRUE | FALSE |
| FALSE | TRUE |
| NULL | NULL |

23.4.8.28. Field aliases

This topic describes how to specify an alias for a field and provides some examples.

A field name in an SQL statement must start with letters and contain digits and underscores (_).

If you have configured a field name that does not conform to the SQL standard (such as User-Agent), you must specify an alias for the field on the field index configuration page. The alias takes effect only for the duration of the SQL statement. The data is still stored under the original field name. You must specify the original name when you perform a search.

You can also specify an alias for a field in an SQL statement if the original name is long.

Sample aliases

| Original field name | Alias |
|-------------------------|-------|
| User-Agent | ua |
| User.Agent | ua |
| 123 | col |
| abceefghijklmnopqrstuvw | a |

23.4.8.29. JOIN operations between Logstores and Relational Database Service (RDS) tables

This topic describes how to join Logstores in Log Service with RDS tables for queries and store the query results in RDS tables.

Procedure

1. Create a VPC.

Create an RDS instance and specify the VPC to host the RDS instance. Then the VPC ID and the RDS instance ID are obtained.

2. Configure a whitelist for the RDS instance.

Add the following CIDR blocks to the whitelist: `100.104.0.0/16` , `11.194.0.0/16` , and `11.201.0.0/16`

3. Create an external store

Run the following statement to create an external store. Replace the parameter values based on your business needs.

```
{
  "externalStoreName":"storeName",
  "storeType":"rds-vpc",
  "parameter":
  {
    "region":"cn-qingdao",
    "vpc-id":"vpc-m5eq4irc1pucp*****"
    "instance-id":"i-m5eeo2whsn*****"
    "host":"localhost",
    "port":"3306",
    "username":"root",
    "password":"*****",
    "db":"scmc"
    "table":"join_meta"
  }
}
```

Parameters

| Parameter | Description |
|-------------|---|
| region | The region where your RDS instance resides. |
| vpc-id | The ID of the VPC where your RDS instance resides. |
| instance-id | The ID of the RDS instance. |
| host | The ID of the ECS instance that is used to access the RDS instance. |
| port | The port of the ECS instance that is used to access the RDS instance. |
| username | The username that is used to log on to the RDS instance. |
| password | The password that is used to log on to the RDS instance. |

| Parameter | Description |
|-----------|--|
| db | The name of the database. |
| table | The name of the table with which the Logstore is joined. |

Note You can join a Logstore with an RDS table that resides only in the China (Beijing), China (Qingdao), and China (Hangzhou) regions.

4. JOIN query.

Log on to the Log Service console. In the **Search & Analyze** search box, run a JOIN statement.

Supported JOIN syntax:

- INNER JOIN
- LEFT JOIN
- RIGHT JOIN
- FULL JOIN

```
[ INNER ] JOIN
LEFT [ OUTER ] JOIN
RIGHT [ OUTER ] JOIN
FULL [ OUTER ] JOIN
```

Note

- You can join Logstores only to external tables.
- In the JOIN statement, you must first specify a Logstore before specifying an external store.
- You must specify the name of the external store instead of the name of an RDS table. The external store name automatically changes into the combination of the RDS database name and the name of the RDS table that you want to join with the Logstore.

Sample JOIN statement:

```
method:postlogstorelogs | select count(1) , histogram(logstore) from log l join join_meta m on l.projectid = cast(m.ikey as varchar)
```

5. Store the query results to the RDS table.

You can use the INSERT statement to insert the query results into the RDS table.

```
method:postlogstorelogs | insert into method_output select cast(methodasvarchar(65535)),count(1)fromloggroupbymethod
```

Sample Python script:

```
# encoding: utf-8
from __future__ import print_function
from aliyun.log import *
from aliyun.log.util import base64_encodestring
from random import randint
import time
import os
from datetime import datetime
endpoint = os.environ.get('ALIYUN_LOG_SAMPLE_ENDPOINT', 'cn-chengdu.log.aliyuncs.com')
accessKeyId = os.environ.get('ALIYUN_LOG_SAMPLE_ACCESSID', '')
accessKey = os.environ.get('ALIYUN_LOG_SAMPLE_ACCESSKEY', '')
logstore = os.environ.get('ALIYUN_LOG_SAMPLE_LOGSTORE', '')
project = "ali-yunlei-chengdu"
client = LogClient(endpoint, accessKeyId, accessKey, token)
## Create an external store
res = client.create_external_store(project, ExternalStoreConfig("rds_store", "region", "rds-vpc", "vpc id", "instance-id", "instance-ip", "port", "username", "password", "db", "table"));
res.log_print()
## Obtain external store details
res = client.get_external_store(project, "rds_store");
res.log_print()
res = client.list_external_store(project, "");
res.log_print();
# Perform the JOIN operation.
req = GetLogStoreLogsRequest(project, logstore, From, To, "", "select count(1) from "+ logstore +" s join meta m on s.p
rojectid = cast(m.ikey as varchar)");
res = client.get_logs(req)
res.log_print();
# Store query results to the RDS table
req = GetLogStoreLogsRequest(project, logstore, From, To, "", "insert into rds_store select count(1) from "+ logstore );
res = client.get_logs(req)
res.log_print();
```

23.4.8.30. Geospatial functions

This topic describes the available geospatial functions in Log Service. You can use these functions to query and analyze log data.

Concept of geometry

Geospatial functions support geometries in the well-known text (WKT) format.

Geometry formats

| Geometry | WKT format |
|-----------------|---|
| Point | POINT (0 0) |
| LineString | LINestring (0 0, 1 1, 1 2) |
| Polygon | POLYGON ((0 0, 4 0, 4 4, 0 4, 0 0), (1 1, 2 1, 2 2, 1 2, 1 1)) |
| MultiPoint | MULTIPOINT (0 0, 1 2) |
| MultiLineString | MULTILINESTRING ((0 0, 1 1, 1 2), (2 3, 3 2, 5 4)) |
| MultiPolygon | MULTIPOLYGON (((0 0, 4 0, 4 4, 0 4, 0 0), (1 1, 2 1, 2 2, 1 2, 1 1)), ((-1 -1, -1 -2, |

| | |
|--------------------|--|
| Geometry | WKT format |
| GeometryCollection | GEOMETRYCOLLECTION (POINT(2 3), LINESTRING (2 3, 3 4)) |

Constructors

Constructor description

| Function | Description |
|---|---|
| ST_Point(double, double) → Point | Returns a geometry point instance with the specified coordinate values. |
| ST_LineFromText(varchar) → LineString | Returns a geometry LineString instance from a WKT representation. |
| ST_Polygon(varchar) → Polygon | Returns a geometry polygon instance from a WKT representation. |
| ST_GeometryFromText(varchar) → Geometry | Returns a geometry instance from a WKT representation. |
| ST_AsText(Geometry) → varchar | Returns the WKT representation of a geometry. |

Operations

| Function | Description |
|---|--|
| ST_Boundary(Geometry) → Geometry | Returns the closure of the combinatorial boundary of a geometry. |
| ST_Buffer(Geometry, distance) → Geometry | Returns the geometry that represents all points whose distance from the specified geometry is shorter than or equal to the specified distance. |
| ST_Difference(Geometry, Geometry) → Geometry | Returns the geometry value that represents the point set difference of the specified geometries. |
| ST_Envelope(Geometry) → Geometry | Returns the bounding rectangular polygon of a geometry. |
| ST_ExteriorRing(Geometry) → Geometry | Returns a line string that represents the exterior ring of the input polygon. |
| ST_Intersection(Geometry, Geometry) → Geometry | Returns the geometry value that represents the point set intersection of two geometries. |
| ST_SymDifference(Geometry, Geometry) → Geometry | Returns the geometry value that represents the point set symmetric difference of two geometries. |

Relationship tests

| Function | Description |
|----------|-------------|
|----------|-------------|

| Function | Description |
|---|---|
| ST_Contains(Geometry, Geometry) → boolean | Returns True if and only if no points of the second geometry lie in the exterior of the first geometry, and at least one point of the interior of the first geometry lies in the interior of the second geometry. Returns False if points of the second geometry are on the boundary of the first geometry. |
| ST_Crosses(Geometry, Geometry) → boolean | Returns True if the specified geometries share some, but not all, interior points in common. |
| ST_Disjoint(Geometry, Geometry) → boolean | Returns True if the specified geometries do not spatially intersect. |
| ST_Equals(Geometry, Geometry) → boolean | Returns True if the specified geometries represent the same geometry. |
| ST_Intersects(Geometry, Geometry) → boolean | Returns True if the specified geometries spatially intersect in two dimensions. |
| ST_Overlaps(Geometry, Geometry) → boolean | Returns True if the specified geometries share space in the same dimension, but are not completely contained by each other. |
| ST_Relate(Geometry, Geometry) → boolean | Returns True if the first geometry is spatially related to the second geometry. |
| ST_Touches(Geometry, Geometry) → boolean | Returns True if the specified geometries have at least one point in common, but their interiors do not intersect. |
| ST_Within(Geometry, Geometry) → boolean | Returns True if the first geometry is completely inside the second geometry. Returns False if the two geometries have points in common at the boundaries. |

Accessors

| Function | Description |
|--|--|
| ST_Area(Geometry) → double | Returns the two-dimensional Euclidean area of a geometry. |
| ST_Centroid(Geometry) → Geometry | Returns the point value that is the mathematical centroid of a geometry. |
| ST_CoordDim(Geometry) → bigint | Returns the coordinate dimension of a geometry. |
| ST_Dimension(Geometry) → bigint | Returns the inherent dimension of a geometry object, which must be less than or equal to the coordinate dimension. |
| ST_Distance(Geometry, Geometry) → double | Returns the minimum two-dimensional Cartesian distance (based on spatial ref) between two geometries in projected units. |
| ST_IsClosed(Geometry) → boolean | Returns True if the start and end points of the linestring are coincident. |

| Function | Description |
|---------------------------------------|---|
| ST_IsEmpty(Geometry) → boolean | Returns True if the specified geometry is an empty geometry, such as geometry collection, polygon, and point. |
| ST_IsRing(Geometry) → boolean | Returns True if and only if the line is closed and simple. |
| ST_Length(Geometry) → double | Returns the length of a LineString or multi-LineString by using Euclidean measurement on a two-dimensional plane (based on spatial ref) in projected units. |
| ST_XMax(Geometry) → double | Returns the X maximum of the bounding box of the geometry. |
| ST_YMax(Geometry) → double | Returns the Y maximum of the bounding box of the geometry. |
| T_XMin(Geometry) → double | Returns the X minimum of the bounding box of the geometry. |
| ST_YMin(Geometry) → double | Returns the Y minimum of the bounding box of the geometry. |
| ST_StartPoint(Geometry) → point | Returns the first point of a geometry LineString instance. |
| ST_EndPoint(Geometry) → point | Returns the last point of a geometry LineString instance. |
| ST_X(Point) → double | Returns the X coordinate of a point. |
| ST_Y(Point) → double | Returns the Y coordinate of a point. |
| ST_NumPoints(Geometry) → bigint | Returns the number of points in a geometry. |
| ST_NumInteriorRing(Geometry) → bigint | Returns the cardinality of the collection of interior rings of a polygon. |

23.4.8.31. Geography functions

This topic describes the syntax of geography functions and provides some examples.

For information about functions that identify the country, province, city, ISP, and the longitude and latitude of a specified IP address, see [IP functions](#).

Geography functions

| Function | Description | Example |
|-----------------|--|---|
| geohash(string) | Returns the geohash value of the specified geographical location. The geographical location is represented by a string in the format of "latitude, longitude". The values for latitude and longitude are separated by a comma. | <pre>select geohash('34.1,120.6')= 'wwjcbdrnzs'</pre> |

| Function | Description | Example |
|------------------|--|---|
| geohash(lat,lon) | Returns the geohash value of the specified geographical location. The geographical location is represented by two parameters that indicate the latitude and longitude. | <code>select geohash(34.1,120.6)= 'wwjcbdrnzs'</code> |

23.4.8.32. JOIN syntax

The JOIN operation joins multiple tables by using one or more fields in the tables. You can join a Logstore created in Log Service with the Logstore itself, with another Logstore, or with an RDS table. This topic describes how to join different Logstores.

Procedure

1. Download the [latest version of the SDK for Python](#).
2. Call the GetProjectLogs operation to query logs.

Sample SDK

```
#!/usr/bin/env python
#encoding: utf-8
import time,sys,os
from aliyun.log.logexception import LogException
from aliyun.log.logitem import LogItem
from aliyun.log.logclient import LogClient
from aliyun.log.getlogsrequest import GetLogsRequest
from aliyun.log.getlogsrequest import GetProjectLogsRequest
from aliyun.log.putlogsrequest import PutLogsRequest
from aliyun.log.listtopicsrequest import ListTopicsRequest
from aliyun.log.listlogstoresrequest import ListLogstoresRequest
from aliyun.log.gethistogramsrequest import GetHistogramsRequest
from aliyun.log.index_config import *
from aliyun.log.logtail_config_detail import *
from aliyun.log.machine_group_detail import *
from aliyun.log.acl_config import *
if __name__=='__main__':
    token = None
    endpoint = "http://cn-hangzhou.log.aliyuncs.com"
    accessKeyId = 'LTAIvKy7U'
    accessKey='6gXLNTLyCfdsfwrewrfhdkfsdfuiwu'
    client = LogClient(endpoint, accessKeyId, accessKey,token)
    logstore = "meta"
    # In the query statements, specify two Logstores, the query time ranges of both Logstores, and the key that you want to use to join the two Logstores.
    req = GetProjectLogsRequest(project,"select count(1) from sls_operation_log s join meta m on s.__date__>'2018-04-10 00:00:00' and s.__date__ < '2018-04-11 00:00:00' and m.__date__>'2018-04-23 00:00:00' and m.__date__ <'2018-04-24 00:00:00' and s.projectid = cast(m.ikey as varchar)");
    res = client.get_project_logs(req)
    res.log_print();
    exit(0)
```

 **Note** For more information about the JOIN syntax and examples, see [Join](#).

23.4.8.33. UNNEST function

This topic describes the UNNEST function.

Scenarios

Log data is typically stored as primitive data types, such as string or number. In certain scenarios, log data may include complex data types, such as arrays, maps, and JSON objects. The UNNEST function can be used to transform complex data types into rows of primitive data types. This simplifies query and analysis.

Example:

```
__source__: 1.1.1.1
__tag__:__hostname__: vm-req-170103232316569850-tianchi111932.tc
__topic__: TestTopic_4
array_column: [1,2,3]
double_column: 1.23
map_column: {"a":1,"b":2}
text_column: Product
```

The values of the `array_column` field are arrays. To obtain the sum of elements of all `array_column` field values, you must traverse all elements of every array.

UNNEST function

| Syntax | Description |
|---|---|
| <code>unnest(array) as table_alias(column_name)</code> | Expands an array into multiple rows. The column name of these rows is <code>column_name</code> . |
| <code>unnest(map) as table(key_name, value_name)</code> | Expands a map into multiple rows. <code>key_name</code> specifies the column name of the keys, and <code>value_name</code> specifies the column name of the values. |

Note The UNNEST function is used to expand arrays or maps. If you want to expand a string, you must transform the string into a JSON object, and then convert the JSON object into an array or map. To do this, you can use the `cast(json_parse(array_column) as array(bigint))` function.

Traverse every element of an array

Expands an array into multiple rows by using the following SQL SELECT statement:

```
* |select array_column, a from log, unnest( cast( json_parse( array_column ) as array( bigint ) ) ) as t( a )
```

The UNNEST function `unnest(cast(json_parse(array_column) as array(bigint))) as t(a)` expands the array into multiple rows. The rows are stored in a derived table referenced as `t`, with the column referenced as `a`.

- Calculate the sum of the elements in an array:

```
* |select sum( a ) from log, unnest( cast( json_parse( array_column ) as array( bigint ) ) ) as t( a )
```

- Perform a GROUP BY operation on all elements of an array:

```
* |select a, count( 1 ) from log, unnest( cast( json_parse( array_column ) as array( bigint ) ) ) as t( a ) group by a
```

Traverse every key and value of a map

- Traverse every key and value of a map:

```
* |select map_column , a,b from log, unnest( cast( json_parse(map_column) as map(vvarchar, bigint) ) ) as t(a,b)
```

- Perform a GROUP BY operation on all keys of a map:

```
* |select key, sum(value) from log, unnest( cast( json_parse(map_column) as map(vvarchar, bigint) ) ) as t(key,value) GROUP BY key
```

Visualize the query results of the histogram and numeric_histogram functions.

- histogram

The histogram function works in a similar manner to the count group by syntax. For more information about the histogram function, see [Map functions](#).

In most cases, the histogram function returns a JSON object. The following is an example:

```
* |select histogram(method)
```

You can use the UNNEST function to expand JSON data into multiple rows. Then the data can be visualized. The following is an example:

```
* |select key , value from( select histogram(method) as his from log ) , unnest(his ) as t(key,value)
```

- numeric_histogram

The numeric_histogram function assigns a column of numeric values into multiple bins. This function is equivalent to a GROUP BY operation that is performed on a numeric value column. For more information about the syntax of the numeric_histogram function, see [Approximate functions](#).

```
* |select numeric_histogram(10,Latency)
```

Use the following SELECT statement to visualize the result:

```
* | select key,value from(select numeric_histogram(10,Latency) as his from log ) , unnest(his) as t(key,value)
```

23.4.9. Machine learning syntax and functions

23.4.9.1. Overview

The machine learning feature of Log Service supports multiple algorithms and calling methods. You can use SELECT statements and machine learning functions to analyze the characteristics of a field or fields within a period of time.

Log Service offers multiple time series analysis algorithms to help you implement time series prediction, time series anomaly detection, time series decomposition, and multi-time series clustering. The algorithms are compatible with standard SQL statements. This greatly simplifies the use of the algorithms and improves the troubleshooting efficiency.

Features

- Supports various smooth operations on single-time series data.
- Supports algorithms related to the prediction, anomaly detection, change point detection, inflection point detection, and multi-period estimation of single-time series data.
- Supports decomposition operations on single-time series data.
- Supports various clustering algorithms of multi-time series data.
- Supports multi-field pattern mining (based on the sequence of numeric data or text).

Limits

- The specified time series data must be sampled based on the same interval.
- The specified time series data cannot contain data repeatedly sampled from the same time point.

| Item | Description |
|---|---|
| Processing capacity of time-series data | Data can be collected from a maximum of 150,000 consecutive time points. If the data volume exceeds the processing capacity, you must aggregate the data or reduce the sampling amount. |
| Clustering capacity of the density-based clustering algorithm | A maximum of 5,000 time series curves, each of which cannot contain more than 1,440 time points. |
| Clustering capacity of the hierarchical clustering algorithm | A maximum of 2,000 time series curves, each of which cannot contain more than 1,440 time points. |

Machine learning functions

| Type | Function | Description |
|--|-----------------------|--|
| Smooth functions | ts_smooth_simple | Uses the Holt Winters algorithm to smooth time series data. |
| | ts_smooth_fir | Uses the finite impulse response (FIR) filter to smooth time series data. |
| | ts_smooth_iir | Uses the infinite impulse response (IIR) filter to smooth time series data. |
| Multi-period estimation functions | ts_period_detect | Forecasts time series data by period. |
| Change point detection functions | ts_cp_detect | Finds intervals with different statistical characteristics from time series data. The interval endpoints are change points. |
| | ts_breakout_detect | Finds the time points when statistics steeply increases or decreases from time series data. |
| Maximum value detection function | ts_find_peaks | Finds the local maximum value of time series data in a specified window. |
| Prediction and anomaly detection functions | ts_predicate_simple | Uses default parameters to model time series data and performs simple time series prediction and anomaly detection. |
| | ts_predicate_ar | Uses an autoregressive (AR) model to model time series data and performs simple time series prediction and anomaly detection. |
| | ts_predicate_arma | Uses an autoregressive moving average (ARMA) model to model time series data and performs simple time series prediction and anomaly detection. |
| | ts_predicate_arima | Uses an autoregressive integrated moving average (ARIMA) model to model time series data and performs simple time series prediction and anomaly detection. |
| | ts_regression_predict | Accurately predicts the trend for a single periodic time series with a certain tendency. |

| Type | Function | Description |
|--|--|--|
| Time series decomposition function | ts_decompose | Uses the Seasonal and Trend decomposition using Loess (STL) algorithm to decompose time series data. |
| Time series clustering functions | ts_density_cluster | Uses a density-based clustering method to cluster multiple pieces of time series data. |
| | ts_hierarchical_cluster | Uses a hierarchical clustering method to cluster multiple pieces of time series data. |
| | ts_similar_instance | Queries curves that are similar to a specified curve. |
| Frequent pattern statistics function | pattern_stat | Mines representative combinations of attributes among the given multi-attribute field samples to obtain the frequent pattern in statistical patterns. |
| Differential pattern statistics function | pattern_diff | Finds the pattern that causes differences between two collections under specified conditions. |
| Root cause analysis function | rca_kpi_search | When a time series metric is abnormal, you can use the root cause analysis function to analyze the dimension attributes that result in the abnormal metric in a timely manner. |
| Correlation analysis functions | ts_association_analysis | Quickly finds the metrics that are correlated with a specified metric among multiple observed metrics in the system. |
| ts_similar | Quickly finds the metrics that are correlated with specified time series data among multiple observed metrics in the system. | |
| Kernel density estimation function | kernel_density_estimation | Uses the smooth peak function to fit the observed data points, thus simulating the real probability distribution curve. |

23.4.9.2. Smooth functions

This topic describes the smooth functions that you can use to smooth and filter specified time series curves. Filtering is the first step to discover the shapes of time series curves.

Functions

| Function | Description |
|------------------|---|
| ts_smooth_simple | Uses the Holt-Winters algorithm to filter time series data. This function is the default smooth function. |
| ts_smooth_fir | Uses a finite impulse response (FIR) filter to filter time series data. |
| ts_smooth_iir | Uses an infinite impulse response (IIR) filter to filter time series data. |

ts_smooth_simple

- Syntax:

```
select ts_smooth_simple(x,y)
```

- The following table lists the parameters of the function.

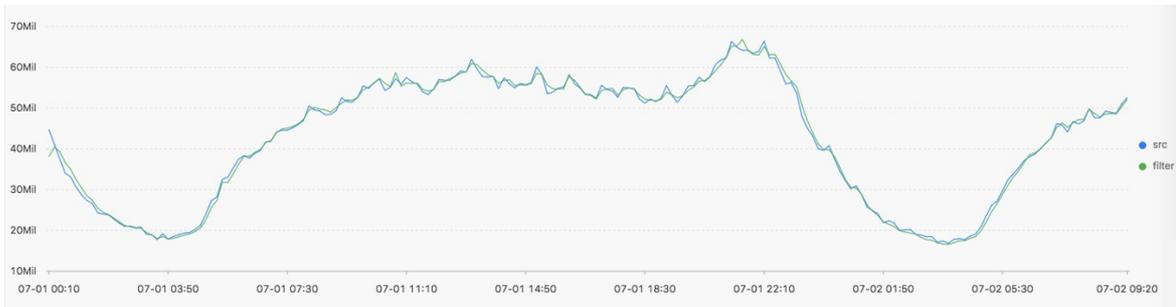
| Parameter | Description | Value |
|-----------|--|--|
| <i>x</i> | The time sequence. The time points along the x axis are sorted in the ascending order. | The Unix timestamp of the time series data. Unit: seconds. |
| <i>y</i> | The sequence of numeric data at each specified time point. | - |

- Example

- The search and analytic statement is shown as follows:

```
* | select ts_smooth_simple(stamp, value) from ( select __time__ - __time__ % 120 as stamp, avg(v) as value from log GROUP BY stamp order by stamp )
```

- The following figure shows the response.



- The following table lists the display items.

| Item | | Description |
|-----------------|----------|--|
| Horizontal axis | unixtime | The Unix timestamp of time series data. Unit: seconds. |
| Vertical axis | src | The unfiltered data. |
| | filter | The filtered data. |

ts_smooth_fir

- Syntax:

- If you cannot determine filter parameters, use built-in window parameters in the following statement:

```
select ts_smooth_fir(x,y,winType,winSize)
```

- If you can determine filter parameters, you can specify the parameters as needed in the following statement:

```
select ts_smooth_fir(x,y,array[])
```

- The following table lists the parameters of the function.

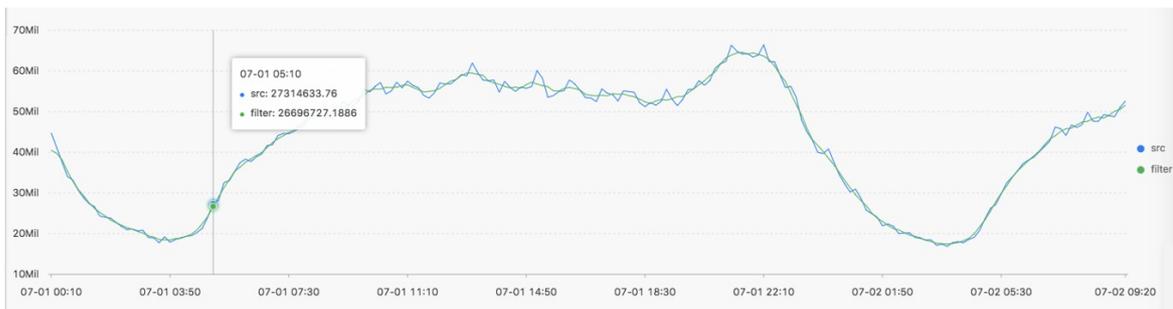
| Parameter | Description | Value |
|-----------|-------------|-------|
|-----------|-------------|-------|

| Parameter | Description | Value |
|----------------|--|--|
| <i>x</i> | The time sequence. The time points along the x axis are sorted in the ascending order. | Each time point is a Unix timestamp. Unit: seconds. |
| <i>y</i> | The sequence of numeric data at each specified time point. | - |
| <i>winType</i> | The type of window used for filtering. | Valid values: <ul style="list-style-type: none"> ◦ rectangle: rectangle window. ◦ hanning: hanning window ◦ hamming: hamming window. ◦ blackman: blackman window. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note We recommend that you select the rectangle window for better display effects. </div> |
| <i>winSize</i> | The length of the filtering window. | The value is of the LONG type. Valid values: 2 to 15. |
| <i>array[]</i> | The parameter used for FIR filtering. | The value is an array where the sum of elements is 1. For example, array[0.2, 0.4, 0.3, 0.1]. |

- Example 1
 - The search and analytic statement is shown as follows:

```
* | select ts_smooth_fir(stamp, value, 'rectangle', 4) from ( select __time__ - __time__ % 120 as stamp, avg(v) as value from log GROUP BY stamp order by stamp )
```

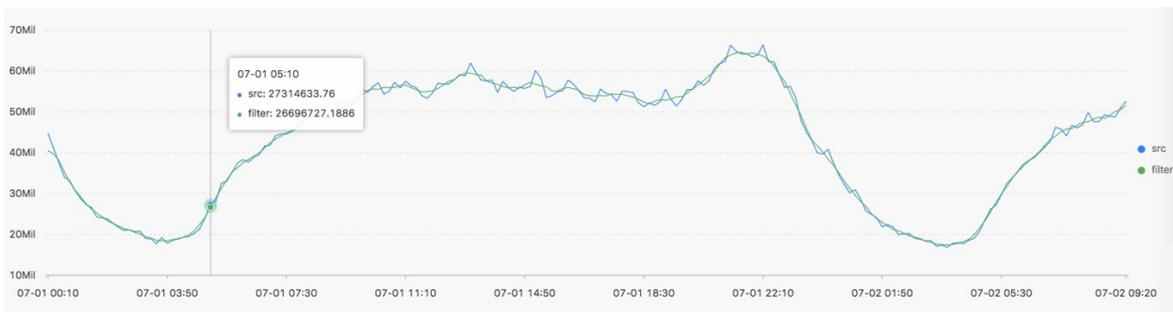
- The following figure shows the response.



- Example 2
 - The search and analytic statement is shown as follows:

```
* | select ts_smooth_fir(stamp, value, array[0.2, 0.4, 0.3, 0.1]) from ( select __time__ - __time__ % 120 as stamp, avg(v) as value from log GROUP BY stamp order by stamp )
```

- The following figure shows the response.



- The following table lists the display items.

| Item | | Description |
|-----------------|----------|--|
| Horizontal axis | unixtime | The Unix timestamp of the time series data. Unit: seconds. |
| Vertical axis | src | The unfiltered data. |
| | filter | The filtered data. |

ts_smooth_iir

- Syntax:

```
select ts_smooth_iir(x, y, array[], array[])
```

- The following table lists the parameters of the function.

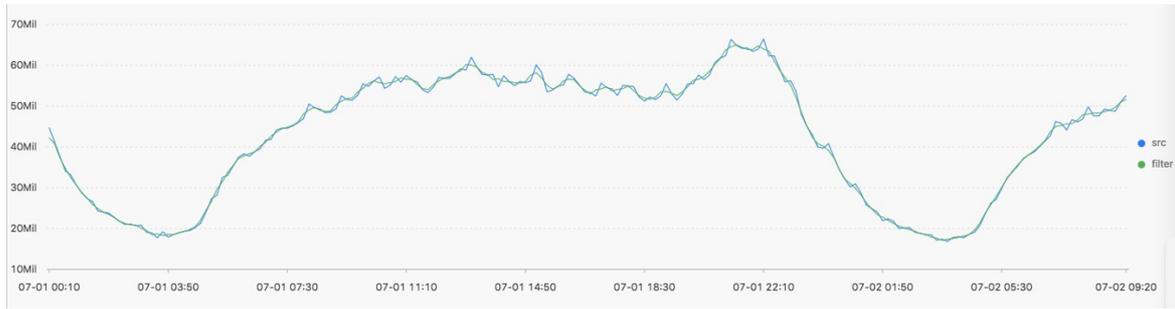
| Parameter | Description | Value |
|-----------|--|--|
| x | The time sequence. The time points along the x axis are sorted in ascending order. | Each time point is a Unix timestamp. Unit: seconds. |
| y | The sequence of numeric data at each specified time point. | - |
| $array[]$ | The parameter used for IIR filtering in terms of x_i . | The value is an array where the sum of elements is 1. The length of the array ranges from 2 to 15. For example, array[0.2, 0.4, 0.3, 0.1]. |
| $array[]$ | The parameter used for IIR filtering in terms of y_{i-1} . | The value is an array where the sum of elements is 1. The length of the array ranges from 2 to 15. For example, array[0.2, 0.4, 0.3, 0.1]. |

- Example

- The search and analytic statement is shown as follows:

```
* | select ts_smooth_iir(stamp, value, array[0.2, 0.4, 0.3, 0.1], array[0.4, 0.3, 0.3]) from ( select __time__ - __time__ % 120 as stamp, avg(v) as value from log GROUP BY stamp order by stamp )
```

- o The following figure shows the response.



- The following table lists the display items.

| Item | | Description |
|-----------------|----------|--|
| Horizontal axis | unixtime | The Unix timestamp of the time series data. Unit: seconds. |
| Vertical axis | src | The unfiltered data. |
| | filter | The filtered data. |

23.4.9.3. Multi-period estimation functions

This topic describes the multi-period estimation functions that are supported by Log Service. You can use the functions to estimate the periods of time series data in different time intervals. You can also extract the periods by performing a series of operations such as Fourier transform (FT).

Functions

| Function | Description |
|---------------------------------|--|
| <code>ts_period_detect</code> | Estimates the periods of time series data that is distributed in different time intervals. |
| <code>ts_period_classify</code> | Uses FT to calculate the periodicity of specified time series curves. This function can be used to identify periodic curves. |

ts_period_detect

Syntax:

```
select ts_period_detect(x,y,minPeriod,maxPeriod)
```

The following table lists the parameters of the function.

| Parameter | Description | Value |
|-----------|--|--|
| <i>x</i> | The time sequence. The points in time along the horizontal axis are sorted in ascending order. | Each point is a UNIX timestamp. Unit: seconds. |
| <i>y</i> | The sequence of numeric data at each specified point in time. | - |

| Parameter | Description | Value |
|------------------|--|--|
| <i>minPeriod</i> | The ratio of the minimum length of the estimated period to the total length of the time series data. | The value must be a decimal number. Value range: (0, 1]. |
| <i>maxPeriod</i> | The ratio of the maximum length of the estimated period to the total length of the time series data. <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> ? Note The value of the <i>maxPeriod</i> parameter must be greater than the value of the <i>minPeriod</i> parameter. </div> | The parameter value must be a decimal number. Value range: (0, 1]. |

Example:

• Query statement

```
* | select ts_period_detect(stamp, value, 0.2, 1.0) from ( select __time__ - __time__ % 120 as stamp, avg(v) as value from m log GROUP BY stamp order by stamp )
```

• Result



The following table lists the display items.

| Display item | Description |
|--------------|--|
| period_id | An array with a length of 1. The element in the array indicates the sequence number of the period. The array [0] indicates the original time series curve. |
| time_series | The sequence of timestamps. |
| data_series | The sequence of data at each timestamp. <ul style="list-style-type: none"> • If the value of period_id is 0, the returned data is the original time series data. • If the value of period_id is not 0, the data returned is filtered time series data. |

ts_period_classify

Syntax:

```
select ts_period_classify(stamp,value,instanceName)
```

The following table lists the parameters.

| Parameter | Description | Value |
|-----------|-------------|-------|
|-----------|-------------|-------|

| Parameter | Description | Value |
|--------------|--|--|
| stamp | The time sequence. The points in time along the horizontal axis are sorted in ascending order. | Each point in time is a UNIX timestamp. Unit: seconds. |
| value | The sequence of numeric data at each specified point in time. | - |
| instanceName | The name of the time series curve. | - |

Example:

- The following query statement is executed:

```
* and h : nu2h05202.nu8 | select ts_period_classify(stamp, value, name) from log
```

- Query result

| line_name | prob | type |
|--------------------------|---------------------|------|
| asg-2z9yn6zf5ewg188pg5 | 1.0 | -1.0 |
| asg-bp1j8anc92p6v5pptgpi | 0.07203669207039314 | 0.0 |
| asg-wz9hse7u4ubopo5d19o | 0.0 | 0.0 |
| asg-bp18oqnl0gg96vy85te4 | 0.05590892692207093 | 0.0 |

The following table lists the display items.

| Display item | Description |
|--------------|---|
| line_name | An array with a length of 1. The element in the array indicates the sequence number of the period. The array [0] indicates the original time series curve. |
| prob | The ratio of the primary period length to the length of the time series curve. Value range: [0, 1]. You can set the value to 0.15 when you perform a test. |
| type | The type of the curve. Valid values: -1, -2, and 0. <ul style="list-style-type: none"> The value -1 indicates that the length of the time series curve is too short (less than 64 points). The value -2 indicates the time series curve has a high fault rate (the fault rate exceeds 20%). The value 0 indicates the time series curve is periodic. |

23.4.9.4. Change point detection functions

This topic describes the change point detection functions in Log Service. You can use the functions to detect the change points in time series data.

The change point detection functions can detect the following two kinds of change points:

- Statistics feature changes within a specified period of time
- Anomalies in time series data

Functions

| Function | Description |
|---------------------------------|--|
| <code>ts_cp_detect</code> | Finds intervals in which data has different statistics features. The interval endpoints are change points. |
| <code>ts_breakout_detect</code> | Finds the time points at which data experiences dramatic changes. |

ts_cp_detect

Syntax:

- If you cannot specify an appropriate time window size, use the following syntax. The default window size used in the function is 10.

```
select ts_cp_detect(x, y, samplePeriod)
```

- To adjust the effect specific to your business environment, you can specify the `minSize` parameter in the following function.

```
select ts_cp_detect(x, y, minSize)
```

The following table lists the parameters of the function.

| Parameter | Description | Value |
|----------------|--|---|
| <i>x</i> | The time sequence. The time points along the x axis are sorted in the ascending order. | Each time point is a Unix timestamp. Unit: seconds. |
| <i>y</i> | The sequence of numeric data at each specified time point. | - |
| <i>minSize</i> | The minimum length of time series data in a continuous interval. | The minimum value is 3 and the maximum value cannot exceed ten percent of the length of the specified time series data. |

Example:

- The search and analytic statement is shown as follows:

```
* | select ts_cp_detect(stamp, value, 3) from (select __time__ - __time__ % 10 as stamp, avg(v) as value from log GROUP BY stamp order by stamp)
```

- The following figure shows the response.



The following table lists the display items.

| Display item | Description |
|-----------------|--|
| Horizontal axis | <code>unixtime</code> The Unix timestamp of time series data, measured in seconds, for example, 1537071480. |

| Display item | | Description |
|---------------|------|--|
| Vertical axis | src | The unfiltered data, such as 1956092.7647745228. |
| | prob | The probability that a time point is a change point. Valid values: 0 to 1. |

ts_breakout_detect

Syntax:

```
select ts_breakout_detect(x, y, winSize)
```

The following table lists the parameters of the function.

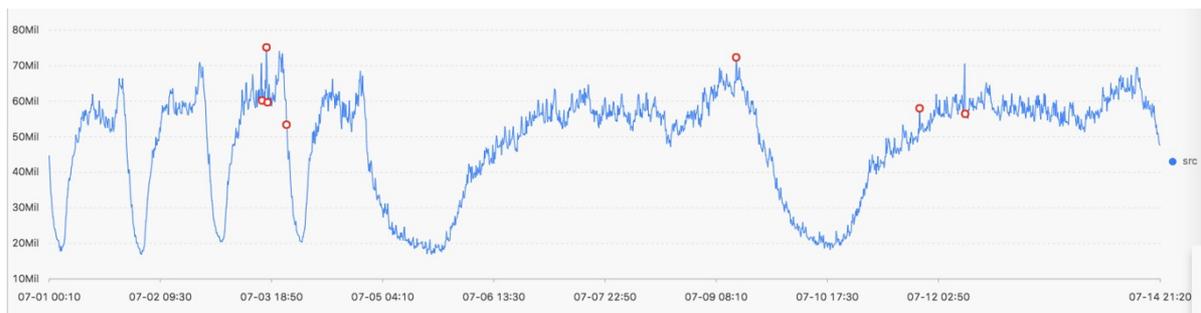
| Parameter | Description | Value |
|----------------|--|---|
| <i>x</i> | The time sequence. The time points along the x axis are sorted in the ascending order. | Each time point is a Unix timestamp. Unit: seconds. |
| <i>y</i> | The sequence of numeric data at each specified time point. | - |
| <i>winSize</i> | The minimum length of time series data in a continuous interval. | The minimum value is 3 and the maximum value cannot exceed ten percent of the length of the specified time series data. |

Example:

- The search and analytic statement is shown as follows:

```
* | select ts_breakout_detect(stamp, value, 3) from (select __time__ - __time__% 10 as stamp, avg(v) as value from log GROUP BY stamp order by stamp)
```

- The following figure shows the response.



The following table lists the display items.

| Display item | | Description |
|-----------------|----------|---|
| Horizontal axis | unixtime | The Unix timestamp of time series data, measured in seconds, for example, 1537071480. |
| Vertical axis | src | The unfiltered data, such as 1956092.7647745228. |
| | prob | The probability that a time point is a change point. Valid values: 0 to 1. |

23.4.9.5. Maximum value detection function

This topic describes the available maximum value detection function in Log Service. You can use the functions to find the local maximum value of time series data in a specified window.

ts_find_peaks

Syntax:

```
select ts_find_peaks(x, y, winSize)
```

The following table lists the parameters of the function.

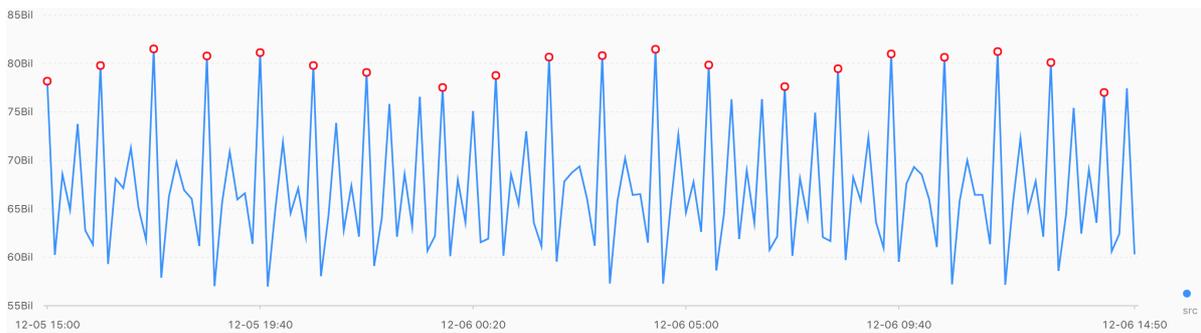
| Parameter | Description | Value |
|----------------|--|--|
| <i>x</i> | The time sequence. The time points along the x axis are sorted in the ascending order. | Each time point is a Unix timestamp. Unit: seconds. |
| <i>y</i> | The sequence of numeric data at each specified time point. | - |
| <i>winSize</i> | The minimum length of the detection window. | The value of the parameter is of the LONG type, ranging from 1 to the length of time series data. We recommend that you set this parameter to ten percent of the actual data length. |

Example:

- The search and analytic statement is shown as follows:

```
* and h : nu2h05202.nu8 and m: NET | select ts_find_peaks(stamp, value, 30) from (select __time__ - __time__ % 10 as stamp, avg(v) as value from log GROUP BY stamp order by stamp)
```

- The following figure shows the response.



The following table lists the display items.

| Display item | Description | |
|-----------------|-------------|---|
| Horizontal axis | unixtime | The Unix timestamp of time series data, measured in seconds, for example, 1537071480. |
| | src | The unfiltered data, such as 1956092.7647745228. |

| Display item | | Description |
|---------------|-----------|--|
| Vertical axis | peak_flag | <p>Indicates whether the numeric value at the time point is the maximum value. Valid values: 1.0 and 0.0.</p> <ul style="list-style-type: none"> 1.0: The numeric value at the time point is the maximum value. 0.0: The numeric value at the time point is not the maximum value. |

23.4.9.6. Prediction and anomaly detection functions

Prediction and anomaly detection functions predict the trend of time series curves and identify the Ksigma and quantiles of the errors between a predicted curve and an actual curve. You can use the functions to detect anomalies.

Functions

| Function | Description |
|------------------------------------|--|
| <code>ts_predicate_simple</code> | Uses default parameters to model time series data and performs prediction and anomaly detection on time series data. |
| <code>ts_predicate_ar</code> | Uses an autoregressive (AR) model to model time series data and performs prediction and anomaly detection on time series data. |
| <code>ts_predicate_arma</code> | Uses an autoregressive moving average (ARMA) model to model time series data and performs prediction and anomaly detection on time series data. |
| <code>ts_predicate_arima</code> | Uses an autoregressive integrated moving average (ARIMA) model to model time series data and performs prediction and anomaly detection on time series data. |
| <code>ts_regression_predict</code> | <p>Accurately predicts the trend for a periodic time series curve.</p> <p>Scenario: This function can be used to predict metering data, network traffic, financial data, and different business data that follows certain rules.</p> |
| <code>ts_anomaly_filter</code> | Filters the anomalies detected from multiple time series curves based on the custom anomaly mode. The anomalies are detected during the anomaly detection. This function helps you find abnormal curves in a timely manner. |

ts_predicate_simple

Syntax:

```
select ts_predicate_simple(x, y, nPred, isSmooth)
```

The following table lists the parameters of the function.

| Parameter | Description | Value |
|-----------|--|---|
| <i>x</i> | The time sequence. The time points along the x axis are sorted in the ascending order. | Each time point is a Unix timestamp. Unit: seconds. |
| <i>y</i> | The sequence of numeric data at each specified time point. | N/A |

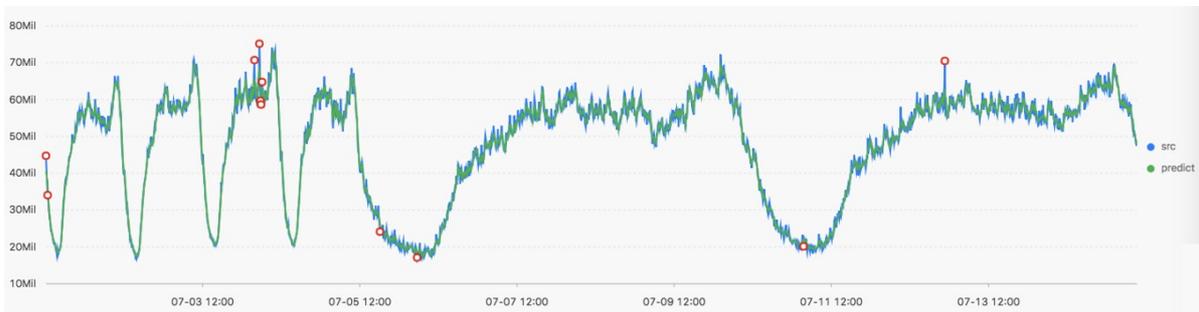
| Parameter | Description | Value |
|-----------------|---|--|
| <i>nPred</i> | The number of points for prediction. | The value is of the LONG data type and must be equal to or greater than 1. |
| <i>isSmooth</i> | Specifies whether to filter the raw data. | The value is of the Boolean type. The default value is True, which indicates to filter raw data. |

Example:

- A search and analytic statement is shown as follows:

```
* | select ts_predicate_simple(stamp, value, 6) from (select __time__ - __time__ % 60 as stamp, avg(v) as value from lo
g GROUP BY stamp order by stamp)
```

- The following figure shows the response.



The following table lists the display items.

| Display item | | Description |
|-----------------|--------------|---|
| Horizontal axis | unixtime | The Unix timestamp of the data. Unit: seconds. |
| Vertical axis | src | The raw data. |
| | predict | The predicted data. |
| | upper | The upper limit of the prediction. The confidence level is 0.85. This value cannot be modified. |
| | lower | The lower limit of the prediction. The confidence level is 0.85. This value cannot be modified. |
| | anomaly_prob | The probability that the point is an anomaly. Valid values: [0, 1]. |

ts_predicate_ar

Syntax:

```
select ts_predicate_ar(x, y, p, nPred, isSmooth)
```

The following table lists the parameters of the function.

| Parameter | Description | Value |
|-----------|--|---|
| <i>x</i> | The time sequence. The time points along the x axis are sorted in the ascending order. | Each time point is a Unix timestamp. Unit: seconds. |

| Parameter | Description | Value |
|-----------------|--|--|
| <i>y</i> | The sequence of numeric data at each specified time point. | N/A |
| <i>p</i> | The order of the AR model. | The value is of the LONG data type. Valid values: [2, 8]. |
| <i>nPred</i> | The number of points for prediction. | The value is of the LONG data type. Valid values: [1, 5 × <i>p</i>]. |
| <i>isSmooth</i> | Specifies whether to filter the raw data. | The value is of the Boolean type. The default value is true, which indicates to filter raw data. |

An example search and analytic statement is shown as follows:

```
* | select ts_predicate_ar(stamp, value, 3, 4) from (select __time__ - __time__ % 60 as stamp, avg(v) as value from log GR OUP BY stamp order by stamp)
```

 **Note** The response is similar to that of the `ts_predicate_simple` function. For more information, see the response of the `ts_predicate_simple` function.

ts_predicate_arma

Syntax:

```
select ts_predicate_arma(x, y, p, q, nPred, isSmooth)
```

The following table lists the parameters of the function.

| Parameter | Description | Value |
|-----------------|--|--|
| <i>x</i> | The time sequence. The time points along the x axis are sorted in the ascending order. | Each time point is a Unix timestamp. Unit: seconds. |
| <i>y</i> | The sequence of numeric data at each specified time point. | N/A |
| <i>p</i> | The order of the AR model. | The value is of the LONG data type. Valid values: [2, 100]. |
| <i>q</i> | The order of the ARMA model. | The value is of the LONG data type. Valid values: [2, 8]. |
| <i>nPred</i> | The number of points for prediction. | The value is of the LONG data type. Valid values: [<i>p</i> , 5 <i>p</i>]. |
| <i>isSmooth</i> | Specifies whether to filter the raw data. | The value is of the Boolean type. The default value is true, which indicates to filter raw data. |

An example search and analytic statement is shown as follows:

```
* | select ts_predicate_arma(stamp, value, 3, 2, 4) from (select __time__ - __time__ % 60 as stamp, avg(v) as value from log GROUP BY stamp order by stamp)
```

 **Note** The response is similar to that of the `ts_predicate_simple` function. For more information, see the response of the `ts_predicate_simple` function.

ts_predicate_arima

Syntax:

```
select ts_predicate_arima(x,y, p, d, q, nPred, isSmooth)
```

The following table lists the parameters of the function.

| Parameter | Description | Value |
|-----------------|--|--|
| <i>x</i> | The time sequence. The time points along the x axis are sorted in the ascending order. | Each time point is a Unix timestamp. Unit: seconds. |
| <i>y</i> | The sequence of numeric data at each specified time point. | N/A |
| <i>p</i> | The order of the AR model. | The value is of the LONG data type. Valid values: [2, 8]. |
| <i>d</i> | The order of the ARIMA model. | The value is of the LONG data type. Valid values: [1, 3]. |
| <i>q</i> | The order of the ARMA model. | The value is of the LONG data type. Valid values: [2, 8]. |
| <i>nPred</i> | The number of points for prediction. | The value is of the LONG type. Valid values: [<i>p</i> , 5 <i>p</i>]. |
| <i>isSmooth</i> | Specifies whether to filter the raw data. | The value is of the Boolean type. The default value is True, which indicates to filter raw data. |

An example search and analytic statement is shown as follows:

```
* | select ts_predicate_arima(stamp, value, 3, 1, 2, 4) from (select __time__ - __time__ % 60 as stamp, avg(v) as value from log GROUP BY stamp order by stamp)
```

 **Note** The response is similar to that of the `ts_predicate_simple` function. For more information, see the response of the `ts_predicate_simple` function.

ts_regression_predict

Syntax:

```
select ts_regression_predict(x, y, nPred, algotype, processType)
```

The following table lists the parameters of the function.

| Parameter | Description | Value |
|-----------|--|---|
| <i>x</i> | The time sequence. The time points along the x axis are sorted in the ascending order. | Each time point is a Unix timestamp. Unit: seconds. |

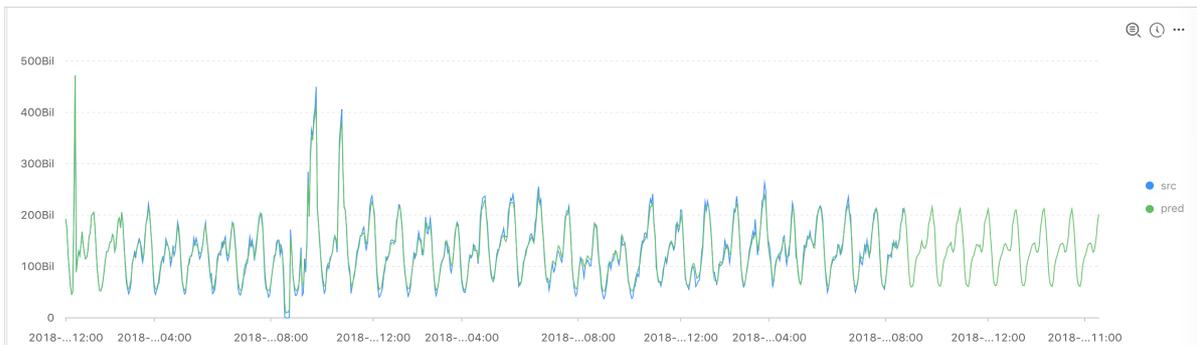
| Parameter | Description | Value |
|--------------------|--|---|
| <i>y</i> | The sequence of numeric data at each specified time point. | N/A |
| <i>nPred</i> | The number of points for prediction. | The value is of the LONG data type. Valid values: [1, 500]. |
| <i>algotype</i> | The algorithm type for prediction. | Valid values: <ul style="list-style-type: none"> • origin: uses the Gradient Boosted Regression Tree (GBRT) algorithm for prediction. • forest: uses the GBRT algorithm for prediction based on the trend component decomposed by Seasonal and Trend decomposition using Loess (STL), and then uses the additive model to sum up the decomposed components and obtains the predicted data. • linear: uses the Linear Regression algorithm for prediction based on the trend components decomposed by STL, and then uses the additive model to sum up the decomposed components and obtains the predicted data. |
| <i>processType</i> | Specifies whether to preprocess the data. | Valid values: <ul style="list-style-type: none"> • 0: no additional data preprocessing is performed. • 1: removes abnormal data before prediction. |

Example:

- A search and analytic statement is shown as follows:

```
* and h : nu2h05202.nu8 and m: NET | select ts_regression_predict(stamp, value, 200, 'origin') from (select __time__ - __time__ % 60 as stamp, avg(v) as value from log GROUP BY stamp order by stamp)
```

- The following figure shows the response.



The following table lists the display items.

| Display item | Description | |
|-----------------|-------------|--|
| Horizontal axis | unixtime | The Unix timestamp of the data. Unit: seconds. |
| Vertical axis | src | The raw data. |
| | predict | The predicted data. |

ts_anomaly_filter

Syntax:

```
select ts_anomaly_filter(lineName, ts, ds, preds, probs, nWatch, anomalyType)
```

The following table lists the parameters of the function.

| Parameter | Description | Value |
|--------------------|--|---|
| <i>lineName</i> | The name of each curve. The value is of the VARCHAR data type. | N/A |
| <i>ts</i> | The time sequence of the curve, which indicates the time of the current curve. The parameter value is an array of time points of the DOUBLE data type sorted in the ascending order. | N/A |
| <i>ds</i> | The actual value sequence of the curve. The parameter value is an array of data points with the same length as the ts parameter value. | N/A |
| <i>preds</i> | The predicted value sequence of the curve. The parameter value is an array of data points with the same length as the ts parameter value. | N/A |
| <i>probs</i> | The sequence of anomaly detection results of the curve. The parameter value is an array of data points with the same length as the ts parameter value. | N/A |
| <i>nWatch</i> | The number of the recently observed actual values on the curve. The value is of the LONG data type. The value must be smaller than the number of time points on the curve. | N/A |
| <i>anomalyType</i> | The type of anomaly to be filtered. The value is of the LONG data type. | Valid values: <ul style="list-style-type: none"> • 0: all anomalies. • 1: positive anomalies. • -1: negative anomalies. |

Example:

- A search and analytic statement is shown as follows:

```
* | select res.name, res.ts, res.ds, res.preds, res.probs
  from (
    select ts_anomaly_filter(name, ts, ds, preds, probs, cast(5 as bigint), cast(1 as bigint)) as res
  from (
    select name, res[1] as ts, res[2] as ds, res[3] as preds, res[4] as uppers, res[5] as lowers, res[6] as probs
  from (
    select name, array_transpose(ts_predicate_ar(stamp, value, 10)) as res
  from (
    select name, stamp, value from log where name like '%asg-%') group by name)) );
```

- The following figure shows the response.

```
| name          | ts                | ds      | preds | probs |
|-----|-----|-----|-----|-----|
| asg-bp1hylzdi2wx7civ0ivk | [1.5513696E9, 1.5513732E9, 1.5513768E9, 1.5513804E9] | [1,2,3,NaN] | [1,2,3,4] | [0,0,1,NaN]
```

23.4.9.7. Time series decomposition function

The time series decomposition function decomposes time series curves into curves that reveal the trend and periodicity of curves.

ts_decompose

Syntax:

```
select ts_decompose(x, y)
```

The following table lists the parameters of the function.

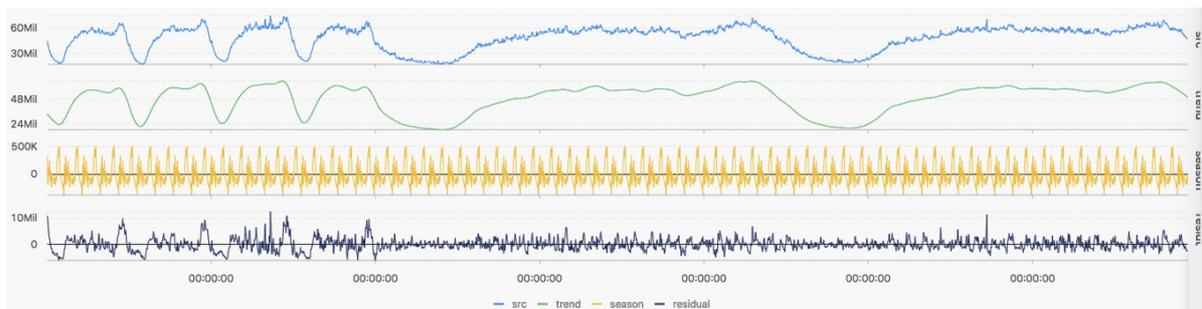
| Parameter | Description | Value |
|-----------|--|---|
| <i>x</i> | The time sequence. The time points along the x axis are sorted in the ascending order. | Each time point is a Unix timestamp. Unit: seconds. |
| <i>y</i> | The sequence of numeric data at each specified time point. | N/A |

Example:

- A search and analytic statement is shown as follows:

```
* |select ts_decompose(stamp, value) from (select __time__ - __time__ % 60 as stamp, avg(v) as value from log GROUP BY stamp order by stamp)
```

- The following figure shows the response.



The following table lists the display items.

| Display item | | Description |
|-----------------|----------|---|
| Horizontal axis | unixtime | The Unix timestamp of the data. Unit: seconds. |
| Vertical axis | src | The raw time series data. |
| | trend | The decomposed data that indicates the trend of the time series data. |
| | season | The decomposed data that indicates the periodicity of the time series data. |

| Display item | | Description |
|--------------|----------|---|
| | residual | The residual data decomposed from the time series data. |

23.4.9.8. Time series clustering functions

You can use time series clustering functions to cluster multiple time series and obtain different curve shapes. Then, you can find the cluster center and identify curves with shapes that are different from other curve shapes in the cluster in a timely manner.

Functions

| Function | Description |
|--------------------------------------|---|
| <code>ts_density_cluster</code> | Uses a density-based clustering method to cluster multiple time series. |
| <code>ts_hierarchical_cluster</code> | Uses a hierarchical clustering method to cluster multiple time series. |
| <code>ts_similar_instance</code> | Queries time series curves that are similar to a specified time series curve. |

ts_density_cluster

Syntax:

```
select ts_density_cluster(x, y, z)
```

The following table lists the parameters of the function.

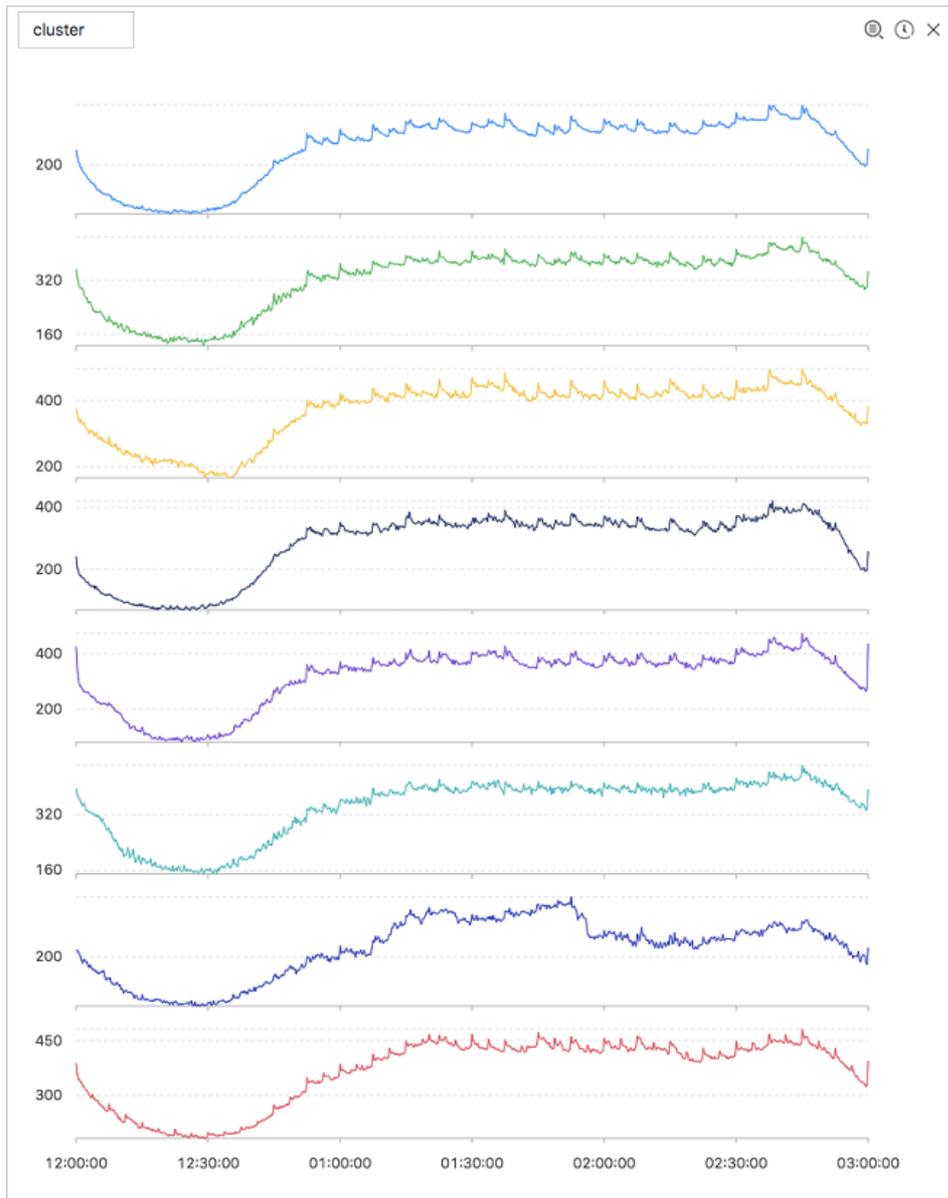
| Parameter | Description | Value |
|-----------|--|--|
| <i>x</i> | The time sequence. The points in time along the horizontal axis are sorted in ascending order. | Each point in time is a UNIX timestamp. Unit: seconds. |
| <i>y</i> | The sequence of numeric data that corresponds to a specified point in time. | - |
| <i>z</i> | The name of the curve corresponding to the data at a specified point in time. | The value is of the string type, for example, machine01.cpu_usr. |

Example

- The query statement:

```
* and (h: "machine_01" OR h: "machine_02" OR h: "machine_03") | select ts_density_cluster(stamp, metric_value, metric_name) from (select __time__ - __time__ % 600 as stamp, avg(v) as metric_value, h as metric_name from log GROUP BY stamp, metric_name order BY metric_name, stamp)
```

- Output result



The following table lists the display items.

| Display item | Description |
|----------------|--|
| cluster_id | The category of the cluster. The value -1 indicates that the cluster is not categorized in any cluster center. |
| rate | The proportion of instances in the cluster. |
| time_series | The timestamp sequence of the cluster center. |
| data_series | The data sequence of the cluster center. |
| instance_names | The instances that are included in the cluster center. |
| sim_instance | The name of an instance in the cluster. |

ts_hierarchical_cluster

Syntax:

```
select ts_hierarchical_cluster(x, y, z)
```

The following table lists the parameters of the function.

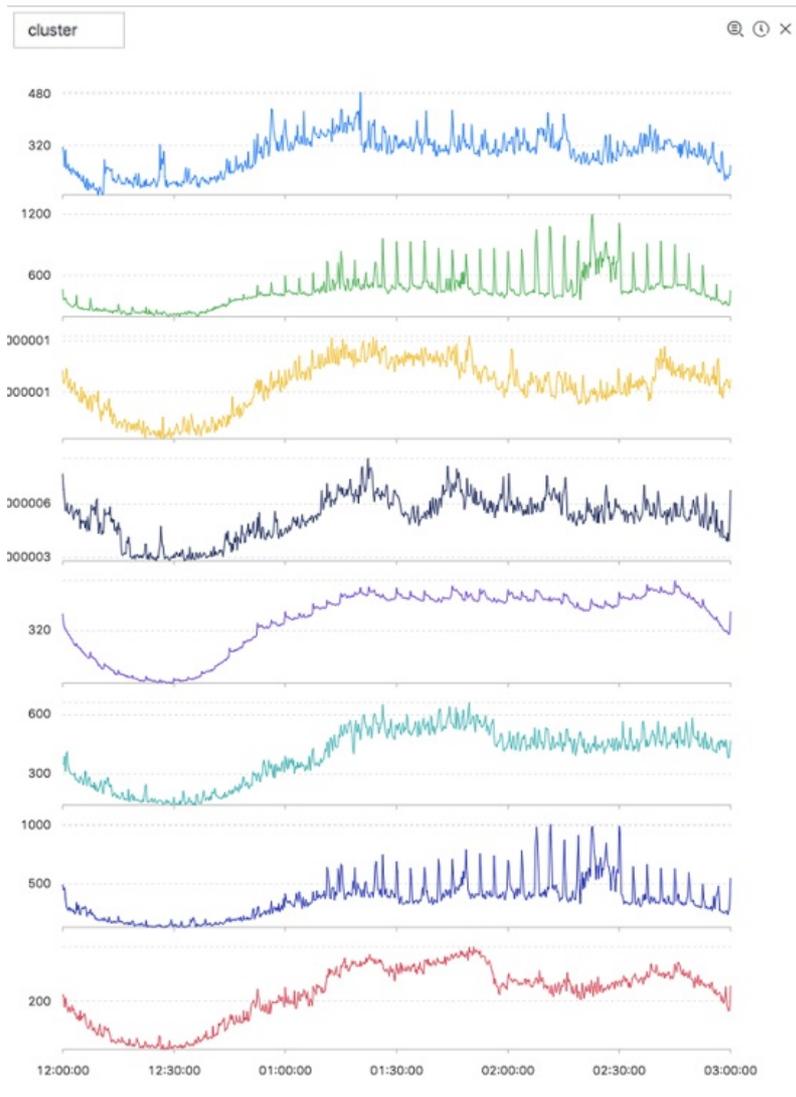
| Parameter | Description | Value |
|-----------|--|--|
| <i>x</i> | The time sequence. The points in time along the horizontal axis are sorted in ascending order. | Each point in time is a UNIX timestamp. Unit: seconds. |
| <i>y</i> | The sequence of numeric data that corresponds to a specified point in time. | - |
| <i>z</i> | The name of the curve corresponding to the data at a specified point in time. | The value is of the string type, for example, machine01.cpu_usr. |

Examples

- The query statement:

```
* and (h: "machine_01" OR h: "machine_02" OR h: "machine_03") | select ts_hierarchical_cluster(stamp, metric_value, metric_name) from ( select __time__ - __time__ % 600 as stamp, avg(v) as metric_value, h as metric_name from log GR OUP BY stamp, metric_name order BY metric_name, stamp )
```

- Output result



The following table lists the display items.

| Display item | Description |
|----------------|--|
| cluster_id | The category of the cluster. The value -1 indicates that the cluster is not categorized in any cluster center. |
| rate | The proportion of instances in the cluster. |
| time_series | The timestamp sequence of the cluster center. |
| data_series | The data sequence of the cluster center. |
| instance_names | The instances that are included in the cluster center. |
| sim_instance | The name of an instance in the cluster. |

ts_similar_instance

Syntax:

```
select ts_similar_instance(x, y, z, instance_name, topK, metricType)
```

The following table lists the parameters of the function.

| Parameter | Description | Value |
|----------------------|---|--|
| <i>x</i> | The time sequence. The points in time along the horizontal axis are sorted in ascending order. | Each point in time is a UNIX timestamp. Unit: seconds. |
| <i>y</i> | The sequence of numeric data that corresponds to a specified point in time. | - |
| <i>z</i> | The name of the curve corresponding to the data at a specified point in time. | The value is of the string type, for example, machine01.cpu_usr. |
| <i>instance_name</i> | The name of a specified curve to be queried. | The value is of the string type, for example, machine01.cpu_usr.  Note The curve to be queried must be an existing one. |
| <i>topK</i> | The maximum number of curves that are similar to the specified curve can be returned. | - |
| <i>metricType</i> | {'shape', 'manhattan', 'euclidean'} . The metric used to measure the similarity between time series curves. | - |

The query statement:

```
* and m: NET and m: Tcp and (h: "nu4e01524.nu8" OR h: "nu2i10267.nu8" OR h: "nu4q10466.nu8") | select ts_similar_instance(stamp, metric_value, metric_name, 'nu4e01524.nu8') from ( select __time__ - __time__ % 600 as stamp, sum(v) as metric_value, h as metric_name from log GROUP BY stamp, metric_name order BY metric_name, stamp )
```

The following table lists the display items.

| Display item | Description |
|---------------|---|
| instance_name | The list of metrics that are similar to the specified metric. |
| time_series | The timestamp sequence of the cluster center. |
| data_series | The data sequence of the cluster center. |

23.4.9.9. Frequent pattern statistics function

The frequent pattern statistics function combines representative attributes in a specified multi-attribute field sample.

pattern_stat

Syntax:

```
select pattern_stat(array[col1, col2, col3], array['col1_name', 'col2_name', 'col3_name'], array[col5, col6], array['col5_name', 'col6_name'], support_score, sample_ratio)
```

The following table lists the parameters of the function.

| Parameter | Description | Value |
|---|--|--|
| <i>array[col1, col2, col3]</i> | A column of character values. | An array of values, for example, array[clientIP, sourceIP, path, logstore]. |
| <i>array['col1_name', 'col2_name', 'col3_name']</i> | The field names of the character values. | An array of field names, for example, array['clientIP', 'sourceIP', 'path', 'logstore']. |
| <i>array[col5, col6]</i> | A column of numeric values. | An array of values, for example, array[Inflow, OutFlow]. |
| <i>array['col5_name', 'col6_name']</i> | The field names of the numeric values. | An array of field names, for example, array['Inflow', 'OutFlow']. |
| <i>support_score</i> | The support ratio of samples for pattern mining. | The value is of the DOUBLE data type. Value range: (0,1]. |
| <i>sample_ratio</i> | The sampling ratio. The default value is 0.1, which indicates that only 10% of the total samples are used. | The value is of the DOUBLE data type. Value range: (0,1]. |

Example:

- Query statement

```
* | select pattern_stat(array[ Category, ClientIP, ProjectName, LogStore, Method, Source, UserAgent ], array[ 'Category', 'ClientIP', 'ProjectName', 'LogStore', 'Method', 'Source', 'UserAgent' ], array[ InFlow, OutFlow ], array[ 'InFlow', 'OutFlow' ], 0.45, 0.3) limit 1000
```

- Display item

| Display item | Description |
|---------------|---|
| count | The number of samples in the current pattern. |
| support_score | The score of the current pattern. The score indicates the degree to which the current pattern is supported. |
| pattern | The content of the pattern. The pattern is organized in the format that is defined by the query conditions. |

23.4.9.10. Differential pattern statistics function

The differential pattern statistics function analyzes differential patterns of specified multi-field samples based on the specified condition. It helps you identify the causes of the differences under the current condition in a timely manner.

pattern_diff

Syntax:

```
select pattern_diff(array_char_value, array_char_name, array_numeric_value, array_numeric_name, condition, supportScore, posSampleRatio, negSampleRatio )
```

The following table lists the parameters of the function.

| Parameter | Description | Value |
|----------------------------|--|--|
| <i>array_char_value</i> | A column of values of the character data type. | An array of values, for example, array[clientIP, sourceIP, path, logstore]. |
| <i>array_char_name</i> | The field names of the values of the character data type. | An array of field names, for example, array['clientIP', 'sourceIP', 'path', 'logstore']. |
| <i>array_numeric_value</i> | A column of values of the numeric data type. | An array of values, for example, array[Inflow, OutFlow]. |
| <i>array_numeric_name</i> | The field names of the values of the numeric data type. | An array of field names, for example, array[originflow', 'OutFlow']. |
| <i>condition</i> | The condition for filtering data. The value True indicates positive samples, and the value False indicates negative samples. | For example, Latency <= 300. |
| <i>supportScore</i> | The support ratio of positive and negative samples for pattern mining. | The value is of the DOUBLE data type. Valid values: (0,1]. |
| <i>posSampleRatio</i> | The sampling ratio of positive samples. The default value is 0.5, indicating that 50% of positive samples are collected. | The value is of the DOUBLE data type. Valid values: (0,1]. |
| <i>negSampleRatio</i> | The sampling ratio of negative samples. The default value is 0.5, indicating that 50% of positive samples are collected. | The value is of the DOUBLE data type. Valid values: (0,1]. |

Example:

- A search and analytic statement is shown as follows:

```
* |select pattern_diff(array[ Category, ClientIP, ProjectName, LogStore, Method, Source, UserAgent ], array[ 'Category', 'ClientIP', 'ProjectName', 'LogStore', 'Method', 'Source', 'UserAgent' ], array[ InFlow, OutFlow ], array[ 'InFlow', 'OutFlow' ], Latency > 300, 0.2, 0.1, 1.0) limit 1000
```

- Display item

| Display item | Description |
|---------------|---|
| possupport | The support ratio of positive samples for the mined patterns. |
| posconfidence | The confidence level of the mined patterns in positive samples. |
| negsupport | The support ratio of negative samples for the mined patterns. |
| diffpattern | The content of the mined patterns. |

23.4.9.11. Root cause analysis function

Log Service provides the alert and analytics features that help you quickly analyze data and locate anomalies of specific subdimensions of a metric. You can use the root cause analysis function to analyze the subdimension attributes that result in anomalies of the monitoring metric.

rca_kpi_search

Syntax

```
select rca_kpi_search(varchar_array, name_array, real, forecast, level)
```

The following table lists the parameters of the function.

| Parameter | Description | Value |
|----------------------|--|--|
| <i>varchar_array</i> | The subdimension attributes. | The parameter value is formatted in an array, for example, array[col1, col2, col3]. |
| <i>name_array</i> | The subdimension attribute names. | The parameter value is formatted in an array, for example, array['col1', 'col2', 'col3']. |
| <i>real</i> | The actual value of each subdimension attribute specified by the varchar_array parameter. | The parameter value is of the DOUBLE data type. Valid values: all real numbers. |
| <i>forecast</i> | The predicted value of each subdimension attribute specified by the varchar_array parameter. | The parameter value is of the DOUBLE data type. Valid values: all real numbers. |
| <i>level</i> | The number of subdimension attributes identified in the returned root cause set. The value 0 indicates that all root causes that are found are returned. | The parameter value is of the LONG data type. Valid values: [0, number of analyzed subdimensions]. The number of analyzed subdimensions is the length of the array specified by the varchar_array parameter. |

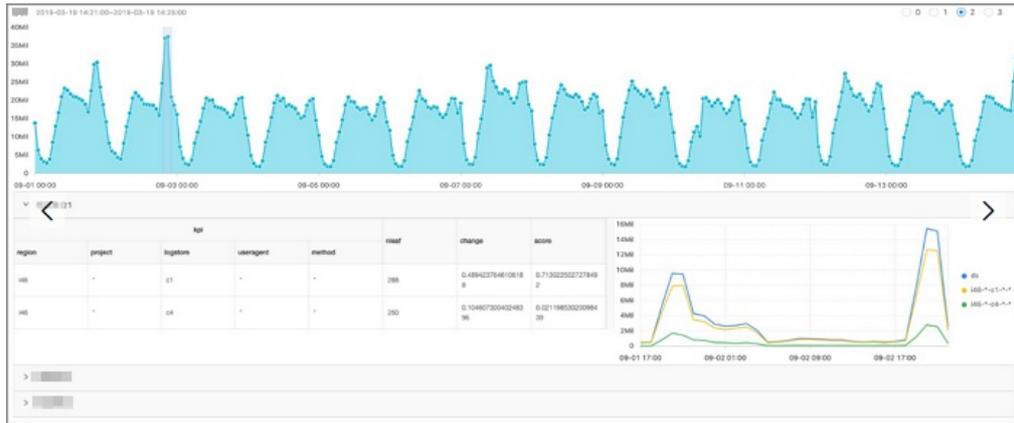
Example:

- The search and analytic statement is shown as follows:

Use a subquery to obtain the actual value and predicted value of each subdimension attribute, and then call the rca_kpi_search function to analyze the root causes of anomalies.

```
* not Status:200 |
select rca_kpi_search(
  array[ ProjectName, LogStore, UserAgent, Method ],
  array[ 'ProjectName', 'LogStore', 'UserAgent', 'Method' ], real, forecast, 1)
from (
  select ProjectName, LogStore, UserAgent, Method,
  sum(case when time < 1552436040 then real else 0 end) * 1.0 / sum(case when time < 1552436040
  then 1 else 0 end) as forecast,
  sum(case when time >=1552436040 then real else 0 end) *1.0 / sum(case when time >= 1552436040
  then 1 else 0 end) as real
  from (
  select __time__ - __time__ % 60 as time, ProjectName, LogStore, UserAgent, Method, COUNT(*) as real
  from log GROUP by time, ProjectName, LogStore, UserAgent, Method )
  GROUP BY ProjectName, LogStore, UserAgent, Method limit 10000000)
```

- The following figure shows the response.



The following figure shows the structured response.

```

{
  "rcSets": [
    {
      "rcItems": [
        {
          "kpi": [{"attr": "xxx", "val": "xxx"}],
          "nleaf": 100,
          "change": 0.524543,
          "score": 0.1454543
        }
      ]
    }
  ]
}
    
```

The following table lists the display items.

| Display item | Description |
|----------------|---|
| <i>rcSets</i> | The root cause sets. Each value of this parameter is an array. |
| <i>rcItems</i> | A specific root cause set. |
| <i>kpi</i> | A specific item in the root cause set. Each item is formatted in an array where each element is a JSON object. The attr parameter indicates the subdimension name, and the val parameter indicates the attribute name under the subdimension. |
| <i>nleaf</i> | The number of leaf nodes that an item (KPI) in the root cause set covers in the original data. <div style="border: 1px solid #ccc; background-color: #e0f2f1; padding: 5px;"> <p>Note Leaf node: the log entry that contains the finest-grained attribute information.</p> </div> |
| <i>change</i> | The ratio of anomalies of leaf nodes in a KPI to the total anomalies in the root cause set that occurred at the same time point. |
| <i>score</i> | The abnormality score of the current KPI. Valid values: [0, 1]. |

The response is formatted in a JSON object as follows:

```
{
  "rcSets": [
    {
      "rcItems": [
        {
          "kpi": [
            {
              "attr": "country",
              "val": "*"
            },
            {
              "attr": "province",
              "val": "*"
            },
            {
              "attr": "provider",
              "val": "*"
            },
            {
              "attr": "domain",
              "val": "download.huya.com"
            },
            {
              "attr": "method",
              "val": "*"
            }
          ]
        },
        {
          "nleaf": 119,
          "change": 0.3180687806279939,
          "score": 0.14436007709620113
        }
      ]
    }
  ]
}
```

23.4.9.12. Correlation analysis functions

You can use a correlation analysis function to find the metrics that are correlated with a specified metric or time series data among multiple observed metrics in the system.

Functions

| Function | Description |
|--------------------------------------|--|
| <code>ts_association_analysis</code> | Quickly finds the metrics that are correlated with a specified metric among multiple observed metrics in the system. |
| <code>ts_similar</code> | Quickly finds the metrics that are correlated with specified time series data among multiple observed metrics in the system. |

ts_association_analysis

Syntax

```
select ts_association_analysis(stamp, params, names, indexName, threshold)
```

The following table lists the parameters of the function.

| Parameter | Description | Value |
|------------------|--|--|
| <i>stamp</i> | The Unix timestamp of the LONG data type. | - |
| <i>params</i> | The metrics to be analyzed, formatted in an array where each element is of the DOUBLE data type. | The parameter value is formatted in an array where each element is of the DOUBLE data type. For example, Latency, QPS, and NetFlow. |
| <i>names</i> | The names of the metrics to be analyzed. | The parameter value is formatted in an array where each element is of the VARCHAR data type. For example, Latency, QPS, and NetFlow. |
| <i>indexName</i> | The name of the target metric. | The parameter value is of the VARCHAR data type, for example, Latency. |
| <i>threshold</i> | The threshold of correlation between the metrics to be analyzed and the target metric. | The parameter value is of the DOUBLE data type. Valid values: [0, 1]. |

Response

- name: the name of the metric that meets the specified correlation condition with the target metric.
- score: the value of correlation between the returned metric and the target metric. Valid values: [0, 1].

Sample statement

```
* | select ts_association_analysis(
  time,
  array[inflow, outflow, latency, status],
  array['inflow', 'outflow', 'latency', 'status'],
  'latency',
  0.1) from log;
```

Sample response

```
| results      |
|-----|
| ['latency', '1.0'] |
| ['outflow', '0.6265'] |
| ['status', '0.2270'] |
```

ts_similar

Syntax 1

```
select ts_similar(stamp, value, ts, ds)
select ts_similar(stamp, value, ts, ds, metricType)
```

The following table lists the parameters of the function.

| Parameter | Description | Value |
|-------------------|---|--|
| <i>stamp</i> | The Unix timestamp of the LONG data type. | - |
| <i>value</i> | The value of the metric to be analyzed. The parameter value is of the DOUBLE data type. | - |
| <i>ts</i> | The time sequence of the specified time series curve. The parameter value is formatted in an array where each element is of the DOUBLE data type. | - |
| <i>ds</i> | The sequence of numeric data of the specified time series curve. | - |
| <i>metricType</i> | The type of correlation between the measured curves. The parameter value is of the VARCHAR data type. | Valid values: SHAPE, RMSE, PEARSON, SPEARMAN, R2, and KENDALL |

Syntax 2

```
select ts_similar(stamp, value, startStamp, endStamp, step, ds)
select ts_similar(stamp, value, startStamp, endStamp, step, ds, metricType )
```

The following table lists the parameters of the function.

| Parameter | Description | Value |
|-------------------|--|-------|
| <i>stamp</i> | The Unix timestamp of the LONG data type. | - |
| <i>value</i> | The value of the metric to be analyzed. This parameter is of the DOUBLE data type. | - |
| <i>startStamp</i> | The start timestamp of the specified time series curve. The parameter value is of the LONG data type. | - |
| <i>endStamp</i> | The end timestamp of the specified time series curve. The parameter value is of the LONG data type. | - |
| <i>step</i> | The time interval between two adjacent data points in a time series. The parameter value is of the LONG data type. | - |
| <i>ds</i> | The sequence of numeric data of the specified time series curve. The parameter is formatted in an array where each element is of the DOUBLE data type. | - |

| Parameter | Description | Value |
|-------------------|---|--|
| <i>metricType</i> | The type of correlation between the measured curves. The parameter value is of the VARCHAR data type. | Valid values: SHAPE, RMSE, PEARSON, SPEARMAN, R2, and KENDALL |

- Response
score: the correlation between the analyzed metric and the specified time series curve. Valid values: [-1, 1].

- Sample statement

```
* | select vhost, metric, ts_similar(time, value, 1560911040, 1560911065, 5, array[5.1,4.0,3.3,5.6,4.0,7.2], 'PEARSON') from log group by vhost, metric;
```

- Sample response

```
| vhost | metric      | score          |
|-----|-----|-----|
| vhost1 | redolog     | -0.3519082537204182 |
| vhost1 | kv_qps      | -0.15922168009772697 |
| vhost1 | file_meta_write | NaN          |
```

23.4.9.13. Kernel density estimation function

Kernel density estimation (KDE) is a non-parametric way to estimate the probability density function of a random variable.

The Kernel density estimation function uses the smooth peak function to fit the observed data points. In this way, the function simulates the real probability distribution curve.

- Syntax

```
select kernel_density_estimation(bigint stamp, double value, varchar kernelType)
```

- Parameters

| Parameter | Description |
|------------|--|
| stamp | The Unix timestamp of observed data. Unit: second. |
| value | The observed value. |
| kernelType | <ul style="list-style-type: none"> ◦ box: rectangle window. ◦ epanechnikov: Epanechnikov curve. ◦ gausener: Gaussian curve. |

- Response

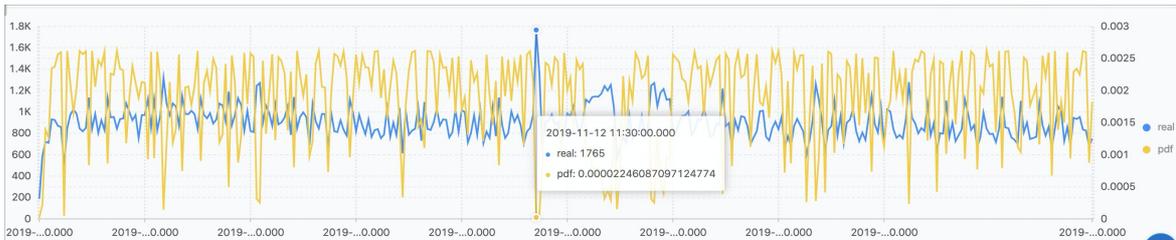
| Display item | Description |
|--------------|--|
| unixtime | The Unix timestamp of observed data. |
| real | The observed value. |
| pdf | The probability of each observed data point. |

- Example

- Sample statement

```
* |
select
  date_trunc('second', cast(t1[1] as bigint)) as time, t1[2] as real, t1[3] as pdf from (
  select kernel_density_estimation(time, num, 'gaussian') as res from (
  select __time__ - __time__ % 10 as time, COUNT(*) * 1.0 as num from log group by time order by time
  ), unnest(res) as t(t1) limit 1000
```

- Response



23.4.10. Advanced analysis

23.4.10.1. Optimize queries

This topic describes how to optimize queries to improve query efficiency.

You can use the following methods to optimize queries:

- Increase the number of shards.
- Reduce the query time range and data volume.
- Repeat queries multiple times.
- Optimize the SQL statement for queries.

Increase the number of shards

More shards represent more computing resources and faster computing speed. You can increase the number of shards to ensure that the average number of log entries to be scanned in each shard does not exceed 50 million. You can increase the number of shards by splitting shards. For more information, see [Split a shard](#).

Note Splitting shards incurs more fees and only accelerates queries of new data. Existing data is still stored in old shards.

Reduce the query time range and data volume

- The larger the time range, the slower the query. If you query data within a year or a month, data is computed by day. To facilitate computing, you can reduce the query time range.
- The larger the data volume, the slower the query. Reduce the amount of data to be queried as much as possible.

Repeat queries multiple times

If you find that the result of a query is inaccurate, you can repeat the query multiple times. The underlying acceleration mechanism ensures that each query uses the previous query result to analyze data. In this way, multiple queries make the query result more accurate.

Optimize the SQL statement for queries

A time-consuming query statement has the following characteristics:

- Performs the GROUP BY operation on string-type columns.
- Performs the GROUP BY operation on more than five fields.
- Includes operations that generates strings.

You can use the following methods to optimize a query statement:

- Avoid operations that generate strings if possible.
 - If you use the date_format function to generate a formatted timestamp, the query is inefficient.

```
* | select date_format(from_unixtime(__time__), '%H_%i') as t, count(1) group by t
```

- If you use the substr() function, strings are generated. We recommend that you use the date_trunc or time_series function in a query statement.
- Avoid performing the GROUP BY operation on string-formatted columns if possible.

Performing the GROUP BY operation on strings may result in a large number of hash calculations, which account for more than 50% of total calculations. Examples:

```
* | select count(1) as pv, date_trunc('hour', __time__) as time group by time
* | select count(1) as pv, from_unixtime(__time__ - __time__%3600) as time group by __time__ - __time__%3600
```

Both query 1 and query 2 count the number of log entries per hour. However, query 1 converts the time into a string, for example, 2017-12-12 00:00:00, and then performs the GROUP BY operation on this string. Query 2 calculates the on-the-hour time value, performs the GROUP BY operation on the result, and then converts the value into a string. Query 1 is less efficient than query 2 because query 1 needs to hash strings.

- List fields alphabetically based on the initial letter when performing the GROUP BY operation on multiple columns.

For example, you need to query 100 million users who are from 13 provinces.

```
Fast: * | select province,uid,count(1)groupby province,uid
Slow: * | select province,uid,count(1)groupby uid,province
```

- Use estimating functions.

Estimating functions provide stronger performance than accurate calculation. In estimation, accuracy is compromised to an acceptable extent for fast calculation.

```
Fast: * |select approx_distinct(ip)
Slow: * |select count(distinct(ip))
```

- Specify only required columns in the SQL statement if possible.

You can specify all columns in the search statement. In the SQL statement, specify only required columns if possible. This will speed up calculation.

```
Fast: * |select a,b,c
Slow: * |select *
```

- Place columns that do not need to be grouped in an aggregate function if possible.

For example, a user ID is associated with a username. Therefore, you can execute the Group By operation on user IDs to analyze data.

```
Fast: * | select userid, arbitrary(username), count(1)groupby userid
Slow: * | select userid, username, count(1)groupby userid,username
```

- Avoid using the IN operator if possible.

If possible, avoid using the IN clause in SQL statements. Instead, use the OR clause.

```
Fast: key : a or key :b or key:c | select count(1)
Slow: * | select count(1) where key in ('a','b')
```

23.4.10.2. Use cases

This topic provides some use cases of log data analysis.

Trigger an alert when the error rate exceeds 40% over the last 5 minutes

Calculate the percentage of 500 Internal Server Error every minute. An alert is triggered when the error rate exceeds 40% over the last 5 minutes.

```
status:500 | select __topic__, max_by(error_count>window_time)/1.0/sum(error_count) as error_ratio, sum(error_count) as total_error from (
select __topic__, count(*) as error_count, __time__ - __time__ % 300 as window_time from log group by __topic__, window_time
)
group by __topic__ having max_by(error_count>window_time)/1.0/sum(error_count) > 0.4 and sum(error_count) > 500
order by total_error desc limit 100
```

Calculate the amount of transferred data and configure alerts

Calculate the amount of transferred data every minute. An alert is triggered when transferred data plunges. Transferred data counted in the last minute does not cover a full minute. The `(max(time) - min(time))` clause is used for normalization to count the average traffic per minute.

```
* | SELECT SUM(inflow) / (max(__time__) - min(__time__)) as inflow_per_minute, date_trunc('minute', __time__) as minute group by minute
```

Calculate the average latency of traffic data in different sizes

Distribute traffic data to multiple bins based on the data size and calculate the average latency of the data in the bins.

```
* | select avg(latency) as latency, case when originSize < 5000 then 's1' when originSize < 20000 then 's2' when originSize < 500000 then 's3' when originSize < 100000000 then 's4' else 's5' end as os group by os
```

Retrieve the percentages of different results

List the number and the percentage of each result for different departments. This query includes subqueries and window functions. The `sum(c) over()` clause indicates the sum of values in all rows.

```
* | select department, c*1.0/sum(c) over () from(select count(1) as c, department from log group by department)
```

Count the number of log entries that meet the query condition

To count the number of URLs based on their characteristics, you can use the CASE WHEN clause or the COUNT_IF clause. The latter clause is simpler.

```
* | select count_if(uri like '%login') as login_num, count_if(uri like '%register') as register_num, date_format(date_trunc('minute', __time__), '%m-%d %H:%i') as time group by time order by time limit 100
```

23.4.10.3. Time field conversion examples

During search and analytics, you often need to process time fields in log data, such as converting a timestamp to another time format. This topic uses some examples to describe how to convert time fields.

A log entry may include multiple time fields, for example:

- `__time__` : the time that you specify when you use the API or SDK to write log data. This field can be used for log data shipping, search, and analytics.
- Original time field in log data: the field that records the time when the log data is generated. This field is in raw logs.

Time fields in different formats are difficult to read. To simplify the read process, you can convert the time format during search and analytics. For example, you can perform the following conversions:

1. [Convert `__time__` to a timestamp](#)
2. [Display `__time__` in a specified format](#)
3. [Convert a timestamp to a specified format](#)

Convert `__time__` to a timestamp

You can use the `from_unixtime` function to convert the `__time__` field to a timestamp.

```
* | select from_unixtime(__time__)
```

Display `__time__` in a specified format

To display the `__time__` field in the format of `YYYY-MM-DD HH:MM:SS`, you can use the `date_format` function.

```
* | select date_format(__time__, '%Y-%m-%d %H:%i:%S')
```

Convert the time in a log to a specified format

To convert the time field in a log to the specified format (`YYYY-MM-DD HH:MM:SS`) and perform the GROUP BY operation on the `YYYY-MM-DD` part, you can use the `date_format` function.

- Sample log entry

```
__topic__:
body_byte_sent: 307
hostname: www.host1.com
http_user_agent: Mozilla/5.0 (iPhone; CPU iPhone OS 10_3_3 like Mac OS X) AppleWebKit/603.3.8 (KHTML, like Gecko)
Mobile/14G60 QQ/7.1.8.452 V1_IPH_SQ_7.1.8_1_APP_A Pixel/750 Core/UIWebView NetType/WIFI QBWebViewType/1
method: GET
referer: www.host0.com
remote_addr: 36.63.1.23
request_length: 111
request_time: 2.705
status: 200
upstream_response_time: 0.225582883754
url: /? k0=v9&
time:2017-05-17 09:45:00
```

- Example SQL statement

```
* | select date_format (date_parse(time,'%Y-%m-%d %H:%i:%S'), '%Y-%m-%d') as day, count(1) as uv group by day order by day asc
```

23.4.11. Visual analysis

23.4.11.1. Analysis graph

23.4.11.1.1. Overview

All search and analytics results can be rendered by using visualized charts.

Prerequisites

- The index feature is enabled and configured. The analytics switches are turned on. For more information, see [Enable the index feature and configure indexes for a Logstore](#).
- An analytic statement is included in a query statement. You cannot use charts to show query results if you do not include an analytic statement in your query statement.

Precautions

When multiple search and analytic statements are being executed in sequence, the **Value Column**, **X Axis**, or **Y Axis** information cannot automatically change based on the search and analytic statement. The X and Y axis information may remain the same as the last search and analytic statement. If this happens, the query results of the current search and analytic statement cannot be automatically displayed in a chart. If the following messages are returned, configure parameters on the **Properties** tab based on the current search and analytic statement:

- The currently selected dimensions are not in the queried results. Check and configure the attributes.
- X-Axis or Y-Axis is not available. Check and configure the attributes.

Chart configurations

On the **Graph** tab, various charts are provided to show query results. You can select a type of chart from the chart bar to show results.

- On the **Graph** tab, you can view the **Chart Preview** and **Data Preview** of query results of the current search and analytic statement. **Chart Preview** is the preview of the query results that are displayed in the specified type of chart. **Data Preview** displays the query results in a table.
- On the **Graph** tab on the right, you can configure the following chart properties:

- **Data Source**: used to set placeholder variables. For example, you configure the drill-down event of Chart A to redirect to the dashboard where Chart B is located. The placeholder variable you configured for Chart B is the same as the variable that you click to trigger the drill-down event. Then the placeholder variable is replaced with the variable you click to trigger the drill-down event and the search and analytic statement of Chart B is executed. For more information, see [Drill-down analysis](#).

This feature is applicable to scenarios where you configure drill-down events to redirect to targeted dashboards.

- **Properties**: used to configure the display properties of a chart, including the X axis, left and right Y axes, margins, font size and other properties. The properties vary with different type of charts.

This feature is applicable to all search and analytics scenarios.

- **Interactive Behavior**: used to configure drill-down events for a chart. After you configure a drill-down event for the chart, you can click the variable value in the chart to trigger the specified drill-down event. For more information, see [Drill-down analysis](#).

This feature is applicable to triggering drill-down events for charts.

23.4.11.1.2. Display query results on a table

Tables are used to sort and display data for quick reference and analysis. All query results that match specified query statements can be rendered into visualized charts. By default, the query results are displayed in a table.

Components

- Table header
- Row
- Column

Where:

- The number of columns can be specified by using a `SELECT` statement.
- The number of rows is calculated based on the number of log entries in a specified time range. The default clause is `LIMIT 100`.

Procedure

1. On the Search & Analysis page, enter a query statement in the search box, specify a time range, and then click **Search & Analyze**.
2. On the **Graph** tab, data is displayed in a table by default. You do not need to click the  icon.
3. On the **Properties** tab, configure the properties of the table.

Properties

| Parameter | Description |
|----------------------------|---|
| Items per Page | The number of log entries to return on each page. |
| Zebra Striping | Specifies whether to display the query results in a zebra-striped table. |
| Transpose Rows and Columns | Specifies whether to transpose rows and columns. |
| Hide Reserved Fields | Specifies whether to hide reserved fields. |
| Disable Sorting | Specifies whether to disable the sorting feature. |
| Disable Search | Specifies whether to disable the search feature. |
| Highlight Settings | The rules for highlighting rows or columns that conform to specified rules. |

23.4.11.1.3. Display query results on a line chart

A line chart is used to analyze the value changes of fields based on an ordered data type. In most cases, this analysis is based on a specified time range.

You can use a line chart to analyze the following change characteristics of field values over a specified period:

- Increment or decrement
- Increment or decrement rate
- Increment or decrement pattern, for example, periodicity
- Peak value and trough value

Line charts are used to analyze field value changes over a time period. You can also use a line chart to analyze the value changes of multiple fields in multiple lines over the same time period. Then, you can analyze the relationships between the different fields. For example, the values of multiple fields can display positive, negative, or inverse trends.

Components

- X-axis
- Left Y-axis
- Right Y-axis (optional)
- Data point
- Line of trend change

- Legend

Procedure

1. On the Search & Analysis page, enter a query statement in the search box, specify a time range, and then click **Search & Analyze**.
2. On the **Graph** tab, click the  icon.
3. On the **Properties** tab, configure the properties of the line chart.

 **Note** In a line chart, each line must contain more than two data points. Otherwise, the data trend cannot be analyzed. We recommend that you select five or fewer lines in a line chart.

Properties

| Parameter | Description |
|---------------------|--|
| X Axis | The sequential data. In most cases, time series is selected. |
| Left Y Axis | The numeric data. You can select one or more fields for the left Y-axis. |
| Right Y Axis | The numeric data. You can select one or more fields for the right Y-axis. The layer of the right Y-axis is higher than that of the left Y-axis. |
| Column Marker | The column on the left or right Y-axis. The column is selected as a histogram. |
| Legend | The position where the legend is located in the chart. Valid values: Top, Bottom, Left, and Right. |
| Format Left Y-axis | The format in which data selected for the left Y-axis and right Y-axis is displayed. |
| Format Right Y-axis | |
| Margin | The distance between an axis and the borders of the chart. Valid values: Top Margin , Bottom Margin , Right Margin , and Left Margin . |

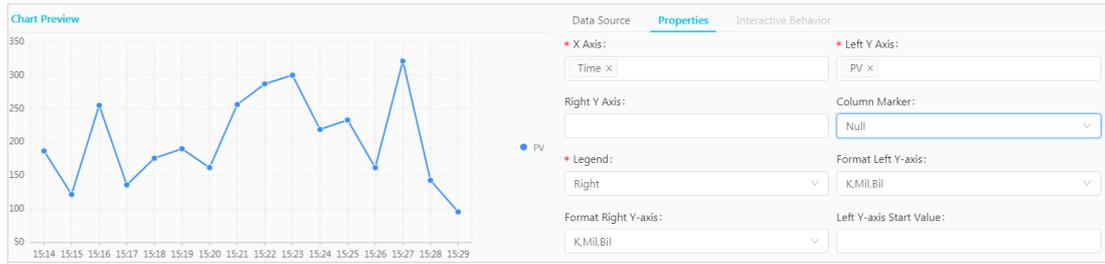
Example of a line chart

To query the page views (PVs) of the IP address `10.0.192.0` in the last 24 hours, execute the following query statement:

```
remote_addr: 10.0.192.0 | select date_format(date_trunc('hour', __time__), '%m-%d %H:%i') as time, count(1) as PV group by time order by time limit 1000
```

Select `time` for the X-axis, `PV` for the left Y-axis, and `Bottom` for Legend. Adjust the margins based on your business requirements.

Line chart



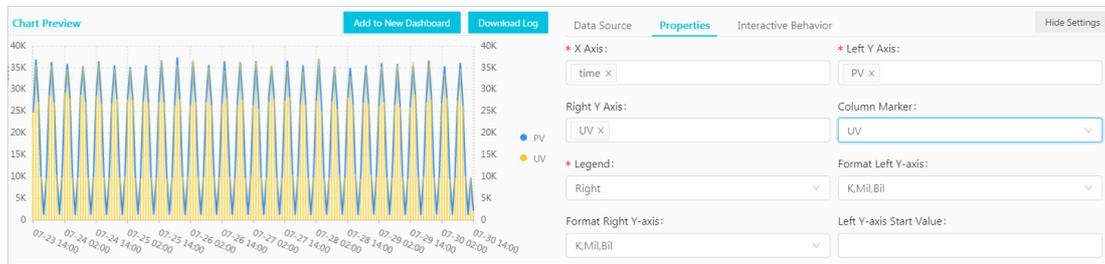
Example of a dual Y-axis line chart

To query the PVs and unique visitors (UVs) in the last 24 hours, execute the following query statement:

```
* | select date_format(date_trunc('hour', __time__), '%m-%d %H:%i') as time, count(1) as PV, approx_distinct(remote_addr) as UV group by time order by time limit 1000
```

Select **time** for the X-axis, **PV** for the left Y-axis, **UV** for the right Y-axis, and **PV** for Column Marker.

Dual Y-axis line chart



23.4.11.1.4. Display query results on a column chart

A column chart uses vertical or horizontal bars to present categorical values. Compared with a line chart, a column chart does not display ordered data, but provides a method to count the number of values in each category.

Components

- X-axis (horizontal)
- Y-axis (vertical)
- Rectangular bar
- Legend

By default, column charts in Log Service use vertical bars. Each rectangular bar has a fixed width and a variable height that indicates a value. You can use a grouped column chart to display the data if multiple columns of data are mapped to the Y-axis.

Procedure

1. On the Search & Analysis page, enter a query statement in the search box, specify a time range, and then click **Search & Analyze**.
2. On the **Graph** tab, click the  icon.
3. On the **Properties** tab, configure the properties of the column chart.

Note Column charts can be used to display query results if the number of returned log entries is less than 20. You can use a **LIMIT** clause to control the number of rectangular bars. Analysis results may not be clearly displayed if the chart contains a large number of rectangular bars. In addition, we recommend that you select less than five fields for the Y-axis.

Properties

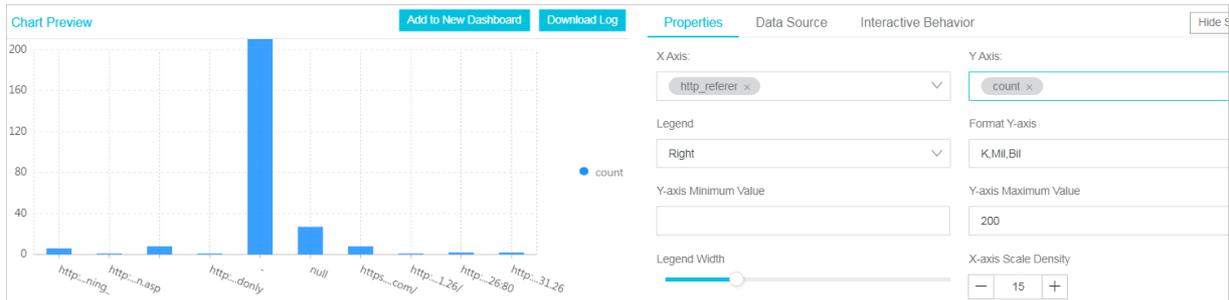
| Parameter | Description |
|-----------|--|
| X Axis | The categorical data. |
| Y Axis | The numeric data. You can select one or more fields for the Y-axis. |
| Legend | The position where the legend is located in the chart. Valid values: Top, Bottom, Left, and Right. |
| Format | The format in which data is displayed on the Y-axis. |
| Margin | The distance between an axis and the borders of the chart. Valid values: Top Margin, Bottom Margin, Right Margin, and Left Margin. |

Example of a simple column chart

To query the number of visits for each `http_referer` in the specified time range, execute the following query statement:

```
* | select http_referer, count(1) as count group by http_referer
```

Select `http_referer` for X-axis and `count` for Y-axis.



Example of a grouped column chart

To query the number of requests and the average bytes for each `http_referer` in the specified time range, execute the following statement:

```
* | select http_referer, count(1) as count, avg(body_bytes_sent) as avg group by http_referer
```

Select `http_referer` for X-axis. Select `count` and `avg` for Y-axis.



23.4.11.1.5. Display query results on a bar chart

A bar chart is a horizontal column chart that is used to analyze the top N values of fields. A bar chart is configured in a similar way to a column chart.

Components

- X-axis (vertical)
- Y-axis (horizontal)
- Rectangular bar
- Legend

Each rectangular bar has a fixed height and a variable width. The variable width indicates a value. You can use a grouped bar chart to display the data if multiple columns of data are mapped to the Y-axis.

Procedure

1. On the Search & Analysis page, enter a query statement in the search box, specify a time range, and then click **Search & Analyze**.
2. On the **Graph** tab, click the  icon.
3. On the **Properties** tab, configure the properties of the bar chart.

Note

- Bar charts can be used to display query results if 20 or fewer log entries are returned. You can use the `LIMIT` clause to control the number of rectangular bars. Analysis results may not be clearly displayed if the chart contains a large number of rectangular bars. You can use the `ORDER BY` clause to analyze the top N values of fields. We recommend that you map less than five columns of data to the Y-axis.
- You can use a grouped bar chart to display query results. However, the values represented by each rectangular bar in a group must be positively or negatively associated with each other.

Properties

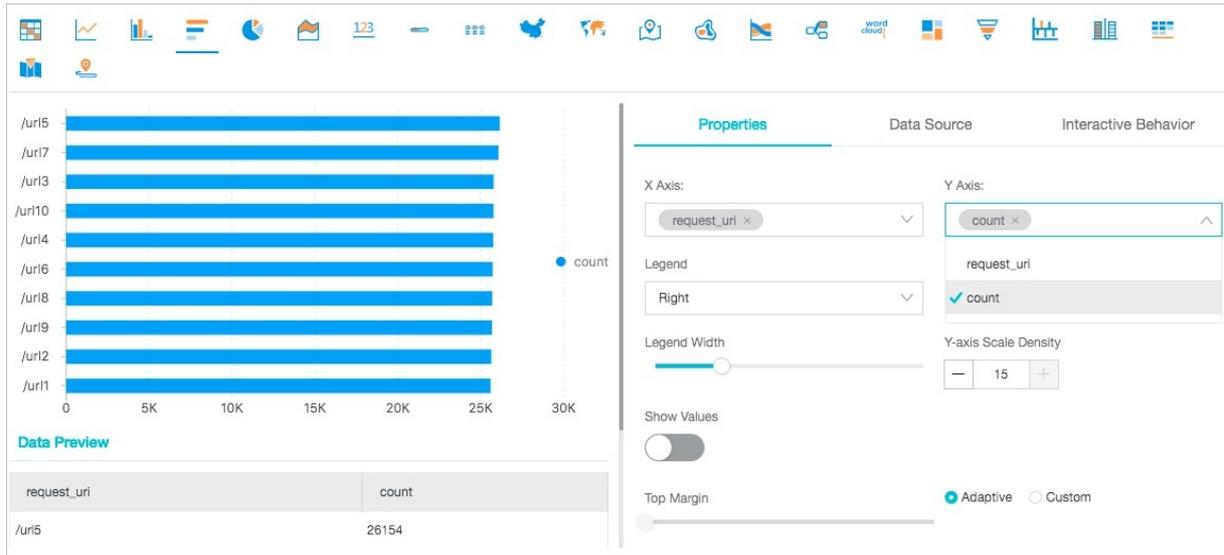
Parameters

| Parameter | Description |
|---------------|--|
| X Axis | The categorical data. |
| Y Axis | The numeric data. You can select one or more fields for the Y-axis. |
| Legend | The position where the legend is located in the chart. Valid values: Top, Bottom, Left, and Right. |
| Format X-axis | The format in which data is displayed on the X-axis. |
| Margin | The distance between an axis and the borders of the chart. Valid values: Top Margin , Bottom Margin , Right Margin , and Left Margin . |

Examples

To analyze the top 10 visited request URIs (`request_uri`), execute the following query statement :

```
* | select request_uri, count(1) as count group by request_uri order by count desc limit 10
```



23.4.11.1.6. Display query results on a pie chart

A pie chart is used to indicate the percentages of different data types and compare different data types based on the arc length of each slice.

Components

- Segment
- Percentage in the text format
- Legend

Types

Log Service provides three types of pie charts: a standard pie chart, donut chart, and polar area chart.

- Standard pie chart

A standard pie chart is divided into multiple segments based on the percentages of various field values. The entire chart displays all field values. Each segment displays the percentage and the numeric value of a field. The sum of percentages from all segments is equal to 100%.

- Donut chart

A donut chart is a standard pie chart with a hollow center. A donut chart has the following benefits:

- In addition to the information that a standard pie chart displays, a donut chart displays the total number of occurrences of all field values.
- You can view the differences between the number of occurrences of the same value in two charts based on the ring length. This is more intuitive than comparing two standard pie charts.

- Polar area chart

A polar area chart is a column chart in the polar coordinate system. Each category of field values is represented by a segment with the same radian. The radius of a segment indicates the number of occurrences of a field value. Compared with a standard pie chart, a polar area chart has the following benefits:

- Standard pie charts are suitable to display query results if 10 or fewer log entries are returned. Polar area charts are suitable for displaying query results if the number of returned log entries ranges from 10 to 30.
- The area is the square of a radius. Therefore, the display of the polar area chart enlarges the differences among multiple types of data. This is best suited for a comparison of similar values.
- A circle can be used to display a periodic pattern. Therefore, you can use a polar area chart to analyze value change characteristics in specified periods, such as weeks and months.

Procedure

1. On the Search & Analysis page, enter a query statement in the search box, specify a time range, and then click **Search & Analyze**.
2. On the **Graph** tab, click the  icon.
3. On the **Properties** tab, configure the properties of the pie chart.

Note

- Donut charts and standard pie charts can be used to display query results if less than 10 log entries are returned. You can use a `LIMIT` clause to control the number of segments. Analysis results may not be clearly displayed if the chart contains a large number of segments of different colors.
- We recommend that you use a polar area chart or column chart if the number of returned log entries exceeds 10.

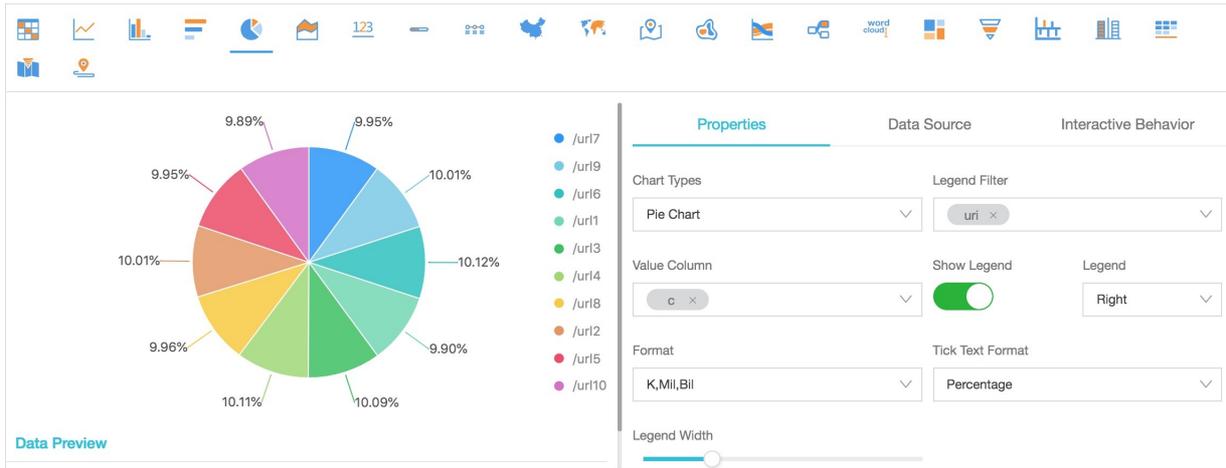
Properties

| Parameter | Description |
|------------------|--|
| Chart Types | The type of the chart. Valid values: Pie Chart, Donut Chart, and Polar Area Chart. Default value: Pie Chart. |
| Legend Filter | The categorical data. |
| Value Column | The values that correspond to different types of data. |
| Show Legend | Specifies whether to show the legend. |
| Legend | The position where the legend is located in the chart. Valid values: Top, Bottom, Left, and Right. You can configure this parameter only after you turn on the Show Legend switch. |
| Format | The format in which data is displayed. |
| Tick Text Format | The format of the tick. |
| Margin | The distance between an axis and the borders of the chart. Valid values: Top Margin , Bottom Margin , Right Margin , and Left Margin . |

Example of a standard pie chart

To analyze the percentages of the `request_uri` field values, execute the following query statement:

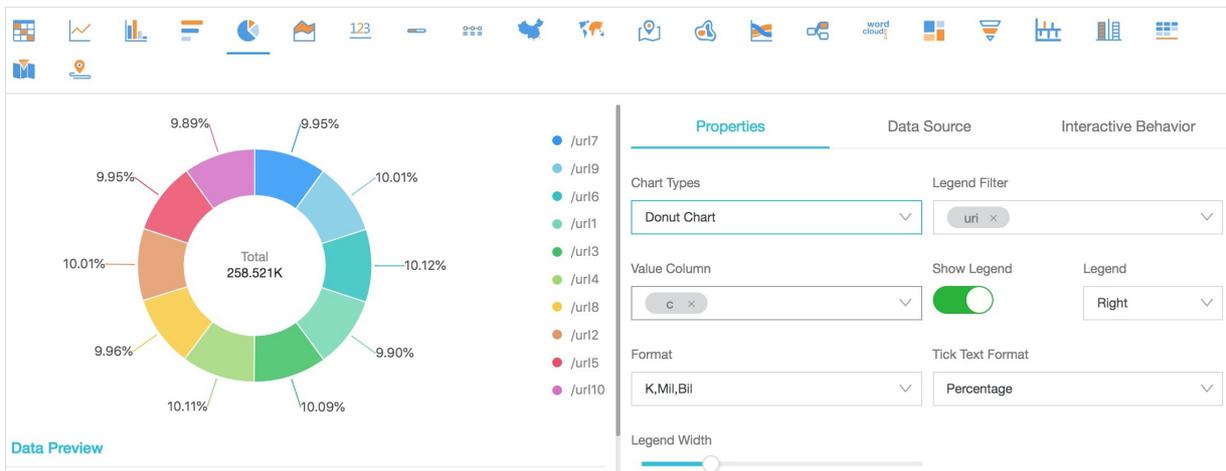
```
* | select requestURI as uri , count(1) as c group by uri limit 10
```



Example of a donut chart

To analyze the percentages of the `request_uri` field values, execute the following query statement:

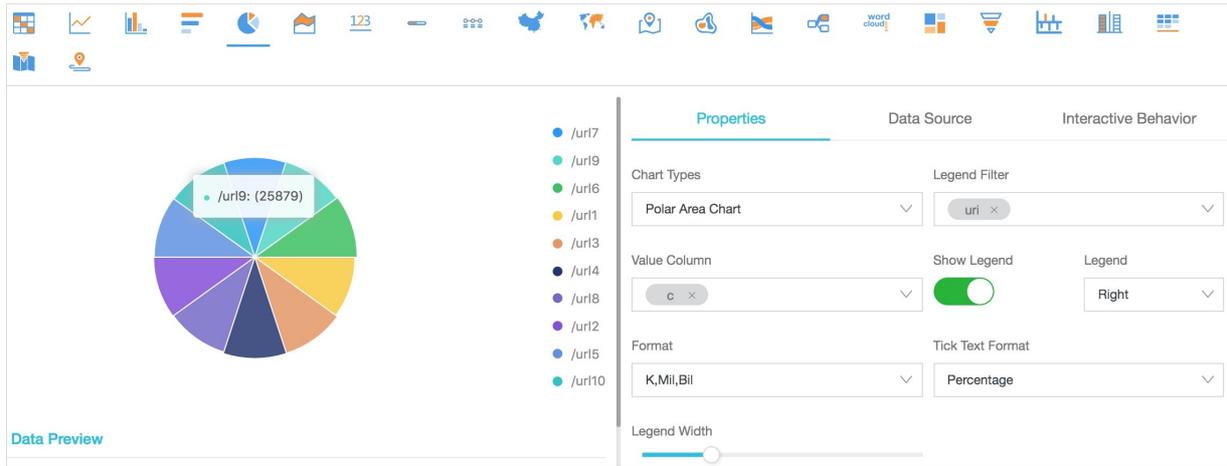
```
* | select requestURI as uri , count(1) as c group by uri limit 10
```



Example of a polar area chart

To analyze the percentages of the `request_uri` field values, execute the following query statement:

```
* | select requestURI as uri , count(1) as c group by uri limit 10
```



23.4.11.1.7. Display query results on an area chart

An area chart is built based on a line chart. The colored section between a line and the axis is an area. The color is used to highlight the trend. Similar to a line chart, an area chart shows the numeric value changes over a specified time period to highlight the overall data trend. Both the line chart and the area chart display the trend and relationship between numeric values instead of displaying specific values.

Components

- X-axis (horizontal)
- Y-axis (vertical)
- Area segment

Procedure

1. On the Search & Analysis page, enter a query statement in the search box, specify a time range, and then click **Search & Analyze**.
2. On the **Graph** tab, click the  icon.
3. On the **Properties** tab, configure the properties of the area chart.

 **Note** In an area chart, a single area segment must contain more than two data points. Otherwise, the data trend cannot be analyzed. We recommend that you select five or fewer area segments in an area chart.

Properties

| Parameter | Description |
|-----------|--|
| X Axis | The sequential data. In most cases, time series is selected. |
| Y Axis | The numeric data. You can select one or more fields for the Y-axis. |
| Legend | The position where the legend is located in the chart. Valid values: Top, Bottom, Left, and Right. |
| Format | The format in which data is displayed. |

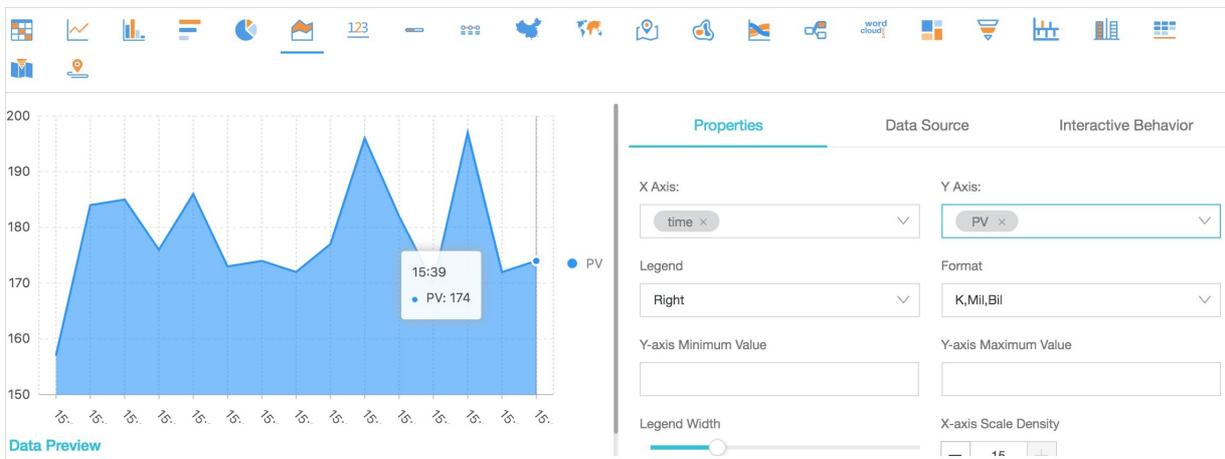
| Parameter | Description |
|-----------|--|
| Margin | The distance between an axis and the borders of the chart. Valid values: Top Margin , Bottom Margin , Right Margin , and Left Margin . |

Example of a simple area chart

To query the page views (PVs) of the IP address `10.0.192.0` in the last 24 hours, execute the following query statement:

```
remote_addr: 10.0.192.0 | select date_format(date_trunc('hour', __time__), '%m-%d %H:%i') as time, count(1) as PV group by time order by time limit 1000
```

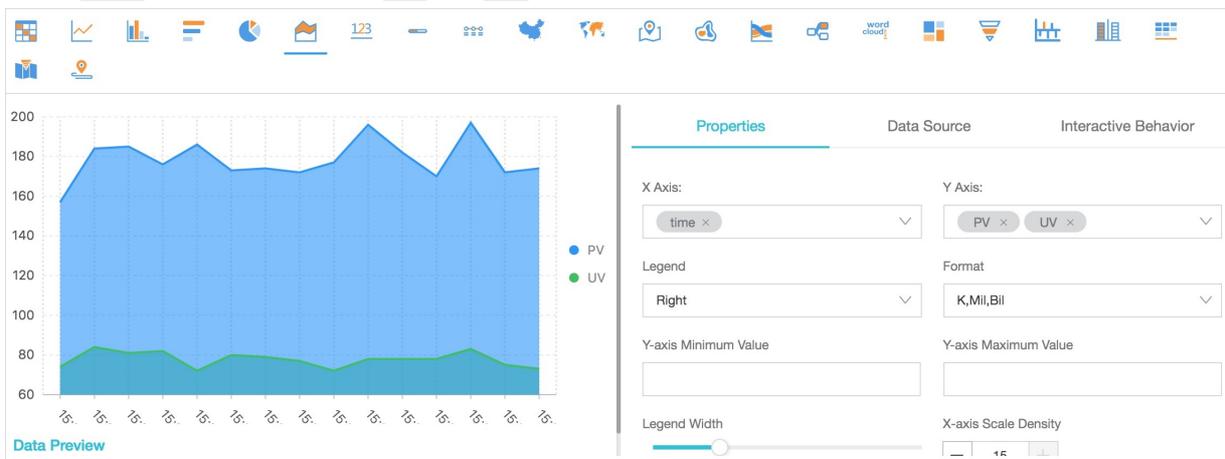
Select `time` for the X-axis and `PV` for the Y-axis.



Example of a cascade chart

```
* | select date_format(date_trunc('hour', __time__), '%m-%d %H:%i') as time, count(1) as PV, approx_distinct(remote_addr) as UV group by time order by time limit 1000
```

Select `time` for the X-axis. Select `PV` and `UV` for the Y-axis.



23.4.11.1.8. Display query results on a single value chart

A single value chart displays a single value.

Single value charts have the following types:

- **Rectangle Frame:** shows a general value.
- **Dial:** shows the difference between the current value and the specified threshold value.
- **Compare Numb Chart:** shows the SQL query results of interval-valued comparison and periodicity-valued comparison functions. For more information about the analytic syntax, see [Interval-valued comparison and periodicity-valued comparison functions](#).

Rectangle Frame is selected by default. A rectangle frame is the most basic method to display data at a specified point. In most cases, it is used to show the key information at a specified point in time. To display a proportional metric, you can select Dial.

Components

- Numeric value
- Chart type

Procedure

1. On the Search & Analysis page, enter a query statement in the search box, specify a time range, and then click **Search & Analyze**.
2. On the **Graph** tab, click the [123](#) icon.
3. On the **Properties** tab, configure the properties of the single value chart.

 **Note** Log Service normalizes data in numeric value-based charts. For example, 230000 is normalized to 230K. You can include [Mathematical calculation functions](#) in query statements to customize numeric formats.

Properties

- The following table lists the parameters of a rectangle frame.

| Parameter | Description |
|------------------------------|---|
| Chart Types | The type of the chart. Select Rectangle Frame. |
| Value Column | The value displayed in the chart. By default, data in the first row of the specified column is displayed. |
| Unit | The unit of data. |
| Unit Font Size | The font size of the unit. You can drag the slider to adjust the font size. Valid values: 10 to 100. Unit: pixels. |
| Description | The description of the value. |
| Description Font Size | The font size of the value description. You can drag the slider to adjust the font size. Valid values: 10 to 100. Unit: pixels. |
| Format | The format in which data is displayed. |
| Font Size | The font size of the value. You can drag the slider to adjust the font size. Valid values: 10 to 100. Unit: pixels. |
| Font Color | The color of the value, unit, and description in the chart. You can use the default color or select a color. |

| Parameter | Description |
|-------------------------|---|
| Background Color | The color of the background. You can use the default color or select another color. |

- The following table describes the parameters of a dial.

| Parameter | Description |
|------------------------------|--|
| Chart Types | The type of the chart. Select Dial to display query results on a dial. |
| Actual Value | The actual value in the chart. By default, data in the first row of the specified column is displayed. |
| Unit | The unit of the value on the dial. |
| Font Size | The font size of the value and unit. Valid values: 10 to 100. Unit: pixels. |
| Description | The description of the value. |
| Description Font Size | The font size of the value description. You can drag the slider to adjust the font size. Valid values: 10 to 100. Unit: pixels. |
| Dial Maximum | The maximum value of the scale in the dial. Default value: 100. |
| Maximum Value Column | The maximum value in the specified column. If you turn on the Use Query Results switch, Dial Maximum is replaced by Maximum Value Column. Then, you can select the maximum value from query results for this parameter. |
| Use Query Results | If you turn on the Use Query Results switch, Dial Maximum is replaced by Maximum Value Column. Then, you can select the maximum value from query results for this parameter. |
| Format | The format in which data is displayed. |
| Colored Regions | The number of segments that divide the dial. Each segment is displayed in a different color. Valid values: 2, 3, 4, and 5. Default value: 3. |
| Region Max Value | The maximum value of the scale in each colored segment of the dial. By default, the maximum value in the last segment is the maximum value on the dial. You do not need to specify this value. <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e0f0ff;"> <p>Note A dial is evenly divided into three colored segments by default. If you change the value of Colored Regions, Region Max Value is not automatically adjusted. You can manually set the maximum value for each colored segment based on your business requirements.</p> </div> |
| Font Color | The color of the value on the dial. |
| Region | The colored segments that divide the dial. A dial is evenly divided into three segments by default. The segments are displayed in blue, yellow, and red. If you set Colored Regions to a value greater than 3, the added segments are displayed in blue by default. You can change the color of each segment. |

| Parameter | Description |
|-------------------|--|
| Show Title | <p>Specifies whether to display the title of a dial when you add the dial to a dashboard. The Show Title feature can show or hide the title of the dial on the dashboard page. This switch is turned off by default.</p> <p>If you turn on the Show Title switch, the title of the dial is not displayed on the current page. You must create or modify a dashboard and view the title on the dashboard page.</p> |

- The following table lists the parameters of a compare numb chart.

| Parameter | Description |
|-----------------------------------|--|
| Chart Types | The type of the single value chart. If you select Compare Numb Chart, query results are displayed on a compare numb chart. |
| Show Value | The value that is displayed in the center of the compare numb chart. In most cases, this value is set to the statistical result that is calculated by the related comparison function in the specified time range. |
| Compare Value | The value that is compared with the threshold. Set the value to the result of the comparison between the statistical results calculated by the related comparison function in the specified time range and the previously specified time range. |
| Font Size | The font size of the show value. Valid values: 10 to 100. Unit: pixels. |
| Unit | The unit of the show value. |
| Unit Font Size | The font size of the unit for the show value. Valid values: 10 to 100. Unit: pixels. |
| Compare Unit | The unit of the compare value. |
| Compare Font Size | The font size of the compare value and unit. Valid values: 10 to 100. Unit: pixels. |
| Description | The description of the show value and its growth trend. |
| Description Font Size | The font size of the description. Valid values: 10 to 100. Unit: pixels. |
| Trend Comparison Threshold | <p>The value that is used to measure the variation trend of the compare value. For example, the compare value is -1.</p> <ul style="list-style-type: none"> ◦ If you set Trend Comparison Threshold to 0, a down arrow that indicates a value decrease is displayed on the page. ◦ If you set Trend Comparison Threshold to -1, it indicates that the value remains unchanged. The system does not display the trend on the page. ◦ If you set Trend Comparison Threshold to -2, an up arrow that indicates a value increase is displayed on the page. |
| Format | The format in which data is displayed. |
| Font Color | The color of the show value and its description. |
| Growth Font Color | The font color of the compare value that is greater than the threshold. |

| Parameter | Description |
|---------------------------|--|
| Growth Background Color | The background color displayed when the compare value is greater than the threshold. |
| Decrease Font Color | The font color displayed when the compare value is less than the threshold. |
| Decrease Background Color | The background color displayed when the compare value is less than the threshold. |
| Equal Background Color | The background color displayed when the compare value is equal to the threshold. |

Examples

To view the number of access requests, execute the following query statements:

- Rectangle frame

* | select count(1) as pv

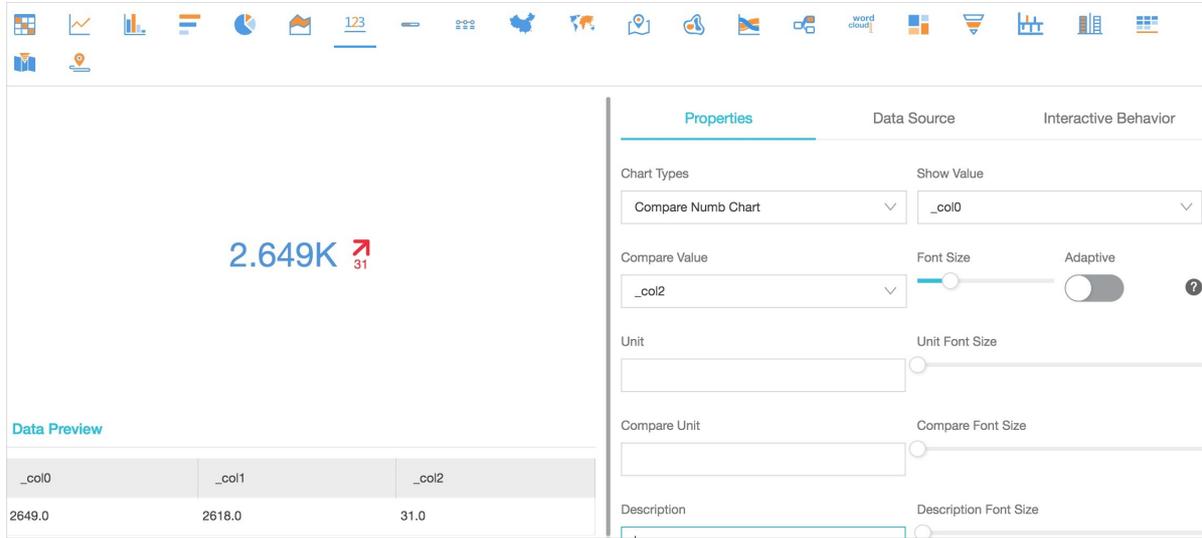
- Dial

* | select count(1) as pv

- Compare numb chart

To view and compare the access requests for today and yesterday, execute the following query statement :

```
* | select diff[1],diff[2], diff[1]-diff[2] from (select compare( pv , 86400) as diff from (select count(1) as pv from log))
```



23.4.11.1.9. Display query results on a progress bar

The progress bar shows the percentage of the actual value of a field to the maximum value of the field. You can configure the properties of the progress bar to adjust its style and set display rules.

Components

- Actual value
- Unit (optional)
- Total value

Procedure

1. On the Search & Analysis page, enter a query statement in the search box, specify a time range, and then click **Search & Analyze**.
2. On the **Graph** tab, click  to select the progress bar.
3. Configure the properties of the progress bar.

Properties

| Parameter | Description |
|--------------|--|
| Actual Value | The actual value in the chart. By default, data in the first row of the specified column is displayed. |
| Unit | The unit of the value in the progress bar. |
| Total Value | The maximum value indicated by the progress bar. Default value: 100. |

| Parameter | Description |
|----------------------|--|
| Maximum Value Column | The maximum value in the specified column. If you turn on the Use Query Results switch, Total Value is replaced by Maximum Value Column . Then, you can select the maximum value from the query results for this parameter. |
| Use Query Results | If you turn on the Use Query Results switch, Total Value is replaced by Maximum Value Column . Then, you can select the maximum value from the query results for this parameter. |
| Edge Shape | The edge shape of the progress bar. |
| Vertical Display | Specifies whether to display the progress bar in vertical display mode. |
| Font Size | The font size of the value in the progress bar. |
| Thickness | The thickness of the progress bar. |
| Background Color | The background color of the progress bar. |
| Font color | The font color of the value in the progress bar. |
| Default Color | The default color of the progress bar. |
| Color Display Mode | The display mode of the progress bar. |
| Start Color | The start color of the progress bar. This parameter is available if Gradient is selected for Color Display Mode . |
| End Color | The end color of the progress bar. This parameter is available if Gradient is selected for Color Display Mode . |
| Display Color | <p>The display color of the progress bar. This parameter is available if Display by Rule is selected for Color Display Mode.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note The value of Actual Value is compared with that of Threshold based on the condition specified by Operator. If the actual value matches the condition specified by Operator, the progress bar is displayed in the color specified by Display Color. Otherwise, the progress bar is displayed in the default color.</p> </div> |
| Operator | The condition that is used to specify the color of the progress bar. This parameter is available if Display by Rule is selected for Color Display Mode . |
| Threshold | The threshold that is used to specify the color of the progress bar. This parameter is available if Display by Rule is selected for Color Display Mode . |

Examples

To view the percentage of a metric or the proportion of data, you can execute the following query statement:

```
* | select diff[1],diff[2] from (select compare( pv , 86400) as diff from (select count(1) as pv from log))
```

The screenshot shows a dashboard with a progress bar indicating 26.3% completion. The progress bar is orange. To the right is a 'Properties' panel with the following settings:

- Actual Value: Unit:
- Total Value: Use Query Results:
- Edge Shape: Rounded Corner Right Angle Vertical Display:
- Font Size: Thickness:
- Background Color: Font Color:

Below the progress bar is a 'Data Preview' table:

| _col0 | _col1 |
|--------|--------|
| 2630.0 | 2630.0 |

If you select **Display by Rule** for Color Display Mode, colors are dynamically displayed based on the specified rule. If the conditions of a rule are not matched, the default color is displayed.

This screenshot is identical to the one above, but the progress bar is green, indicating that a rule-based color display mode is active. The 'Properties' panel settings remain the same.

23.4.11.1.10. Display query results on a map

You can color and mark a map to display geographic data. Log Service provides three types of maps: map of China, world map, and AMap. The display modes of an AMap include the anchor point and heat map. You can include specific functions in query statements to display analysis results as maps.

Components

- Map canvas
- Color area

Properties

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|----------------------|--|
| Location information | The location information that is recorded in logs. The information is displayed in one of the following dimensions based on the map type: <ul style="list-style-type: none"> Provinces (Map of China) Country (World Map) Longitude/Latitude (AMap) |
| Value Column | The data volume of the location information. |

Procedure

- On the Search & Analysis page, enter a query statement in the search box, specify a time range, and then click **Search & Analyze**.
 - To display query results on a map of China, include the `ip_to_province` function in a query statement.
 - To display query results on a world map, include the `ip_to_country` function in a query statement.
 - To display query results on an AMap, include the `ip_to_geo` function in a query statement.
- On the Graph tab, click the  icon.
- Configure the properties of the map.

Map of China

To display query results on a map of China, you can execute the following query statement that includes the `ip_to_province` function:

- SQL statement

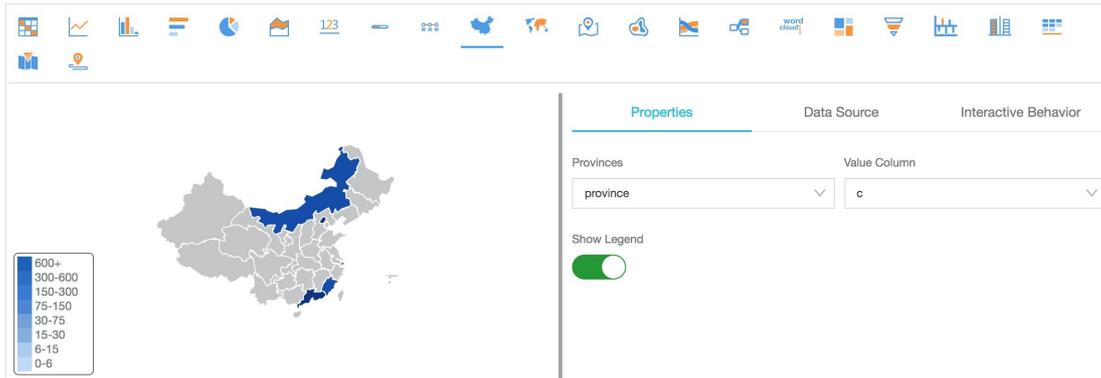
```
* | select ip_to_province(remote_addr) as address, count(1) as count group by address order by count desc limit 10
```

- Dataset

| address | count |
|--------------|-------|
| Guangdong | 163 |
| Zhejiang | 110 |
| Fujian | 107 |
| Beijing | 89 |
| Chongqing | 28 |
| Heilongjiang | 19 |

Select address for *Provinces* and count for *Value Column*.

Map of China



World Map

To display query results in a world map, you can execute the following query statement that includes the `ip_to_country` function:

- SQL statement

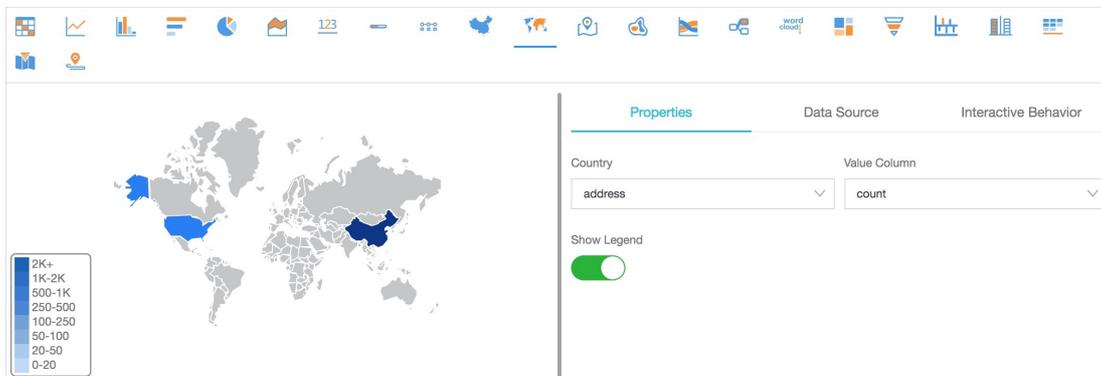
```
* |select ip_to_country(remote_addr) as address, count(1) as count group by address order by count desc limit 10
```

- Dataset

| address | count |
|---------------|-------|
| China | 8354 |
| United States | 142 |

Select address for *Country* and count for *Value Column*.

World Map



AMap

To display query results on an AMap, you can execute the following query statement that includes the `ip_to_geo` function. The address column in the dataset contains the latitude and longitude values, which are separated by a comma (.). If the longitude and latitude values are contained in two separate columns named lng and lat, you can use the `concat('lat', ',', lng')` function to combine the two columns into one column.

- SQL statement

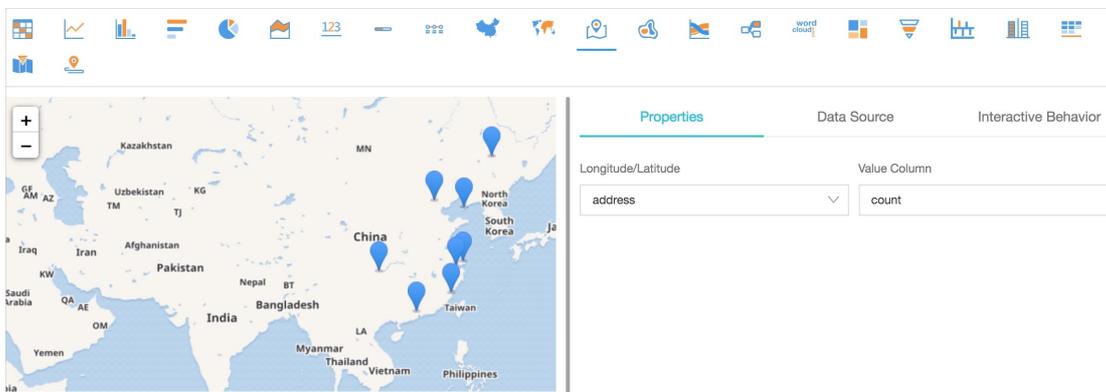
```
* |select ip_to_geo(remote_addr) as address, count(1) as count group by address order by count desc limit 10
```

- Dataset

| address | count |
|--------------------|-------|
| 39.9289,116.388 | 771 |
| 39.1422,117.177 | 724 |
| 29.5628,106.553 | 651 |
| 30.2936,120.161420 | 577 |
| 26.0614,119.306 | 545 |
| 34.2583,108.929 | 486 |

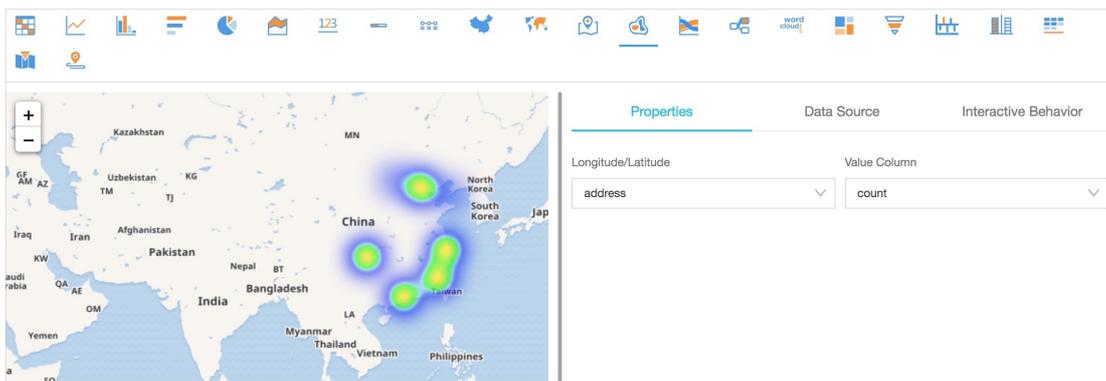
Select address for *Longitude/Latitude* and count for *Value Column*.

AMap: Anchor points



By default, the display mode of the anchor points is returned. If the anchor points are densely distributed on the map, you can switch the display mode to heat map.

AMap: Heat map



23.4.11.1.11. Flow chart

The flow chart, also known as ThemeRiver, is a stacked area chart around a central axis. The banded branches with different colors indicate different categorical data. The band width indicates the numeric value. The time information of the data is mapped to the X-axis by default. A flow chart can display the data of three parameters.

You can select the line chart or column chart for the Chart Types parameter. The column chart is stacked by default. Each category of data starts from the top of the last column of categorical data.

Components

- X-axis (horizontal)
- Y-axis (vertical)
- Band

Procedure

1. On the Search & Analysis page, enter a query statement in the search box, specify a time range, and then click **Search & Analyze**.
2. On the **Graph** tab, click the  icon.
3. On the **Properties** tab, configure the properties of the flow chart.

Properties

| Parameter | Description |
|------------------|--|
| Chart Types | The type of the chart. Valid values: Line Chart, Area Chart, and Column Chart. Default value: Line Chart. |
| X Axis | The sequential data. In most cases, time series is selected. |
| Y Axis | The numeric data. You can configure one or more fields on the Y-axis. |
| Aggregate Column | The field information required to be aggregated as the third point for comparison. |
| Legend | The position where the legend is located in the chart. Valid values: Top, Bottom, Left, and Right. |
| Format | The format in which data is displayed. |
| Margin | The distance between an axis and the borders of the chart. Valid values: Top Margin , Bottom Margin , Right Margin , and Left Margin . |

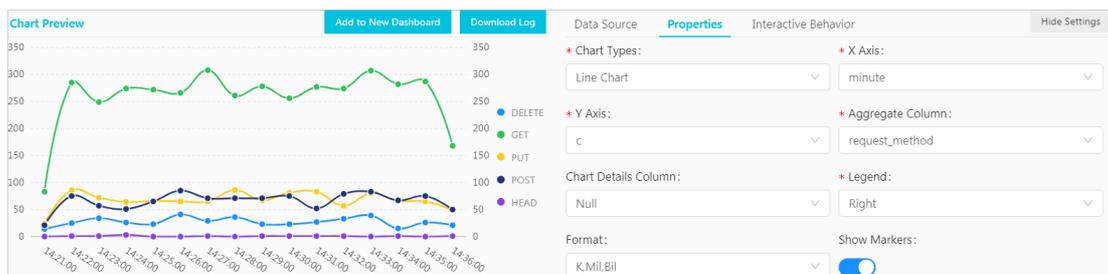
Examples

The flow chart is suitable to display data of three parameters, including the time, categories, and numeric values. In this example, you can execute the following query statement:

```
* | select date_format(from_unixtime(__time__ - __time__% 60), '%H:%i:%S') as minute, count(1) as c, request_method group by minute, request_method order by minute asc limit 100000
```

Select **minute** for the X-axis, **c** for the Y-axis, and **request_method** for Aggregate Column.

Flow chart



23.4.11.1.12. Display query results in a Sankey diagram

A Sankey diagram is a specific type of flow chart. It is used to describe the flow from one set of values to another set of values.

Sankey diagrams are applicable to scenarios such as network traffic flows. A Sankey diagram contains the values of three fields: `source`, `target`, and `value`. The `source` and `target` fields describe the source and target nodes and the `value` field describes the flows from the `source` node to the `target` node.

Features

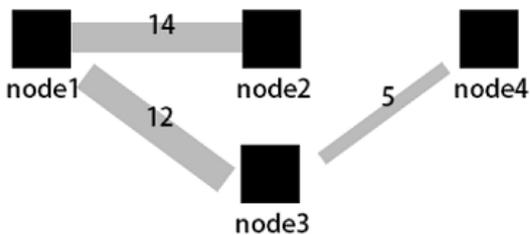
A Sankey diagram has the following features:

- The start flow is equal to the end flow. The overall width of all main edges is equal to the total sum of all branch edges. This allows you to manage and maintain a balanced flow of all traffic.
- The edge width in a row represents the traffic distribution in a specific status. The edge width in a row represents the distribution of traffic.
- The width of an edge between two nodes represents the flow volume of a status.

The following table lists the data that can be displayed in a Sankey diagram.

| source | target | value |
|--------|--------|-------|
| node1 | node2 | 14 |
| node1 | node3 | 12 |
| node3 | node4 | 5 |
| ... | ... | ... |

The following figure shows the data relationships in a Sankey diagram.



Components

- Node
- Edge

Procedure

1. On the Search & Analysis page, enter a query statement in the search box, specify a time range, and then click **Search & Analyze**.
2. On the Graph tab, click the  icon.
3. On the **Properties** tab, configure the properties of the Sankey diagram.

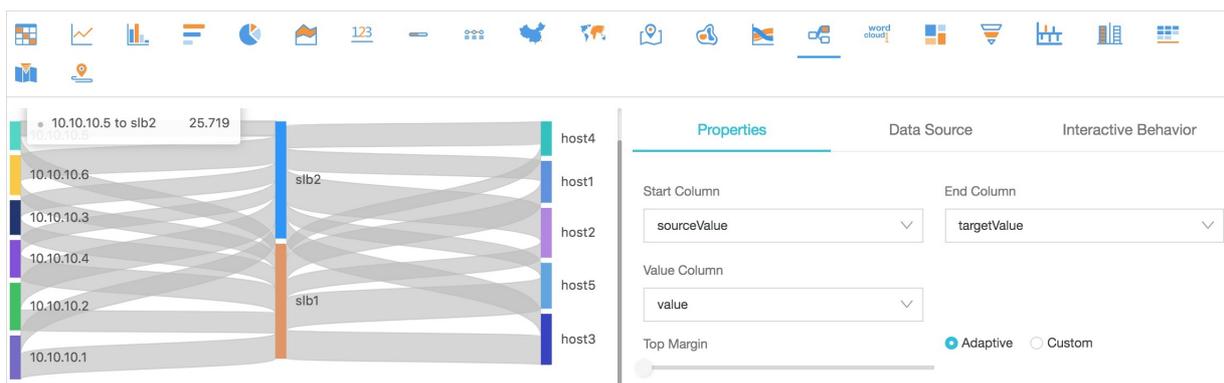
Properties

| Parameter | Description |
|--------------|--|
| Start Column | The start node. |
| End Column | The end node. |
| Value Column | The value that indicates the flow volume from the start node to the end node. |
| Margin | The distance between an axis and the borders of the chart. Valid values: Top Margin , Bottom Margin , Right Margin , and Left Margin . |

Example of a Sankey diagram

If a log entry contains the `source`, `target`, and `value` fields, you can use a **Nested subqueries** statement to obtain the sum of all `streamValue` values.

```
* | select sourceValue, targetValue, sum(streamValue) as streamValue from (select sourceValue, targetValue, streamValue, __time__ from log group by sourceValue, targetValue, streamValue, __time__ order by __time__ desc) group by sourceValue, targetValue
```



23.4.11.13. Display query results on a word cloud

A word cloud visualizes text data. It is a cloud-like and colored image that consists of words. You can use a word cloud to display a large amount of text data. The font size or color of a word indicates the significance of the word. This allows you to identify the most significant words in an efficient way.

Components

The words in a word cloud are sorted.

Procedure

1. On the Search & Analysis page, enter a query statement in the search box, specify a time range, and then click **Search & Analyze**.
2. On the **Graph** tab, click the  icon.
3. On the **Properties** tab, configure the properties of the word cloud.

Properties

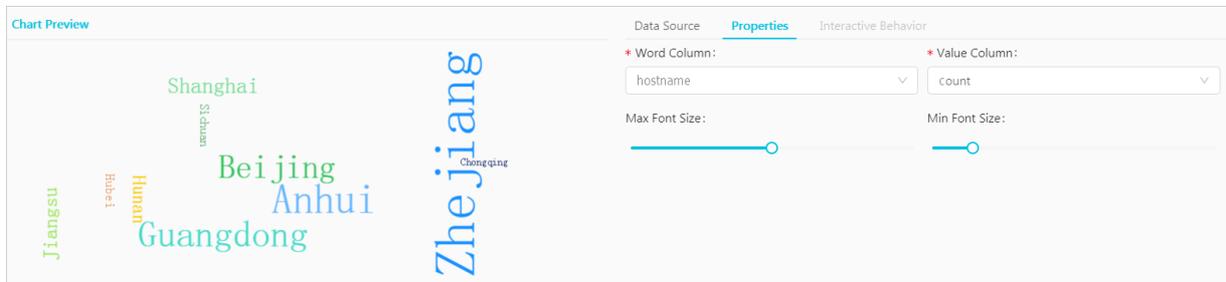
| Parameter | Description |
|--------------|---|
| Word Column | The words to be displayed. |
| Value Column | The numeric value that corresponds to a word. |
| Font Size | The font size of a word. <ul style="list-style-type: none"> The minimum font size ranges from 10 pixels to 24 pixels. The maximum font size ranges from 50 pixels to 80 pixels. |

Examples

To query the distribution of host names in NGINX logs, execute the following query statement :

```
* | select domain, count(1) as count group by domain order by count desc limit 1000
```

Select `hostname` for Word Column and `count` for Value Column.



23.4.11.1.14. Display query results on a treemap chart

A treemap chart includes multiple rectangles that represent the data volume. A larger rectangle area represents a larger proportion of the categorical data.

Components

Sorted rectangles

Procedure

1. On the Search & Analysis page, enter a query statement in the search box, specify a time range, and then click **Search & Analyze**.
2. On the Graph tab, click the  icon.
3. On the Properties tab, configure the properties of the treemap chart.

Properties

| Parameter | Description |
|---------------|---|
| Legend Filter | The field that includes categorical data. |
| Value Column | The numeric value of a field. A greater field value represents a larger rectangle area. |

Examples

To query the distribution of host names in NGINX logs, execute the following query statement :

```
* | select hostname, count(1) as count group by hostname order by count desc limit 1000
```

Select host name for **Legend Filter** and count for **Value Column** .

23.4.11.2. Dashboard

23.4.11.2.1. Overview

A dashboard provided by Log Service is a platform where you can analyze data in real time. You can add multiple charts to a dashboard for data analysis. Each chart is a visualized search and analytic statement.

A dashboard allows you to view the charts of multiple search and analytic statements at one time. When you open or refresh the dashboard, the statements of the charts run automatically.

After you add a chart to a dashboard, you can configure **Drill-down analysis** for the chart. Then you can click the chart on the dashboard to further analyze data and obtain more fine-grained analysis results.

Limits

- You can create a maximum of 50 dashboards for a project.
- Each dashboard can contain a maximum of 50 analysis charts.

Features

A dashboard has two modes: display mode and edit mode.

- **Configure the display mode of a dashboard**

In the display mode, you can configure multiple display settings on the dashboard page.

- **Dashboard:** You can specify the time range, the automatic refresh interval, full screen, and the display mode of the title for the dashboard, configure alerts for all charts on the dashboard, and filter chart data based on the **Configure and use a filter on a dashboard of a Logstore**.
- **Chart:** You can view the analysis details of a specified chart, specify the time range and configure alerts for the chart, download logs and the chart, and check whether **drill-down analysis** is configured for the chart.

- **Edit mode**

In the edit mode, you can change the configurations of the dashboard and charts.

- **Dashboard:** You can use a dashboard as a canvas and add **Markdown chart**, custom charts, text, icons, and other chart elements to the dashboard. You can also add lines between chart elements that are self-adaptive to the positions of the charts. You can also add **Configure and use a filter on a dashboard of a Logstore**, which can be used to filter chart data in the display mode. In addition, you can configure display gridlines to help arrange chart elements such as icons in an orderly manner.
- **Chart:** You can also edit a chart on the dashboard. You can modify the statement, properties, and interactive behavior such as **drill-down analysis** of the chart.

23.4.11.2.2. Create and delete a dashboard

This topic describes how to create and delete a dashboard in a Logstore. In the Log Service console, you can run a query statement and visualize the query result in a chart. After you complete the configurations of the chart, you can add the chart to a dashboard. Each dashboard can display up to 50 charts, which support multiple formats and custom settings.

Prerequisites

The indexing feature of the Logstore is enabled and configured. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

Create a dashboard

1. [Log on to the Log Service console](#).
2. Click a project name.
3. Click the  icon next to the name of the Logstore, and then select **Search & Analysis**.
4. Enter a query statement in the search box, and then click **Search & Analyze**.
5. On the **Graph** tab that appears, configure the chart properties on the **Properties** tab.
6. (Optional)Set a placeholder variable.

For example, you have configured a drill-down event for another chart. This drill-down event redirects you to the current dashboard. You have also specified a placeholder variable for the query statement of the preceding chart. When you click a chart value to trigger the drill-down event, you are redirected to the current dashboard. The placeholder variable is replaced by the chart value and the current dashboard is refreshed by the new query statement. For more information, see [Drill-down analysis](#).

- i. Click the **Data Source** tab, and then select a part of the query statement in the **Query** field.
- ii. Click **Generate Variable** to generate a placeholder variable.

iii. Set the parameters in the **Variable Config** section.

| Parameter | Description |
|----------------------|---|
| Variable Name | The name of the placeholder variable. If the name of the placeholder variable is the same as the variable specified in the chart, the placeholder variable is replaced with the chart value when the drill-down event is triggered. |
| Default Value | The default value of the placeholder variable in the current dashboard. |
| Matching mode | You can select Global Match or Exact Match. |
| Result | The query statement that contains the specified variable. |

Data Source
Properties
Interactive Behavior
Hide Settings

Query: Generate Variable

```
* | SELECT date_format(__time__ - __time__ % 60, '%H:%i:%s') as time, count(1) as count GROUP BY time ORDER BY time
```

Select the query statement to generate a placeholder variable. You can configure a drill-down configuration to replace the variable. For how to use dashboards, please refer to the documentation ([Help](#))

Variable Config:

* Variable Name:

* Default Value:

* Matching Mode:

Global Match
v
✕

Result

```
* | SELECT date_format(__time__ - __time__ % ${interval}, '%H:%i:%s') as time, count(1) as count GROUP BY time ORDER BY time
```

💬
🔍

7. Configure a drill-down event.

After you configure a drill-down event, you can click the chart on the dashboard for a more detailed analysis. For example, you can be redirected to another dashboard or a saved search. For more information, see [Drill-down analysis](#).

- i. Click the **Interactive Behavior** tab.
- ii. Select an **Event Action**.

iii. Set the parameters of the selected event action.

8. Click **Add to New Dashboard** and specify the dashboard name and chart name.

| Parameter | Description |
|-----------------------|---|
| Operation | <ul style="list-style-type: none"> ◦ Add to Existing Dashboard: Add the chart to an existing dashboard. ◦ Create Dashboard: Create a dashboard and then add the chart to the dashboard. |
| Dashboards | Select an existing dashboard name. <div style="background-color: #e0f2f1; padding: 5px; margin-top: 5px;"> ? Note This parameter is required only when you set the Operation parameter to Add to Existing Dashboard. </div> |
| Dashboard Name | Enter a dashboard name. <div style="background-color: #e0f2f1; padding: 5px; margin-top: 5px;"> ? Note This parameter is required only when you set the Operation parameter to Create Dashboard. </div> |
| Chart Name | Enter a name for the current chart. The chart name is displayed as the chart title in the dashboard. |

9. Click **OK**.

You can add multiple charts to a dashboard.

Delete a dashboard

You can delete a dashboard if you no longer need it. However, you cannot recover a deleted dashboard.

1. Log on to the Log Service console, and then click the destination project name.

- In the left-side navigation pane, click the **Dashboard** icon.
- Click the  icon next to the dashboard that you want to delete, and then select **Delete**.

23.4.11.2.3. Configure the display mode of a dashboard

This topic describes how to configure the display mode of a dashboard. By default, you can view all charts in a dashboard in the display mode. When you configure the display mode, you can add chart elements, enable automatic refresh, and set the title display mode.

Set a query time range for a dashboard of a Logstore

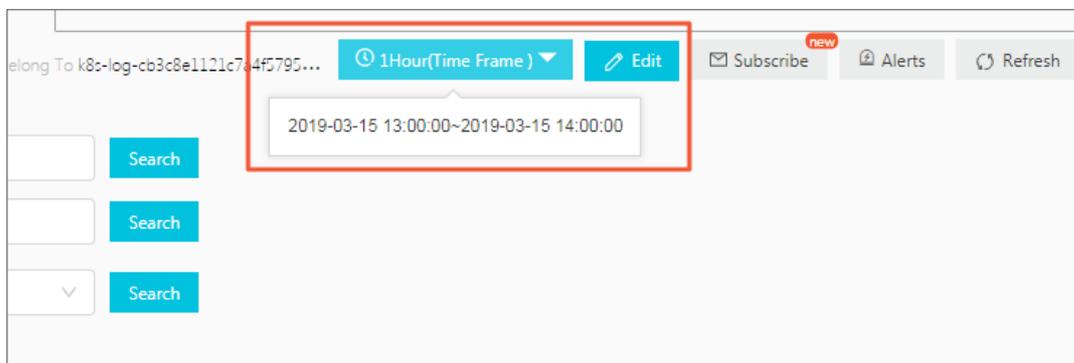
By default, all charts in a dashboard use the query time range that is set for the dashboard. For more information about how to set a query time range for a single chart, see [Set a query time range for a chart](#).

Note On the dashboard page, you can click **Time Range** to specify a time range for a query. The specified query time range is used only for the current query. The next time you open the dashboard, the system will display the analysis results in the default query time range.

- [Log on to the Log Service console](#).
- Click a project name.
- In the left-side navigation pane, click the **Dashboard** icon.
- Click the  icon next to the dashboard, and then select **Details** from the drop-down list.
- Click **Time Range** to set a time range.

You can set one of the following time ranges:

- Relative: queries log data obtained in a time range of 1 minute, 5 minutes, 15 minutes, or other time ranges that end with the current time, accurate to the second. For example, if the current time is 19:20:31 and you select 1Hour as the relative time, the charts on the dashboard display the analysis results of the log data queried from 18:20:31 to 19:20:31.
 - Time Frame: If you select or customize a time range less than one hour (for example, 1 minute, 5 minutes, and 15 minutes), log data obtained in the time range that ends with the current time is queried, accurate to the minute. If you select or customize a time range greater than one hour, log data obtained on the hour before the current time is queried. For example, if the current time is 19:20:31 and you select 1Hour as the time frame, the charts on the dashboard display the analysis results of the log data queried from 18:00:00 to 19:00:00.
 - Custom: queries log data obtained in a specified time range.
- Move the pointer over the **Time Range** button to confirm the specified time range.



Switch to the edit mode

Click **Edit** to switch to the edit mode of the dashboard. In the edit mode, you can add [Markdown chart](#), custom charts, text, icons, and other chart elements to the dashboard. For more information, see [Edit mode](#).

Set alerts

On the dashboard page, choose **Alerts > Create** to create an alert. Choose **Alerts > Modify** to modify an alert. An alert must be associated with one or more charts.

For more information, see [Configure alerts](#).

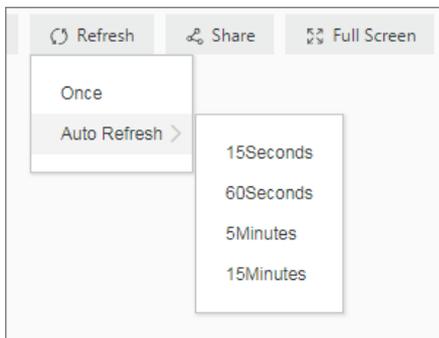
Set a refresh method

You can manually refresh the dashboard, or set an automatic refresh interval for the dashboard.

- To manually refresh the dashboard, choose **Refresh > Once**.
- To set an automatic refresh interval for the dashboard, choose **Refresh > Auto Refresh**, and then select an interval.

The **Auto Refresh** interval can be 15 seconds, 60 seconds, 5 minutes, or 15 minutes.

Note If your browser is inactive, the automatic refresh interval may be inaccurate.



Share a dashboard

To share a dashboard with authorized users, click **Share** to copy the link of the dashboard page and then send the link to the users. The shared dashboard page uses the settings of the dashboard at the time of sharing. The settings include the time range of charts and chart title format.

Note Before you share the dashboard with other users, you must grant relevant permissions to them.

Display a dashboard in full screen

Click **Full Screen**. Then the charts on a dashboard are displayed in full screen.

Set the chart title format

On the dashboard page, click **Title Configuration**. Available title formats include:

- Single-line Title and Time Display
- Title Only
- Time Only

Reset the query time range

To restore the default query time ranges of all charts on the dashboard, click **Reset Time**.

Select chart view

- View analysis details of a chart

To view analysis details of a chart, move the pointer over the More icon in the upper-right corner of the chart, and then select **View Analysis Details**. The corresponding Search & Analysis page appears, showing the query statement and property settings.

- Set the query time range for a chart

To set the query time range for a chart, move the pointer over the More icon in the upper-right corner of the chart, and then select **Select Time Range**. The settings are valid only for the current chart.

- Set an alert for a chart

To set an alert for a chart, move the pointer over the More icon in the upper-right corner of the chart, and then select **Create Alert**. For more information, see [Configure alerts](#).

- Download log analysis results of a chart

To download analysis results of a chart, move the pointer over the More icon in the upper-right corner of the chart, and then select **Download Log**.

- Download a chart

To download a chart, move the pointer over the More icon in the upper-right corner of the chart, and then select **Download Chart**.

- Check whether a drill-down event is configured for a chart

To check whether a drill-down event is configured for a chart, move the pointer over the More icon in the upper-right corner of the chart. Then, check the color of the hand icon at the bottom of the shortcut menu. If the icon is red, a drill-down event is configured for the chart. If the icon is gray, no drill-down event is configured for the chart.

 **Note** Different charts in a dashboard have different shortcut menus. For example, you cannot use the shortcut menu of a custom chart or Markdown chart to view analysis details because they are not analysis charts.

23.4.11.2.4. Edit mode

You can click the Edit button on the dashboard page to enter the edit mode. Then you can change the configurations of the dashboard and charts on the dashboard.

- Dashboard:
 - You can modify the dashboard name in the upper-left corner of the page.
 - You can use a dashboard as a canvas and add [Markdown chart](#), custom charts, text, icons, and other chart elements to the dashboard.
 - You can add lines between chart elements. The lines are adaptive to the positions of the charts.
 - You can add a filter to the dashboard. Then you can filter chart data in the display mode. For more information, see [Configure and use a filter on a dashboard of a Logstore](#).
 - In addition, you can configure display grid lines to arrange chart elements such as icons in order.
 - You can use the menu bar to manage the chart property settings on the dashboard. For example, you can perform the add, delete, and cancel operations on a chart. You can also configure the size and location of a chart and move a chart to the top or bottom of the dashboard.
- Chart: You can edit a chart on the dashboard. For example, you can modify the statement, properties, and interactive behavior such as [drill-down analysis](#) of a chart.

 **Note** All changes to the dashboard take effect only after you click **Save** in the upper-right corner of the page.

Chart elements

In the edit mode of a dashboard, you can add the following chart elements:

- Common icons

Log Service allows you to display common icons on a dashboard page. You can drag an icon from the menu bar and drop the icon to a specified position.

- Text

You can drag the text icon from the menu bar and drop the icon to a specified position. Double-click the text box to modify the text.

- Markdown chart

You can add a Markdown chart to a dashboard and edit the chart with the Markdown syntax.

Drag the Markdown icon from the menu bar and drop the icon to a specified position. Select **Edit** from the More icon. Then you can set the Markdown content.

- Filter

You can add a filter to a dashboard. Then you can add search conditions or replace placeholder variables in query statements.

Click the filter icon in the dashboard menu bar. On the page that appears, configure the filter based on your needs. By default, the filter is in the upper-left corner of the dashboard page. To modify the settings of a filter, you can select **Edit** on the More icon in the upper-right corner of the page.

- Customize SVG

You can upload a Scalable Vector Graphics (SVG) file to a dashboard. Click the SVG icon in the menu bar. On the page that appears, click the box or drag an SVG file to the box to upload the file.

 **Note** The size of an SVG file cannot exceed 10 KB.

- Customize image's HTTP link

You can upload the HTTP link of an image to a dashboard. Click the Customize image's HTTP link icon in the menu bar. On the page that appears, enter the HTTP link of an image and click **OK**.

Layout

In the edit mode, all charts and chart elements are placed on a canvas. You can drag and scale each chart, except for the lines. The horizontal width of a chart is up to 1,000 units. Each unit is equal to $\text{current browser width}/1,000$. The vertical height is unlimited and each unit is equal to 1 pixel. Before you arrange a chart on the dashboard, you can click **Display gridlines** to better arrange the position of the chart and the spacing between charts.

You can perform the following operations to arrange a chart on the dashboard:

- Adjust the position of a chart

- You can drag a chart and drop the chart to a specified position.
- After you specify a chart, you can click **L** and **T** to adjust the chart position.

- Adjust the width and height of a chart

- After you specify a chart, you can drag the chart in the lower-right corner to resize the chart.
- After you specify a chart, you can also specify the **W** and **H** parameters to resize the chart.

- Add lines to connect charts

You can add a directional line between two charts. When you adjust the position and size of the charts, the line automatically moves to display the relative position between the two charts.

After you select a chart, four squares appear on the border of the chart. These squares indicate the starting point of the connection line. Press and hold one square, and the area where the end point of the connection line is automatically displayed. Move the pointer to this position to connect the two charts.

- You can move a chart to the top or bottom of the dashboard. After you select a chart, you can use the menu bar to move the chart to the top or bottom.

Chart configurations

In the edit mode of a dashboard, you can perform the following operations on chart elements:

- **Edit**: modifies the query statement, properties, and interactive behavior such as [drill-down analysis](#) of a chart.
 - i. In the upper-right corner of the dashboard page, click **Edit**.
 - ii. Find the target chart and choose  > **Edit**.
 - iii. Modify the query statement of the chart, **Properties**, **Data Source**, or **Interactive Behavior**.
 - iv. Click **Preview**, and then click **OK**.
 - v. In the upper-right corner of the dashboard page, click **Save**.
- **Copy**: creates a copy of the specified chart and retains all configurations.
 - i. In the upper-right corner of the dashboard page, click **Edit**.
 - ii. Find the target chart and choose  > **Copy**.
 - iii. Drag the chart copy and drop it to a specified position, and then set the top margin, left margin, and size of the copy.
 - iv. In the upper-right corner of the dashboard page, click **Save**.
- **Delete**: deletes the specified chart from the dashboard.
 - i. In the upper-right corner of the dashboard page, click **Edit**.
 - ii. Find the target chart and choose  > **Delete**.
 - iii. In the upper-right corner of the dashboard page, click **Save**.

23.4.11.2.5. Drill-down analysis

This topic describes how to configure drill-down analysis for a chart of a Logstore. When you add a chart to a dashboard, you can modify the configurations in the drill-down list to obtain deeper data analysis results.

Prerequisites

- The index feature of the Logstore is enabled and configured. For more information, see [Enable the index feature and configure indexes for a Logstore](#).
- A saved search, a dashboard, and a custom link to which you want to be redirected are configured.
- A placeholder variable is specified in the saved search or the query statement of a chart added to a dashboard to which you are redirected if you want to add a variable. For more information, see [Saved search](#) and [Create and delete a dashboard](#).

Context

Drilling is essential for data analysis. This feature allows you to analyze data in a fine-grained or coarse-grained way. This feature includes drill-up and drill-down analysis. You can implement drill-down analysis to gain a deeper insight into data and make a better decision.

Log Service supports drill-down analysis of the following charts: tables, line charts, column charts, bar charts, pie charts, individual value plots, area charts, and treemap charts.

Procedure

1. [Log on to the Log Service console](#).
2. Click the target project name.

3. Click the  icon next to the name of the Logstore, and then select **Search & Analysis** from the drop-down list.
4. Enter a query statement in the search box, set a time range, and then click **Search & Analyze**.
5. On the **Graph** tab that appears, select a **chart type**, and then configure the parameters on the **Properties** tab of the chart.
6. Click the **Interactive Behavior** tab. On this tab, configure the **Event Action** for drill-down analysis.
 - o **Disable**: disables drill-down analysis.
 - o **Open Logstore**: sets the drill-down event to open a Logstore.

If you configure a filter statement on the **Interactive Behavior** tab, the filter statement automatically fills the **Search & Analyze** search box of the redirected Logstore page when you click a value on the chart.

| Parameter | Description |
|-------------------------------------|--|
| Select Logstore | The name of the Logstore to which you want to be redirected. |
| Open in New Tab | If you turn on this switch, the specified Logstore is opened on a new tab when the interactive behavior is triggered. |
| Time Range | <p>The query time range of the Logstore to which you are redirected. Valid values:</p> <ul style="list-style-type: none"> ■ Default: The default time range (15 minutes, accurate to the second) is used for a query statement of the redirected Logstore. ■ Inherit table time: The time range that a statement queries in the redirected Logstore is the time range specified for the chart when the event is triggered. ■ Relative: The time range that a statement queries in the redirected Logstore is accurate to the second. ■ Time Frame: The time range that a statement queries in the redirected Logstore is accurate to the minute, hour, or day. <p>Default value: Default.</p> |
| Inherit Filtering Conditions | If you turn on the Inherit Filtering Conditions switch, the filtering conditions added to the dashboard are synchronized to a query statement of the specified Logstore. The filtering conditions are added before the query statement by using the logical AND operator. |
| Filter | <p>Enter a Filter Statement on the Filter tab. The filter statement can contain the Optional Parameter Fields.</p> <p>If you configure a filter statement on the Filter tab, the filter statement automatically fills the Search & Analyze search box of the redirected Logstore page when you click a chart value on the dashboard.</p> |

- o **Open Saved Search**: sets the drill-down event to execute a saved search.

You can configure variables and a filter statement for a chart at the same time. When you click a value on the chart:

- If you configure a variable for the chart, the variable value that you click replaces the placeholder variable that you have configured for the saved search and the saved search is executed for drill-down analysis.
- If you configure a filter statement, the filter statement is added to the saved search to which you are redirected.

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|------------------------------|---|
| Select Saved Search | The name of the saved search to which you want to be redirected. For more information about how to configure a saved search, see Saved search . |
| Open in New Tab | If you turn on this switch, the specified saved search is opened on a new tab when the interactive behavior is triggered. |
| Time Range | <p>The time range of the saved search to which you want to be redirected. Valid values:</p> <ul style="list-style-type: none"> ▪ Default: The default time range (15 minutes, accurate to the second) is used for the saved search to which you are redirected. ▪ Inherit table time: The time range of the saved search is the query time range of the chart that you configure on the dashboard. ▪ Relative: The time range of the saved search is accurate to the second. ▪ Time Frame: The time range of the saved search is accurate to the minute, hour, or day. <p>Default value: Default.</p> |
| Inherit Filtering Conditions | If you turn on the Inherit Filtering Conditions switch, the filtering conditions added on the dashboard are synchronized to the saved search of the specified Logstore. The filtering conditions are added before the saved search by using the logical AND operator. |
| Filter | <p>Click the Filter tab, and then enter a Filter Statement. The filter statement can contain the Optional Parameter Fields.</p> <p>If you configure a filter statement on the Filter tab, the Filter Statement is added to the saved search when you click a chart value on the dashboard.</p> |
| Variable | <p>Click the Variable tab, click Add Variable, and then set the following parameters:</p> <ul style="list-style-type: none"> ▪ Replace Variable Name: the variable that triggers the drill-down event. When you click this variable, you are redirected to the specified saved search. ▪ Replace the value in the column: the column where the value that you want to replace data with is located. To process multiple columns, you can specify the current column and other columns. The current column is the column on which you want to perform drill-down analysis. Specify the current column in the Replace the value in the column field. Other columns can be the fields in the chart for which you configure drill-down analysis. <p>If the variable name of the saved search is the same as the added variable, the variable of the saved search is replaced with the chart value that triggers the drill-down event. This helps you use the saved search for analysis in an efficient way.</p> <div style="background-color: #e6f2ff; padding: 10px; border: 1px solid #d9e1f2;"> <p> Note</p> <ul style="list-style-type: none"> ▪ If you add a variable for drill-down analysis, you must first configure a placeholder variable for the saved search to which you want to be redirected. ▪ You can add up to five placeholder variables for a saved search. </div> |

- **Open Dashboard:** sets the drill-down event to open a dashboard.

Analysis charts on a dashboard are visualized query results. When you click a chart value on the source dashboard:

- If you configure a variable for the chart and set a placeholder variable for the chart on the destination dashboard to which you want to be redirected, the placeholder variable is replaced with the chart value that you click.
- If you configure a filter statement, the filter statement is added to the query statement of the chart on the destination dashboard.

| Parameter | Description |
|------------------------------|--|
| Select Dashboard | The name of the dashboard to which you want to be redirected. For more information, see Create and delete a dashboard . |
| Open in New Tab | If you turn on this switch, the specified dashboard is opened on a new tab when interactive behavior is triggered. |
| Time Range | <p>Set the time range for the dashboard to which you want to be redirected. Valid values:</p> <ul style="list-style-type: none"> ■ Default: After you are redirected to the dashboard by clicking a chart value on the current dashboard, the selected dashboard uses its original time range. ■ Inherit table time: After you are redirected to the selected dashboard, the time range of the chart on the selected dashboard is the time range of the chart specified on the source dashboard when the event is triggered. ■ Relative: After you are redirected to the selected dashboard, set the time range of the selected dashboard to the specified relative time. ■ Time Frame: After you are redirected to the selected dashboard set the time range of the selected dashboard to the specified time frame. <p>Default value: Default.</p> |
| Inherit Filtering Conditions | If you turn on the Inherit Filtering Conditions switch, the filtering conditions added on the current dashboard are synchronized to the dashboard to which you are redirected. The filtering conditions are added before the query statement by using the logical AND operator. |
| Filter | <p>Click the Filter tab, and then enter a Filter Statement. The filter statement can contain the Optional Parameter Fields.</p> <p>If you have set the Filter, a filtering condition is added to the selected dashboard when you click a chart value on the current dashboard. The filtering condition is the specified Filter Statement.</p> |

| Parameter | Description |
|-----------|--|
| Variable | <p>Click the Variable tab, click Add Variable, and then set the following parameters:</p> <ul style="list-style-type: none"> ▪ Replace Variable Name: the variable that triggers drill-down analysis. When you click this variable, you are redirected to the selected dashboard. ▪ Replace the value in the column: the column where the value that you want to replace data with is located. To process multiple columns, you can specify the Default Column and other columns. The Default Column is the current column on which you want to perform drill-down analysis. Other columns can be fields in the chart for which you configure drill-down analysis. <p>If the variable in the query statement of the chart on the selected dashboard is the same as the added variable, the variable is replaced with the chart value that triggers the drill-down event. This changes the query statement of the chart on the selected dashboard.</p> <div style="background-color: #e0f2f7; padding: 10px; border: 1px solid #ccc;"> <p> Note</p> <ul style="list-style-type: none"> ▪ If you add a variable for drill-down analysis, you must first configure a placeholder in the query statement for the selected dashboard to which you want to be redirected. ▪ You can add up to five variables. </div> |

- **Custom HTTP Link:** sets the drill-down event to open a custom HTTP link.

The path in the HTTP link is the path of the destination file in the server. When you add optional parameter fields to the path and click the chart value on the dashboard, the added parameter fields are replaced with the chart value. At the same time, you are redirected to the new HTTP link.

| Parameter | Description |
|---------------------------|--|
| Enter Link | The destination address to which you want to be redirected. |
| Optional Parameter Fields | By clicking an optional parameter variable, you can replace a part of the HTTP link with the chart value that triggers a drill-down event. |

7. Click **Add to New Dashboard** and set the dashboard name and chart name.

Example

You want to store collected NGINX access logs in a Logstore named accesslog, display the relevant analysis results on a dashboard named RequestMethod, and display the page view (PV) changes over time on a dashboard named destination_drilldown. You can configure drill-down analysis for the table of request methods, add the table to the RequestMethod dashboard, and then configure the drill-down event that redirects you to the destination_drilldown dashboard. When you click a request method on the table on the RequestMethod dashboard, you are redirected to the destination_drilldown dashboard. Then you can view the PV changes on the dashboard.

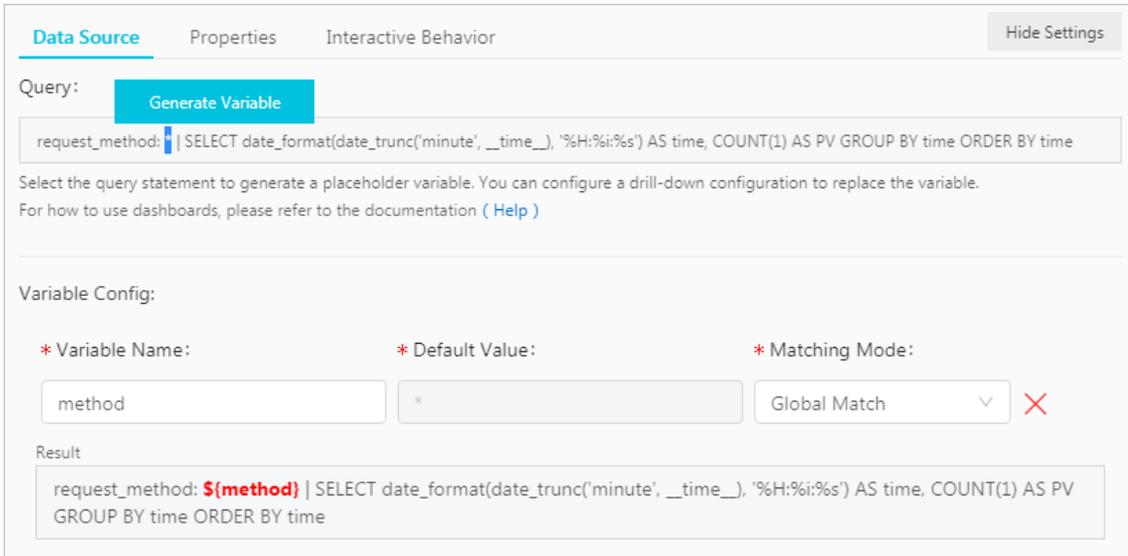
1. Set the dashboard to which you want to be redirected.
 - i. Filter log data by request method and analyze how the PV of each request method changes over time.

The query statement is shown as follows:

```
request_method: * | SELECT date_format(date_trunc('minute', __time__), '%H:%i:%s') AS time, COUNT(1) AS PV GROUP BY time ORDER BY time
```

- ii. Use a line chart to display the query result and add the line chart to the dashboard named destination_drilldown.

Before you add the line chart to the dashboard, specify the asterisk (*) to generate a placeholder variable and set the variable name to method. If the variable name of another chart for which you configure drill-down analysis is method, when you click the variable value on the chart to trigger drill-down analysis, the asterisk (*) is replaced with the value that you click and the query statement of the line chart is performed.



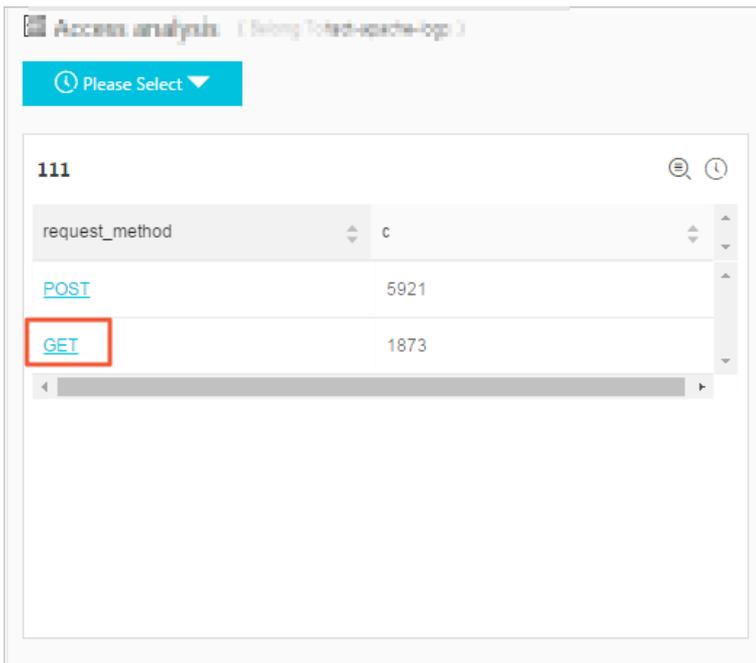
- 2. Configure drill-down analysis for a chart and add the chart to the dashboard.

- i. On the Search & Analysis page, enter a SQL statement to query the number of NGINX access log entries of each request method, and display the result in a table.

The query statement is shown as follows:

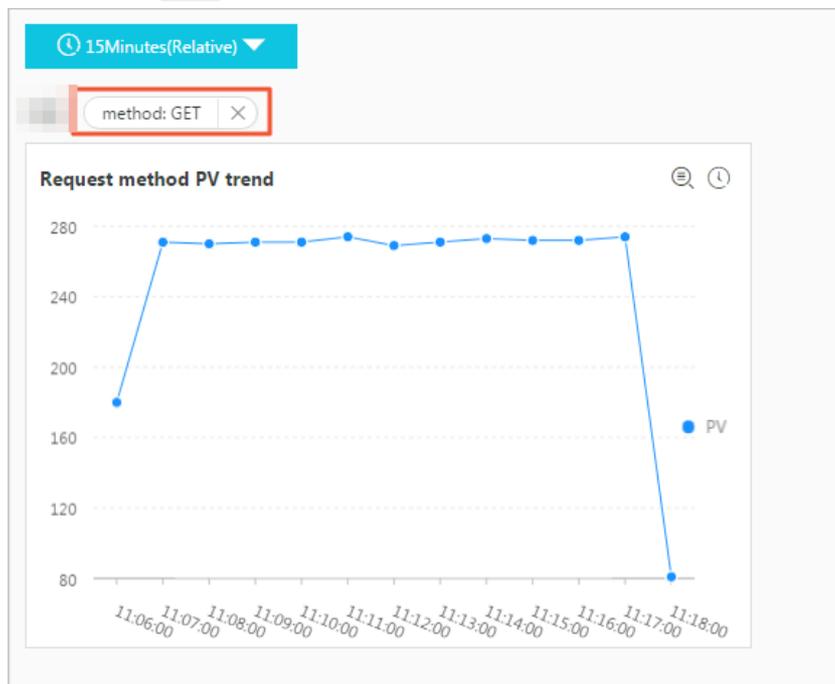
```
*|SELECT request_method, COUNT(1) AS c GROUP BY request_method ORDER BY c DESC LIMIT 10
```

- ii. Configure drill-down analysis for the request_method column in the table.
- iii. Click the GET request on the RequestMethod dashboard.



iv. Redirected to the destination_drilldown dashboard.

You are redirected to the dashboard configured in step . The asterisk (*) in the query statement is replaced with GET . The dashboard shows how the PV of the GET request changes over time.



23.4.11.2.6. Configure and use a filter on a dashboard of a Logstore

This topic describes how to configure and use a filter on a dashboard of a Logstore. Filters help you refine search results or replace placeholder variables with specified values.

Prerequisites

- The index feature of the Logstore is enabled and configured. For more information, see [Enable the index feature and configure indexes for a Logstore](#).
- A dashboard is created. For more information, see [Create and delete a dashboard](#). A placeholder variable is specified for charts on the dashboard if the filter type is set to Replace Variable.

Context

Each chart on a dashboard is a visualized query statement. When you configure a filter on a dashboard, the specified filtering condition or variables apply to all charts on the dashboard. You can configure the following two types of filters:

- Filter: Add key-value pairs as a filtering condition before the query statement [search query] . The new query statement is key:value AND [search query] , which means to search the result of the original query statement for log entries that contain key:value . For the Filter type, you can select or enter multiple key-value pairs. When you select multiple key-value pairs as filtering conditions, the logical OR operator is used between the pairs.
- Replace Variable: Specify a variable. If the dashboard contains a chart configured with the same placeholder variable, the placeholder variable in the query statement of the chart is replaced with the selected value.

Components

Each filter chart has one or more filters. Each filter consists of the following two components:

- The key, which indicates a filter operation.

- The values of the key.

Procedure

1. Log on to the Log Service console.
2. Click a project name.
3. In the left-side navigation pane, click the **Dashboard** icon.
4. In the dashboard list, click the name of the target dashboard.
5. On the dashboard page, click **Edit** to enter the edit mode.
6. Click the  icon, and then set the filter parameters. Click **OK**.

Parameters of a filter

| Parameter | Description |
|--------------------------|---|
| Filter Name | The name of the filter. |
| Display Settings | <p>Valid values:</p> <ul style="list-style-type: none"> ◦ Title: specifies to add a title for a filter. You can turn on the Title switch to add a title for a filter. ◦ Border: specifies to add borders for a filter. You can turn on the Border switch to add borders for a filter. ◦ Background: specifies to add a white background for a filter. You can turn on the Background switch to add a white background for a filter. |
| Type | <p>The filter type. Valid values:</p> <ul style="list-style-type: none"> ◦ If you select Filter, a List Item is a value of a key that is used to filter the results of a query statement. You can set multiple values for a key. After the filter takes effect, you can select one or more values on the dashboard to filter query results based on your needs. ◦ If you select Replace Variable, a List Item is the value that replaces a specified variable. You can set multiple values for a variable. After the filter takes effect, you can select one or more values on the dashboard to filter query results based on your needs. |
| Key | <ul style="list-style-type: none"> ◦ For the Filter type, the Key parameter specifies the key in the filtering condition. ◦ For the Replace Variable type, the Key parameter specifies the variable. <div style="background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note The variable must be the same as the placeholder variable that you specified for charts. Otherwise, the placeholder variable cannot be replaced.</p> </div> |
| Alias | The alias of the key. This parameter is available only when you select Filter . After you set an alias for a key, the alias is displayed in the filter on the dashboard. |
| Global filter | <p>This switch indicates whether to filter the specified values in all fields. This switch is turned off by default. The switch is available only when you select the Filter type.</p> <ul style="list-style-type: none"> ◦ To filter the specified values in all fields, turn on the Global filter switch. ◦ To filter the specified values in specified keys, turn off the Global filter switch. |
| Static List Items | The values of the key that is used as a filtering condition. Click the plus sign and enter a value for the key in the text box. |

| Parameter | Description |
|------------------------------|--|
| Add Dynamic List Item | The dynamic value of the key retrieved by running the specified query statement. Turn on the Add Dynamic List Item switch, select a Logstore, and turn on the Inherit Filtering Conditions switch (specifies whether to include the filtering condition in the query statement). Enter a query statement in the search box, specify a time range, and then click Search to preview the dynamic values. |

Examples

You can use a filter to modify the query statements of charts on a dashboard and replace placeholder variables in the charts on the dashboard. Each chart represents a query statement in the format of `[search query] | [sql query]`. The filter query is appended to the original query statement to filter data.

- If the filter type is Filter, the key-value pairs to be filtered are added before `[search query]` to form a new query statement by using the logical `AND` operator. The new query statement is `key:value AND [search query]`.
- If the filter type is Replace Variable, the filter queries all charts that contain the specified placeholder variables on the dashboard and replaces the placeholder variables with the selected values.

Example 1: Use different time granularities to analyze logs

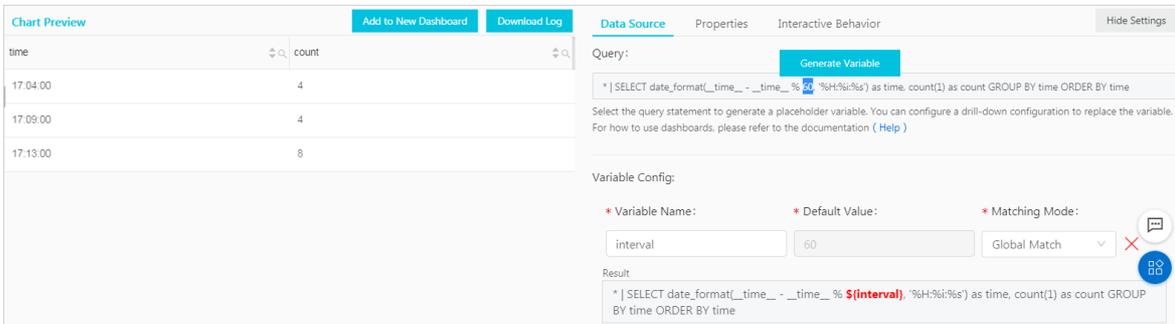
After you collect NGINX access logs, you need to query and analyze these logs in real time.

You can use a query statement to view the number of page views (PVs) per minute. However, if you want to view the number of PVs per second, you must modify the value of `__time__ - __time__ % 60` in the query statement. To simplify this operation, you can use a filter to replace variables in the query statement.

1. Use the following query statement to view the number of PVs per minute:

```
* | SELECT date_format(__time__ - __time__ % 60, '%H:%i:%s') as time, count(1) as count GROUP BY time ORDER BY time
```

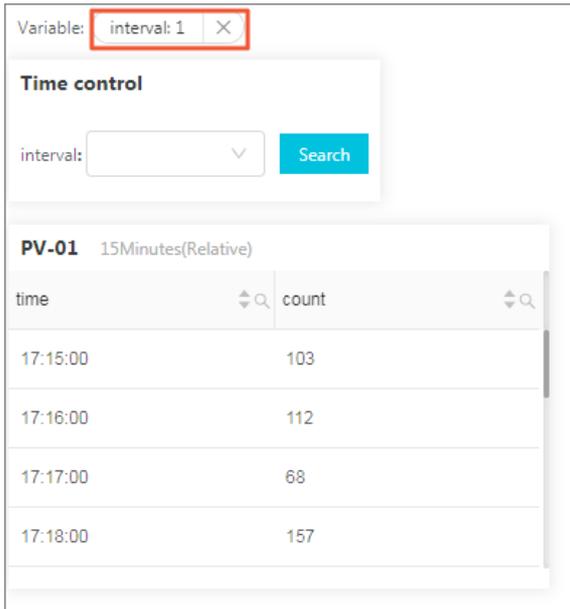
2. Add the chart to a dashboard, and select `60` generate a variable with the name `interval`.



3. Add a filter on the dashboard page.
 - o **Type:** Select **Replace Variable**.
 - o **Key:** Enter `interval`.
 - o **Static List Items:** Add `1` and `120` as values of the key. The default unit is seconds.
4. In the Filter section of the dashboard, select `1` from the Interval drop-down list to view statistics by second.

The following statement shows the query statement after the placeholder variable is replaced:

```
* | SELECT date_format(__time__ - __time__ % 1, '%H:%i:%s') as time, count(1) as count GROUP BY time ORDER BY time
```



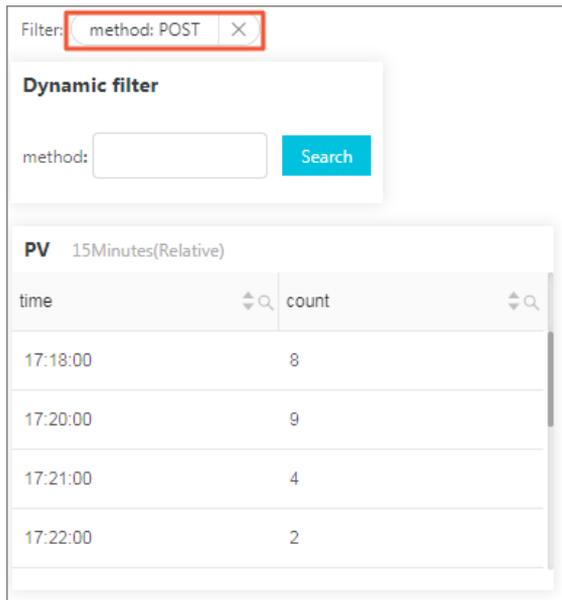
Example 2: Switch between request methods

You can add dynamic values to a filter to dynamically switch between request methods. In example 1, the query statement starts with an asterisk (*), which means no condition is set to filter the query results and all logs are queried. You can add another filter to view the PV data of different request methods .

1. Add a filter on the dashboard and turn on the **Add Dynamic List Item** switch.
Set the parameters as follows:
 - o **Type:** Select Filter.
 - o **Key:** Enter `method` .
 - o **Select Logstore:** Select the Logstore to which the dashboard belongs.
 - o **Add Dynamic List Item:** Enter a query statement to obtain the relevant dynamic values, and then click OK.
2. In the Filter section of the dashboard, select `POST` from the method drop-down list.

Only the PV data whose `method` is `POST` is displayed in the chart. The query statement is changed as follows:

```
(*) and (method: POST) | SELECT date_format(__time__ - __time__ % 60, '%H:%i:%s') as time, count(1) as count GROUP BY time ORDER BY time
```



23.4.11.2.7. Markdown chart

Log Service allows you to add a Markdown chart to a dashboard. In the Markdown chart, you can insert images, links, videos, and other elements to make your dashboard page more intuitive.

Context

You can create different Markdown charts based on your business needs. Markdown charts can make a dashboard more informative. You can insert text such as background information, chart description, notes, extension information, and custom images in a Markdown chart. You can also insert saved searches or dashboard links of other projects to redirect to other query pages.

Scenarios

You can insert links in a Markdown chart to redirect to other dashboard pages of the current project. You can also insert an image corresponding to each link for better recognition. In addition, you can use a Markdown chart to describe parameters of analysis charts.

Procedure

1. [Log on to the Log Service console.](#)
2. Click the destination project name.
3. In the left-side navigation pane, click the **Dashboard** icon.
4. In the dashboard list, click the name of the destination dashboard.
5. In the upper-right corner of the dashboard page, click **Edit** to enter the edit mode.
6. In edit mode, drag the Markdown icon  from the menu bar to the specified location to create a Markdown chart.
7. Click the created Markdown chart, find the More icon in the upper-right corner of the chart, and click **edit**.

| Parameter | Description |
|-------------|---|
| Chart Name | The name of the Markdown chart. |
| Show Border | Specifies whether to show the borders of a Markdown chart. You can turn on the Show Border switch to show the borders of a Markdown chart. |

| Parameter | Description |
|-----------------|---|
| Show Title | Specifies whether to show the title of a Markdown chart. You can turn on the Show Title switch to show the title of a Markdown chart. |
| Show Background | Specifies whether to show the background of a Markdown chart. You can turn on the Show Background switch to show the background of a Markdown chart. |
| Query Binding | Specifies whether to bind a query statement to a Markdown chart. You can turn on the Query Binding switch and bind a query statement to a Markdown chart. Then, query results are dynamically displayed in the Markdown chart. |

8. (Optional) Bind a query statement.

- i. Select the destination Logstore and enter a query statement in the search box. A query statement consists of a search statement and an analytic statement in the format of `search statement|analytic statement`.
- ii. On the Search & Analysis page, click **15 Minutes(Relative)** to set the time range for the query.
- iii. Click **Search** to display the first values of the returned query result.
- iv. Click the plus sign next to a field to insert the corresponding query result to **Markdown Content**.

9. Edit **Markdown Content**.

Enter Markdown content in the **Markdown content** column. Then, data preview is displayed in real time in the **Show Chart** column on the right. You can modify the Markdown content based on the data preview.

Modify a Markdown chart

- Modify the location and size of a Markdown chart
 - i. Click **Edit** in the upper-right corner of the **Dashboard** page.
 - ii. Drag the Markdown icon to the specified location on the dashboard and drag the lower-right corner of the chart to adjust its size.
 - iii. Click **Save** in the upper-right corner of the dashboard page to save the modification.
- Modify the title of a Markdown chart
 - i. Click **Edit** in the upper-right corner of the **Dashboard** page.
 - ii. Click the specified Markdown chart, find the More icon in the upper-right corner of the chart, and click **Edit**.
 - iii. Enter a new title in the **Chart name** field and then click **OK**.
 - iv. Click **Save** in the upper-right corner of the dashboard page. In the dialog box that appears, click **OK**.
- Modify the content of a Markdown chart
 - i. Click **Edit** in the upper-right corner of the **Dashboard** page.
 - ii. Click the specified Markdown chart, find the More icon in the upper-right corner of chart, and then click **Edit**.
 - iii. Modify the chart content, and then click **OK**.
 - iv. Click **Save** in the upper-right corner of the dashboard page. In the dialog box that appears, click **OK**.
- Delete a Markdown chart
 - i. Click **Edit** in the upper-right corner of the **Dashboard** page.
 - ii. Click the specified Markdown chart, find the More icon in the upper-right corner of chart, and then click **Delete**.
 - iii. Click **Save** in the upper-right corner of the dashboard page. In the dialog box that appears, click **OK**.

Common Markdown syntax

- Title

Markdown syntax:

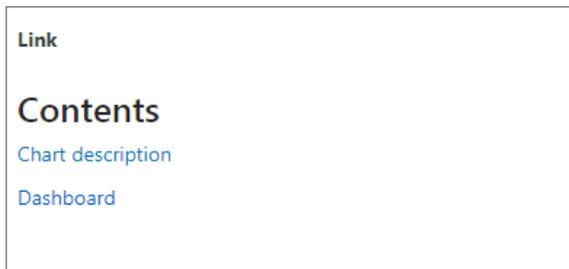
```
# Level 1 heading
## Level 2 heading
### Level 3 heading
```

- Link

Markdown syntax:

```
### Contents
[Chart description](https://xxx)
[Dashboard](https://xxx)
```

Link preview



- Image

Markdown syntax:

```
<div align=center>
![ Alt txt][id]
With a reference later in the document defining the URL location
[id]: https://octodex.github.com/images/dojocat.jpg "The Dojocat"
```

- Special tag

Markdown syntax:

```
---
__Advertisement :)__
==some mark== `some code`
> Classic markup: :wink: :crush: :cry: :tear: :laughing: :yum:
>> Shortcuts (emoticons): :-) 8-) ;)
__This is bold text__
*This is italic text*
---
```

23.5. Alerts

23.5.1. Overview

Log Service enables you to configure alerts for charts on a dashboard to monitor the service status in real time.

You can configure alerts on the **Search & Analysis** page of a Logstore or on a **Dashboard** page. When you configure an alert, you must configure the alert name, trigger condition, notification method, and other parameters. After you [Configure alerts](#), Log Service checks the query results on the dashboard at an interval and sends an alert notification if the check results meet the specified conditions. In this way, Log Service facilitates real-time monitoring of the service status.

Limits

| Item | Description |
|------------------------|---|
| Associated charts | The number of charts that can be associated with an alert ranges from 1 to 3. |
| String | If the length of a string exceeds 1,024 characters, only the first 1,024 characters are computed during a query. |
| Conditional expression | <ul style="list-style-type: none"> The conditional expression must be 1 to 128 characters in length. The conditional expression is evaluated based on the first 100 log entries returned for a query. The conditional expression can be evaluated up to 1,000 times. If the conditional expression is not matched, the alert is not triggered. |
| Search Period | The time range of each query statement cannot exceed 24 hours. |

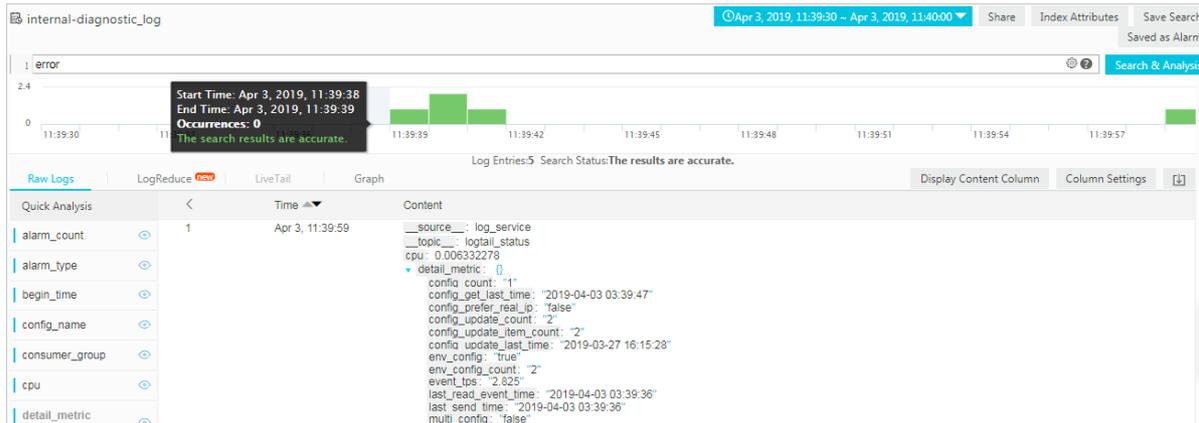
Query statement in an alert

An alert is configured based on analysis charts on a dashboard. An analysis chart is a visualized query result of a query statement. A query statement can include a search statement and an analytic statement.

- If you use only a search statement for a query, log data that matches the search condition is returned.
- If you include search and analytic statements for a query, log data that matches the search condition is analyzed before being returned.
- Search statement

For example, you want to query the data that contains "error" information in the last 15 minutes. The search statement is error. A total of 154 log entries are retrieved. Each log entry consists of key-value pairs. You can set an alert rule for the value of a key.

Note If over 100 log entries are returned for a query, only the first 100 log entries are used for evaluating the conditions set in an alert. An alert is triggered when any of the first 100 log entries returned meets the conditions.

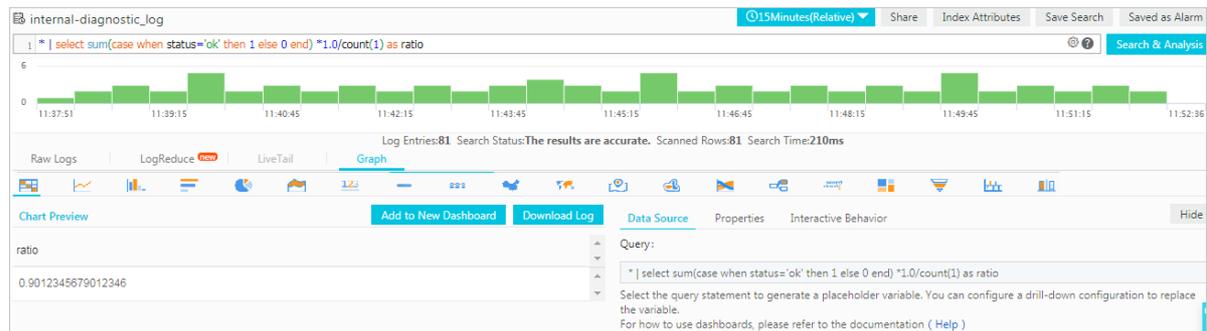


- Search and analytic statement

For example, the following statement queries the ratio of log entries whose status field value is ok to all log entries. For more information about query syntax, see [Query syntax](#).

```
* | select sum(case when status='ok' then 1 else 0 end) *1.0/count(1) as ratio
```

If you set the trigger condition of an alert to `ratio < 0.9`, the alert is triggered when the ratio of log entries whose status field value is ok to all log entries is less than 90%.



23.5.2. Configure an alarm

23.5.2.1. Configure alerts

Log Service allows you to configure alerts on the Search & Analysis page of a Logstore or on a dashboard page. If the trigger condition of an alert is met, the alert is triggered and a notification is sent to specified recipients.

Prerequisites

- Log data is collected.
- The indexing feature is enabled and indexes are configured. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

Context

Alerts are configured based on analysis charts. When you view an analysis chart, you can add the chart to a dashboard and configure an alert for the chart. You can also configure an alert for the existing charts on a dashboard.

- Create a chart and configure an alert for the chart

You can save the results of a query statement as a chart on a dashboard, and configure an alert for the chart. When you configure an alert on the Search & Analysis page, you must specify the name of the dashboard on which the chart is saved and the chart name.



- Configure an alert for existing charts on a dashboard.

You can configure an alert for one or more charts on a dashboard at a time. When you configure an alert for multiple charts, you can specify a conditional expression for each chart and combine the conditional expressions into the trigger condition for the alert.

This topic describes how to configure an alert for existing charts on a dashboard.

Note If an alert is configured for a chart on a dashboard and you update the search and analytic statement of the chart, you must update the search and analytic statement in the alert configuration. For more information, see [Modify an alert](#).

For information about example alert configurations, see [FAQ about alerts](#).

Procedure

1. [Log on to the Log Service console.](#)
2. In the **Projects** section, click the name of a project.
3. In the left-side navigation pane, click the **Dashboard** icon.
4. Click the target dashboard name.
5. In the upper-right corner of the dashboard, choose **Alerts > Create**.
6. In the Create Alert wizard, configure an alert and click **Next**.

The following table describes the configuration parameters of an alert.

| Parameter | Description |
|--------------------------|--|
| Alert Name | The name of the alert. The name must be 1 to 64 characters in length. |
| Associated Chart | <p>The chart that is associated with the alert.</p> <p>The Search Period parameter specifies the time range of log data that Log Service reads when you query data. You can select a relative time or a time frame. For example, if you set Search Period to 15 minutes (relative) and start the query at 14:30:06, Log Service reads the log data that was written from 14:15:06 to 14:30:06. If you set Search Period to 15 minutes (time frame) and start the query at 14:30:06, Log Service reads the log data that was written from 14:15:00 to 14:30:00.</p> <p>To associate the alert with multiple charts, you must separately add and configure the charts. The number before the chart name indicates the sequence number of the chart in the alert configuration. You can use the sequence number to associate a chart in the trigger condition.</p> |
| Frequency | The time interval at which Log Service executes the alert. |
| Trigger Condition | <p>The conditional expression that determines whether to trigger the alert. If the condition is met, Log Service sends an alert notification based on the specified Check Frequency and Notification Interval.</p> <p>For example, you can enter <code>pv%100 > 0 && uv > 0</code> in the Trigger Condition field.</p> <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p>Note In the conditional expression, you can use <code>\$sequence number</code> to differentiate charts. For example, you can use <code>\$0</code> to identify the chart whose sequence number is 0.</p> </div> |
| Advanced | |

| Parameter | Description |
|---------------------------------------|---|
| Notification Trigger Threshold | <p>If the number of times that the trigger condition is met exceeds this threshold and the specified Notification Interval elapses, Log Service sends an alert notification to the specified recipients.</p> <p>The default value of Notification Trigger Threshold is 1. This value indicates that each time the specified Trigger Condition is met, Log Service checks Notification Interval to determine whether to send notifications.</p> <p>You can set a custom value. This way, Log Service sends an alert notification to the specified recipient only after the trigger condition is met multiple times. For example, if you set the value to 100, Log Service checks Notification Interval only after the trigger condition is met 100 times. If the Notification Trigger Threshold and Notification Interval are reached, Log Service sends an alert notifications to the specified recipients. The overall count is then reset to zero. If Log Service fails to check log data due to exceptions such as network failures, the overall count does not change.</p> |
| Notification Interval | <p>The time interval at which Log Service sends an alert notification.</p> <p>If the number of times that the trigger condition is met exceeds the specified Notification Trigger Threshold and the specified notification interval elapses, Log Service sends an alert notification to the specified recipients. If you set this parameter to 5 minutes, you can receive up to one alert notification every 5 minutes. The default value is No Interval.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note You can use the Notification Trigger Threshold and Notification Interval parameters to control the number of alert notifications that you receive.</p> </div> |

7. Set the notification method.

Notifications can be sent to a custom webhook address in a specified format. To use this notification method, you must set the **Request URL**, **Request Method**, and **Request Content** parameters. For more information, see [Notification methods](#).

- **Request URL**: a custom webhook address, for example, `https://webhook.com/notify`.
- **Request Method**: the request method. Request methods include GET, PUT, POST, DELETE, and OPTIONS.
- **Request Content**: the content of the notification. The content must be 1 to 500 characters in length. Template variables are supported.

8. Click OK.

23.5.2.2. Grant permissions on alerts to a RAM user

This topic describes how to grant a RAM user the permissions to enable the alerting feature.

Context

Grant a RAM user the permissions only to create and modify alerts. Create a custom authorization policy, and apply the policy to the RAM user. For more information, see Procedure in this topic.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an administrator.

For more information, see [Log on to the Log Service console](#).

2. [Create a RAM role](#).
3. [Create a permission policy](#).

Use the following policy and replace the `<Project name>` with the actual project name.

```
{
  "Version": "1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "log:CreateLogStore",
        "log:CreateIndex",
        "log:UpdateIndex"
      ],
      "Resource": "acs:log:*:*:project/<Project name>/logstore/internal-alert-history"
    },
    {
      "Effect": "Allow",
      "Action": [
        "log:CreateDashboard",
        "log:CreateChart",
        "log:UpdateDashboard"
      ],
      "Resource": "acs:log:*:*:project/<Project name>/dashboard/*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "log:*"
      ],
      "Resource": "acs:log:*:*:project/<Project name>/job/*"
    }
  ]
}
```

4. [Create a user](#).
5. [Create a RAM user group](#).
6. [Add a RAM user to a RAM user group](#)
7. [Grant permissions to a RAM role](#).

23.5.2.3. Configure alert notification methods

This topic describes how to configure the notification methods for an alert in Log Service. Log Service supports the following alert notification methods: custom webhooks and DingTalk chatbot webhooks.

Custom webhooks

You can set the alert notification method to WebHook-Custom. If an alert is triggered, Log Service sends an alert notification to the custom webhook URL.

 **Note** The timeout period for this method is 5 seconds. If no response is received within 5 seconds after a notification is sent, it means that the notification has failed.

1. When you [Configure alerts](#), select **WebHook-Custom** from the **Notifications** drop-down list.
2. Set the following parameters.

| Parameter | Description |
|-----------------------|---|
| Request URL | The custom webhook URL. |
| Request Method | The method that is used to send the notification. Request available methods include GET, POST, DELETE, PUT, and OPTIONS. The default request header is Content-Type: application/json; charset=utf-8. If you need to add request headers, click Add Request Headers . |
| Content | The notification content is specified by default. The value must be 1 to 500 characters in length. You can specify a custom value. Template variables are supported. For more information, see Template variables . |

3. Click **Submit**.

DingTalk chatbot webhooks

If you set the notification method to WebHook-DingTalk Bot, Log Service sends alert notifications by using a DingTalk chatbot to the DingTalk group to which a specified webhook URL points. The chatbot can also remind the specified contacts of the alert notifications.

 **Note** Each DingTalk chatbot can send a maximum of 20 alert notifications per minute.

1. Create a DingTalk chatbot.
 - i. Open DingTalk and go to a DingTalk group.
 - ii. In the upper-right corner of the chat window, click the **Group Settings** icon and choose **Group Assistant > Add Robot**.
 - iii. In the **ChatBot** dialog box, click the + icon in the **Add Robot** section.
 - iv. In the Robot details dialog box, select **Custom (Custom message services via Webhook)** and click **Add**.
 - v. In the **Add Robot** dialog box, enter a **chatbot name** and select **security options** based on your business requirements. Then, select **I have read and accepted DingTalk Custom Robot Service Terms of Service** and click **Finished**.

 **Note** We recommend that you set the **Security Settings** parameter to **Custom Keywords**. You can set a maximum of 10 keywords. Each message must contain at least one keyword. We recommend that you set a keyword to **Alert**. 关于安全设置, 更多信息, 请参见 [钉钉开放平台](#).

- vi. Click **Copy** to copy the webhook URL.
2. Configure a notification method in the Log Service console.
 - i. When you **Configure alerts**, select **WebHook-DingTalk Bot** from the **Notifications** drop-down list.

ii. Set the following parameters.

| Parameter | Description |
|-------------|--|
| Request URL | The webhook URL of the DingTalk chatbot. Paste the webhook URL that you copied in Step 1 . |
| Title | The alert topic. The value must be 1 to 100 characters in length. You can specify a custom value. Template variables are supported. For more information, see Template variables . |
| Recipients | The group members whom you want to remind to read the alert notification. Valid values: None, All, and Specified Members. If you select Specified Members , enter the mobile phone numbers of the group members in the Tagged List field. Separate multiple mobile phone numbers with commas (,). |
| Content | <p>The notification content is specified by default. You can modify the content based on your business requirements. The value must be 1 to 500 characters in length. You can specify a custom value. Template variables are supported. For more information, see Template variables.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 10px; margin-top: 10px;"> <p>Note If you need to remind a group member to read the alert notification, use the <code>@<Mobile phone number of the group member></code> syntax in the Content field.</p> </div> |

iii. Click **Submit**.

Template variables

You must set **Content** for each notification method. In the notification content, you can reference some template variables in the `${fieldName}` format for the alert. When Log Service sends an alert notification, Log Service replaces the template variables referenced in the **Content** field with the actual values. For example, Log Service replaces `${Project}` with the name of the project to which the alert belongs.

Note You must reference valid variables. If a referenced variable does not exist or is invalid, Log Service processes the variable as a null string. If the value of a referenced variable is of the object type, the value is converted and displayed as a JSON string.

The following table describes all available variables and how to reference these variables for an alert.

| Variable | Description | Example | Reference example |
|-----------|---|--|--|
| Aliuid | The ID of the Apsara Stack tenant account to which the project belongs. | 1234567890 | The alert that is configured by the user <code>\${Aliuid}</code> is triggered. |
| Project | The project to which the alert belongs. | my-project | The alert that is configured in the project <code>\${Project}</code> is triggered. |
| AlertID | The alert ID | 0fdd88063a611aa114938f9371daeeb6-1671a52eb23 | The ID of the triggered alert is <code>\${AlertID}</code> . |
| AlertName | The name of the alert. The name must be unique in a project. | alert-1542111415-153472 | The <code>\${AlertName}</code> alert is triggered. |

| Variable | Description | Example | Reference example |
|------------------|---|--|---|
| AlertDisplayName | The display name of the alert rule. | My alert | The \${AlertDisplayName} alert is triggered. |
| Condition | The conditional expression that triggers the alert. Each variable in the conditional expression is replaced by the value that triggers the alert. The value is enclosed in brackets []. | [5] > 1 | The conditional expression that triggers the alert is \${Condition}. |
| RawCondition | The original conditional expression that triggers the alert is \${RawCondition}. | count > 1 | The trigger condition is \${RawCondition}. |
| Dashboard | The name of the dashboard with which the alert is associated. | mydashboard | The alert rule is associated with the \${Dashboard} dashboard. |
| DashboardUrl | The URL of the dashboard with which the alert is associated. | https://sls.console.aliyun.com/next/project/myproject/dashboard/mydashboard | The URL of the dashboard that is associated with the alert is \${DashboardUrl}. |
| FireTime | The time when the alert is triggered. | 2018-01-02 15:04:05 | The alert is triggered at \${FireTime}. |
| FullResultUrl | The URL that is used to query the history records that the alert was triggered. | https://sls.console.aliyun.com/next/project/myproject/logsearch/internal-alert-history?endTime=1544083998&queryString=AlertID%3A9155ea1ec10167985519fccede4d5fc7-1678293caad&queryTimeType=99&startTime=1544083968 | Click \${FullResultUrl} to view details. |

| Variable | Description | Example | Reference example |
|----------|---|---|---|
| Results | <p>The parameters and results of each log data query. The value is of the array type. For more information, see Fields in alert log entries.</p> <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p> Note A maximum of 100 alert notifications can be sent.</p> </div> | <pre>[{ "EndTime": 1542507580, "FireResult": { "__time__": "1542453580", "count": "0" }, "LogStore": "test- logstore", "Query": "* SELECT COUNT(*) as count", "RawResultCount": 1, "RawResults": [{ "__time__": "1542453580", "count": "0" }], "StartTime": 1542453580 }]</pre> | <p>The first query starts at <code>\$(Results[0].StartTime)</code> and ends at <code>\$(Results[0].EndTime)</code>. The alert has been triggered <code>\$(Results[0].FireResult.count)</code> times.</p> <div style="border: 1px solid #add8e6; padding: 5px; margin-top: 10px;"> <p> Note In this example the value 0 is the sequence number of a chart. For more information, see How can I view the serial number of a chart?</p> </div> |

23.5.3. Modify and view an alarm

23.5.3.1. Modify an alert

This topic describes how to modify an alert. After you create an alert, you can modify the alert and then update the alert. To modify an alert associated with a search statement, you can directly modify the search statement.

Precautions

- You can modify only search statements with which alerts are associated. You cannot modify search statements to search and analytic statements, which are in the format of `search statement|analytic statement`.

For example, after you associate the `request_method:GET` statement with an alert, you can modify the statement to `error`, but you cannot modify the statement to `error|select count(1) as c`.

- To modify an alert, you can click **Modify Settings** on the **Alert Overview** page, or choose **Alerts > Modify** on the associated dashboard.

Modify the search statement associated with an alert

If you associate a search statement with an alert, you can modify the search statement to modify the alert.

1. [Log on to the Log Service console](#).
2. Click a project name.
3. In the left-side navigation pane, click the **Dashboard** icon.
4. In the dashboard list, click the name of the target dashboard.
5. On the dashboard, choose **Alerts > Modify**.
6. Find the search statement, and then click .

You can modify only search statements with which alerts are associated. You cannot modify search statements to search and analytic statements, which are in the format of `search statement|analytic statement`.

7. On the dialog box that appears, enter a new search statement, click **Preview**, and then click **OK** after the search statement is verified.
8. Modify other parameters specific to your environment, such as **Frequency** and **Trigger Condition**, and then click **Next**.
9. Set the notification method, and then click **Submit**.

Modify the chart associated with an alert

After you create an alert, you can modify the chart associated with the alert to modify the alert.

1. In the dashboard list, click the name of the target dashboard.
2. On the dashboard, choose **Alerts > Modify**.
3. Find the **Associated Chart**, and then click  next to **Query**.
4. On the dialog box that appears, enter a new query statement, click **Preview**, and then click **OK** after the query statement is verified.
5. Modify other parameters specific to your environment, such as **Frequency** and **Trigger Condition**, and then click **Next**.
6. Set the notification method.
7. Click **Submit**. The new settings take effect immediately.

23.5.3.2. View history alerts

This topic describes how to view history alerts in the Log Service console. Log Service records alerts as log data and creates a dashboard to display alert details.

View history alerts in the Logstore

When you create an alert, Log Service creates a Logstore named **internal-alert-history** for the project to which the alert belongs. A log entry is generated and written to the Logstore each time the alert rule is executed, regardless of whether the alert is triggered. For more information about the fields in the log entry, see [Fields in alert log entries](#).

 **Note** The Logstore does not incur fees and cannot be deleted or modified. Each alert log entry is retained in the Logstore for seven days.

1. [Log on to the Log Service console](#).
2. Click a project name.

3. Click the  icon next to the **internal-alert-history** Logstore, and then select **Search & Analysis**.
4. On the page that appears, query alert log entries based on your needs.

View history alerts on the dashboard

After you create an alert, Log Service creates a dashboard named **internal-alert-analysis** for the project to which the alert belongs. The dashboard displays the statistics of all previous alerts, including the number of triggered alerts, percentage of successful alerts and notifications, and top 10 alerts whose alert rules are executed.

 **Note** The dashboard cannot be deleted or modified.

1. In the left-side navigation pane, click the **Dashboard** icon.
2. Click **Alert History Statistics** to open the dashboard page.

On the **Alert History Statistics** dashboard, the details of history alerts are displayed, including whether the alerts are triggered, why the alerts are triggered, error information, and other information.

23.5.3.3. Manage an alert

This topic describes how to manage an alert after you create the alert.

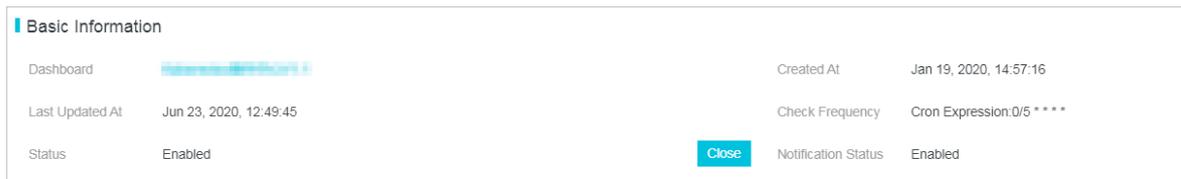
Context

You can disable, enable, modify, and delete the alert, or view the details of the alert such as the update time.

View the details of an alert

1. [Log on to the Log Service console](#).
2. Click a project name.
3. In the left-side navigation pane, click the **Alerts** icon.
4. In the alert list, click the name of the target alert.

On the **Alert Overview** page, you can view the details of the alert, such as the dashboard, creation time, last update time, check frequency, alert status, and notification status.



The screenshot shows the 'Basic Information' section of an alert overview page. It contains the following details:

- Dashboard:** 
- Last Updated At:** Jun 23, 2020, 12:49:45
- Status:** Enabled
- Created At:** Jan 19, 2020, 14:57:16
- Check Frequency:** Cron Expression: 0/5 * * * *
- Notification Status:** Enabled

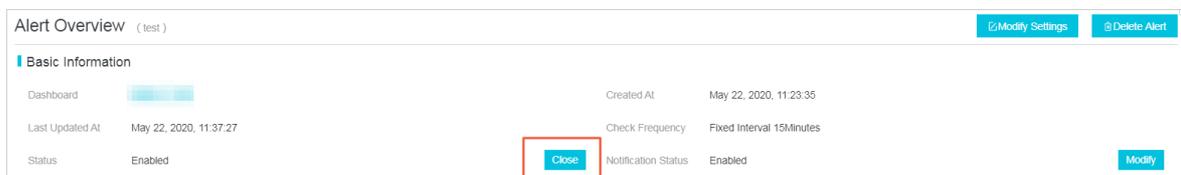
There is a blue 'Close' button next to the 'Status' field.

Disable and enable an alert

After you create an alert, you can disable or enable the alert. If you disable an alert, Log Service no longer checks the alert or send alert notifications.

1. In the left-side navigation pane, click the **Alerts** icon.
2. In the alert list, click the name of the target alert.

On the **Alert Overview** page, click **Enable** or **Close** in the **Alert Status** field.



The screenshot shows the 'Alert Overview' page for an alert named '(test)'. It includes 'Basic Information' and 'Alert Status' sections. The 'Alert Status' section has a 'Close' button highlighted with a red box. Other details include:

- Dashboard:** 
- Last Updated At:** May 22, 2020, 11:37:27
- Status:** Enabled
- Created At:** May 22, 2020, 11:23:35
- Check Frequency:** Fixed Interval 15Minutes
- Notification Status:** Enabled

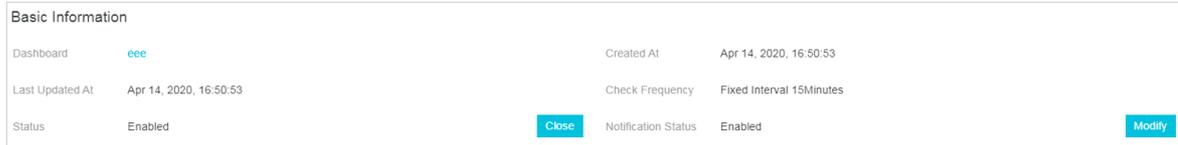
Buttons for 'Modify Settings', 'Delete Alert', 'Close', and 'Modify' are visible.

Disable and enable alert notifications

You can disable notifications of enabled alerts. After you disable notifications of an alert, alert notifications are not sent even if the trigger condition is met.

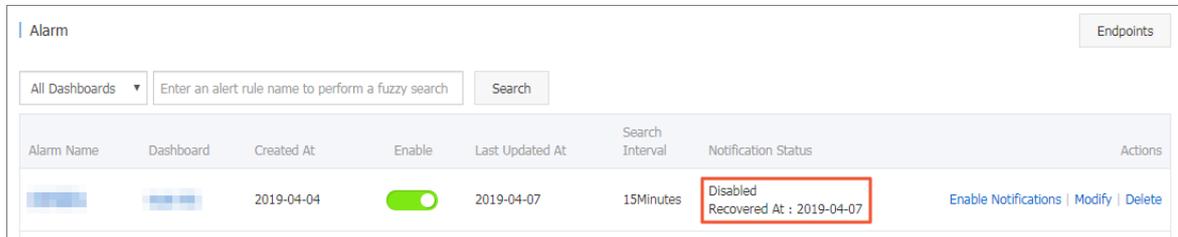
1. In the left-side navigation pane, click the **Alerts** icon.
2. In the alert list, click the name of the target alert.

On the **Alert Overview** page, click **Modify** in the **Notification Status** field.



3. Set the time range for which notifications are disabled, and then click **Confirm**.

After you disable alert notifications, you can view the time when alert notifications resume in the **Notification Status** field. You can click **Modify** in the **Notification Status** field to manually resume alert notifications.



Delete an alert

You cannot recover a deleted alert. Proceed with caution when you delete an alert.

1. In the left-side navigation pane, click the **Alerts** icon.
2. In the alert list, click the name of the target alert.
3. On the **Alert Overview** page, click **Delete Alert**.
4. In the dialog box that appears, click **OK**.

23.5.4. Relevant syntax and fields for reference

23.5.4.1. Conditional expression syntax of an alert

This topic describes how to configure a conditional expression for an alert in Log Service. An alert is triggered if the conditional expression configured for the alert is met.

In determining whether the conditional expression of an alert is met, the results of query statements configured for the alert are used as the inputs and the log fields in the conditional expression are used as the variables. If the conditional expression is met, the alert is triggered.

Limits

- Negative numbers must be enclosed with parentheses (), for example, $x + (-100) < 100$.
- Numeric data is treated as 64-bit floating-point numbers. If a comparison is performed, errors may occur.
- A variable can contain only letters and digits and must start with a letter.
- A conditional expression can be up to 128 characters in length.
- A conditional expression can be evaluated up to 1,000 times. If an alert is configured for multiple charts and the conditional expression of the alert is not met after 1,000 times of evaluation, the alert is not triggered.
- A maximum of three charts can be associated with an alert.

- An alert is triggered if and only if the Boolean value of its conditional expression is true. For example, the result of the expression `100+100` is 200, which cannot trigger the alert.
- `true` , `false` , `$` , and `.` are reserved and cannot be used as variables.

Basic syntax

The following table lists the syntax supported in a conditional expression.

| Syntax | Description | Examples |
|----------------------|--|--|
| Basic operators | Supports the addition operator (+), subtraction operator (-), multiplication operator (*), division operator (/), and modulus operator (%). | <code>x*100+y>200</code> <code>x%10>5</code> |
| Comparison operators | Supports eight comparison operators, including the greater than operator (>), greater than or equal to operator (>=), less than operator (<), less than or equal to operator (<=), equal to operator (==), not equal to operator (!=), match operator (=~), and mismatch operator (!~). <div style="border: 1px solid #add8e6; padding: 5px; background-color: #e6f2ff;"> <p> Note</p> <ul style="list-style-type: none"> • Backslashes (/) must be escaped. • Regular expressions support the RE2 syntax. </div> | <code>x >= 0</code> <code>x < 100</code> <code>x <= 100</code> <code>x == 100</code> <code>x == "foo"</code> Regular expression: <code>x =~ "\\w +"</code> |
| Logical operators | Supports the logical operators AND (&&) and OR (). | <code>x >=0&& y <=100</code> <code>x > 0 y > 0</code> |
| Logical negation | Supports the logical negation operator (!). | <code>!(a < 1 && a > 100)</code> |
| Numeric constants | Numeric constants are processed as 64-bit floating-point numbers. | <code>x > 100</code> |
| String constants | String constants are formatted in a string enclosed in single quotation marks (''). | <code>foo == 'string'</code> |
| Boolean constants | Supports true and false. | <code>(x > 100) == true</code> |
| Parentheses | Parentheses () can be used to enforce precedence order. | <code>x*(y+100)>100</code> |
| contains function | Determines whether a substring is included in a string. For example, if the result of the <code>contains(field, 'xxxx')</code> function is true, the field string includes the <code>xxxx</code> substring. | <code>contains(foo, 'hello')</code> |

Conditional expression for results of multiple query statements

- Syntax

If you configure an alert for multiple charts, the variables in a conditional expression must be prefixed. In this way, you can specify from which query result to obtain the corresponding values of the variables when it evaluates your expression. The format is `$N.fieldname` , where N is the sequence number of the query statement. You can configure up to three query statements in an alert. Therefore the value range of N is 0 to 2. For example, `$0.foo` indicates to obtain the value of the `foo` field returned from the first query statement. If you configure one query statement in an alert, you do not need to specify the prefix.

 **Note** How can I view the sequence number of a query statement?

In the Alert Configuration step, the **Associated Chart** parameter specifies the sequence number of each query statement (chart). The first query statement is numbered 0, the second query statement is numbered 1, and the third statement is numbered 2.

- Evaluate a conditional expression

If multiple query results are returned, the variables specified in the conditional expression determine how to use the results to evaluate the conditional expression. For example, if you configure three query statements in an alert, x, y, and z log entries are returned when you execute each of the three statements. The conditional expression that you configure for the alert is `$0.foo > 100 && $1.bar < 100`. Then the first two query results are used to evaluate the conditional expression. A maximum of $x \times y$ times of evaluation (or 1,000 if $x \times y$ is greater than 1,000) is performed. If the conditional expression is met within the maximum times of evaluation, true is returned. Otherwise, false is returned.

Operations

 **Note**

- 64-bit floating-point numbers are used in a conditional expression.
- Each string constant must be enclosed in single quotation marks (') or double quotation marks (""), for example, 'string', and "string".
- Boolean values include true and false.

| Operator | Operation | | |
|--|---|--|--|
| | Operation between variables | Operation between non-string constants and variables | Operation between string constants and variables |
| Basic operators, including the addition operator (+), subtraction operator (-), multiplication operator (*), division operator (/), and modulus operator (%) | The left and right operands are converted to numbers before being operated. | | Unsupported. |

| Operator | Operation | | |
|---|---|---|---|
| | Operation between variables | Operation between non-string constants and variables | Operation between string constants and variables |
| <p>Comparison operators, including the greater than operator (>), greater than or equal to operator (>=), less than operator (<), less than or equal to operator (<=), equal to operator (==), not equal to operator (!=)</p> | <p>Operations are performed based on the following priorities:</p> <ol style="list-style-type: none"> 1. The left and right operands are converted to numbers before being operated based on the numerical order. If the conversion fails, 2. operands are operated based on the alphabetical order of strings. | <p>The left and right operands are converted to numbers before being operated based on the numerical order.</p> | <p>The left and right operands are operated based on the alphabetical order of strings.</p> |
| <p>Regular expression match operator (=~) and regular expression mismatch operator (!~)</p> | <p>The left and right operands are operated based on the alphabetical order of strings.</p> | <p>Unsupported.</p> | <p>The left and right operands are operated based on the alphabetical order of strings.</p> |
| <p>Logical operators, including AND (&&) and OR ()</p> | <p>These two operators cannot be directly used on the fields in query results. The left and right operands must be sub-expressions, and the values of the sub-expressions are of the Boolean type.</p> | | |
| <p>Logical negation (!)</p> | <p>This operator cannot be directly used on the fields in query results. The left and right operands must be sub-expressions, and the values of the sub-expressions are of the Boolean type.</p> | | |
| <p>contains function</p> | <p>The left and right operands are operated based on the alphabetical order of strings.</p> | <p>Unsupported.</p> | <p>The left and right operands are operated based on the alphabetical order of strings.</p> |
| <p>Parentheses ()</p> | <p>Parentheses () enforce precedence order.</p> | | |

23.5.4.2. Fields in alert log entries

After you configure an alert, Log Service automatically creates a Logstore to store log entries that are generated when alert rules are executed and notifications are sent. This topic describes fields in alert log entries.

Fields

| Field | Description | Example |
|------------------|--|--|
| AlertDisplayName | The display name of an alert. | Test alert rules |
| AlertID | The unique ID of an alert. | 0fdd88063a611aa114938f9371daeeb6-1671a52eb23 |
| AlertName | The unique name of an alert in a project. | alert-1542111415-153472 |
| Condition | The conditional expression configured for an alert. | \$0.count > 1 |
| Dashboard | The dashboard associated with an alert. | my-dashboard |
| FireCount | The cumulative times that an alert has been triggered since the last notification was sent. | 1 |
| Fired | Indicates whether an alert was triggered. Valid values: true and false. | true |
| LastNotifiedAt | The time when the last notification was sent. The time is displayed in a Unix timestamp. | 1542164541 |
| NotifyStatus | The status of a notification. Valid values: <ul style="list-style-type: none"> Success: indicates that a notification was successfully sent. Failed: indicates that a notification failed to be sent. NotNotified: indicates that no notification was sent. PartialSuccess: indicates that the notification sending partially succeeded. | Success |
| Reason | The reason that a notification failed to be sent or no notification was sent. | result type is not bool |

| Field | Description | Example |
|---------|--|--|
| Results | The parameters and results of each log data query. The value is of the array type. For information about parameters in the Results field, see Parameters in the Result field . | <pre>[{ "EndTime": 1542334900, "FireResult": null, "LogStore": "test-logstore", "Query": "* select count(1) as count", "RawResultCount": 1, "RawResults": [{ "__time__": "1542334840", "count": "0" }], "StartTime": 1542334840 }]</pre> |
| Status | The execution result of an alert. Valid values: Success and Failed. | Success |

Parameters in the Result field

| Parameter | Description | Example |
|-------------|---|---|
| Query | The query statement that is configured in an alert. | * select count(1) as count |
| LogStore | The target Logstore of a query. | my-logstore |
| StartTime | The time when a query starts. | 2019-01-02 15:04:05 |
| StartTimeTs | The time when a query starts. The time is in the Unix timestamp format. | 1542334840 |
| EndTime | The time when a query ends. | 2019-01-02 15:19:05 |
| EndTimeTs | The time when a query ends. The time is in the Unix timestamp format. The actual query time range is <code>[StartTime, EndTime)</code> . | 1542334900 |
| RawResults | The raw query result. The parameter value is formatted in an array where each element is a log entry. The maximum length of the array is 100. | <pre>[{ "__time__": "1542334840", "count": "0" }]</pre> |

| Parameter | Description | Example |
|----------------|--|---|
| RawResultsAsKv | The query result that is formatted in key-value pairs. Note This parameter can only be used as a template variable. It is not saved to a Logstore. | [foo:0] |
| RawResultCount | The number of log entries that are returned in the RawResults parameter. | 1 |
| FireResult | The log entry that records the triggering of an alert. If an alert is not triggered, the parameter value is null. | <pre>{ "__time__": "1542334840", "count": "0" }</pre> |
| FireResultAsKv | The log entry that records the triggering of an alert, formatted in key-value pairs. Note This parameter can only be used as a template variable. It is not saved to a Logstore. | [foo:0] |

23.6. Real-time consumption

23.6.1. Overview

Log Service allows you to consume log data by using multiple methods.

After data is collected to LogHub, you can consume the log data by using two methods. The following table describes the methods.

| Method | Scenarios | Timeliness | Retention period |
|--------------------------------|--------------------------|------------|------------------|
| Real-time consumption (LogHub) | Real-time computing | Real time | Custom |
| Indexing and query (LogSearch) | Online query of hot data | Real time | Custom |

Real-time consumption

LogHub allows you to pull log data and consume the data in real time. The following procedure describes how log data in a shard is consumed:

1. Obtain a cursor based on the start time and end time of data consumption.
 2. Read log data based on the cursor and step parameters and return the position of the next cursor.
 3. Move the cursor to continuously consume log data.
- Consume log data by using an SDK

You can use Log Service SDKs in multiple programming languages such as Java, Python, and Go to consume log data.

- Consume log data by using consumer groups

Log Service provides an advanced method that allows you to consume logs by using consumer groups. A consumer group is a lightweight computing framework that allows multiple consumers to consume data from a Logstore at the same time. The consumers in a consumer group are automatically allocated shards. Data is consumed in order based on the time when it is written to the Logstore. In addition, the consumers can use checkpoints to resume consumption from a breakpoint. You can use consumer group SDKs in multiple programming languages such as Go, Python, and Java to consume log data.

- Log consumption by using real-time stream processing systems

- Use [Spark Streaming clients](#) to consume log data.
- Use [Storm spouts](#) to consume log data.
- Use [Flink Connector](#) to consume log data.

- Log consumption by using open-source services

Use [Flume](#) to consume log data and import log data to Hadoop file system (HDFS) instances.

Log search and analytics

- You can query log data in the Log Service console.
- You can also query log data by using an SDK or the API of Log Service. Log Service provides an HTTP-based RESTful API. You can call all log query API operations that are provided by Log Service.

23.6.2. Consume log data

Log Service provides SDKs in various programming languages, such as Java, Python, and Go. You can use an SDK to consume log data.

Use an SDK to consume log data

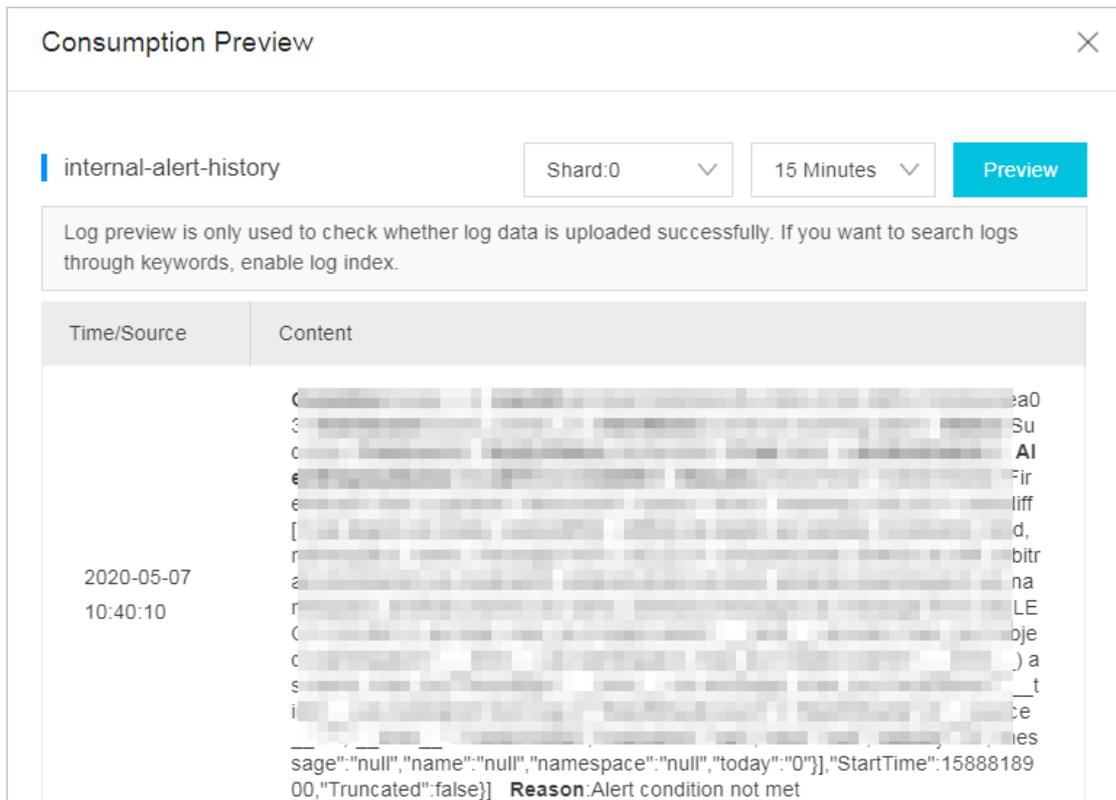
The following example shows how to use the SDK for Java to consume log data in a shard:

```
Client client = new Client(host, accessId, accessKey);
String cursor = client.GetCursor(project, logStore, shardId, CursorMode.END).GetCursor();
System.out.println("cursor = " + cursor);
try {
    while (true) {
        PullLogsRequest request = new PullLogsRequest(project, logStore, shardId, 1000, cursor);
        PullLogsResponse response = client.pullLogs(request);
        System.out.println(response.getCount());
        System.out.println("cursor = " + cursor + " next_cursor = " + response.getNextCursor());
        if (cursor.equals(response.getNextCursor())) {
            break;
        }
        cursor = response.getNextCursor();
        Thread.sleep(200);
    }
}
catch (LogException e) {
    System.out.println(e.GetRequestId() + e.GetErrorMessage());
}
```

Preview log data in the Log Service console

Log preview is a way of log data consumption. To preview log data that is stored in a Logstore in the Log Service console, perform the following steps:

1. [Log on to the Log Service console.](#)
2. In the Projects section, click the target project.
3. In the Logstore list, find the target Logstore, click the  icon next to the Logstore, and then select **Consumption Preview**.
4. In the Consumption Preview dialog box, select a shard, set a time range, and then click **Preview**.
The log preview page displays the log data of the first 10 packets in the specified time range.



23.6.3. Consumption by consumer groups

23.6.3.1. Use consumer groups to consume log data

Log consumption by using consumer groups

Consumer groups allow you to focus on the business logic during log data consumption. You do not need to consider factors such as Log Service implementation, load balancing among consumers, and failovers that may be introduced when you use an SDK to consume log data.

Terms

The following table describes the terms of consumer groups and consumers.

| Term | Description |
|----------------|--|
| consumer group | A consumer group consists of multiple consumers. Each consumer in a consumer group consumes different data in a Logstore. |
| consumer | The consumers in a consumer group consume data from specified data sources. Each consumer name in a consumer group must be unique. |

A Logstore has multiple shards. A consumer library allocates shards to consumers in a consumer group based on the following principles:

- Each shard can be allocated to one consumer.
- Each consumer can consume data from multiple shards.

After a new consumer joins a consumer group, the shards allocated to the consumer group are reallocated to each consumer for load balancing. The reallocation is based on the preceding principles and cannot be viewed by users.

A consumer library stores checkpoints. This allows consumers to resume consumption from a breakpoint and avoid repetitive consumption after a program fault is resolved.

Procedure

Log consumption by using consumer groups is implemented in Java or Python. The following procedure takes Java as an example to describe how consumer groups consume log data.

1. Add Maven dependencies.

```
<dependency>
  <groupId>com.google.protobuf</groupId>
  <artifactId>protobuf-java</artifactId>
  <version>2.5.0</version>
</dependency>
<dependency>
  <groupId>com.aliyun.openservices</groupId>
  <artifactId>loghub-client-lib</artifactId>
  <version>0.6.16</version>
</dependency>
```

2. Create a file named main.java.

```

import com.aliyun.openservices.loghub.client.ClientWorker;
import com.aliyun.openservices.loghub.client.config.LogHubConfig;
import com.aliyun.openservices.loghub.client.exceptions.LogHubClientWorkerException;
public class Main {
    // The endpoint of Log Service.
    private static String sEndpoint = "cn-hangzhou.log.aliyuncs.com";
    // The name of a Log Service project.
    private static String sProject = "ali-cn-hangzhou-sls-admin";
    // The name of a Logstore.
    private static String sLogstore = "sls_operation_log";
    // The name of a consumer group.
    private static String sConsumerGroup = "consumerGroupX";
    // The AccessKey pair of an Apsara Stack tenant account or RAM user that is used to consume data.
    private static String sAccessKeyId = "";
    private static String sAccessKey = "";
    public static void main(String[] args) throws LogHubClientWorkerException, InterruptedException {
        // The second parameter is the consumer name. Each consumer name in a consumer group must be unique. However, the names of consumer groups can be the same. When different consumers start multiple processes on multiple servers to consume the data of a Logstore, you can use a server IP address to identify a consumer. The ninth parameter maxFetchLogGroupSize indicates the maximum number of log groups that are retrieved from the server at a time. Valid values: 1 to 1000. You can use the default value or specify a value based on your requirements.
        LogHubConfig config = new LogHubConfig(sConsumerGroup, "consumer_1", sEndpoint, sProject, sLogstore, sAccessKeyId, sAccessKey, LogHubConfig.ConsumePosition.BEGIN_CURSOR);
        ClientWorker worker = new ClientWorker(new SampleLogHubProcessorFactory(), config);
        Thread thread = new Thread(worker);
        // The ClientWorker instance automatically runs after the thread is executed and extends the Runnable interface
        .
        thread.start();
        Thread.sleep(60 * 60 * 1000);
        // The shutdown function of the ClientWorker instance is called to exit the consumption instance. The associated thread is automatically stopped.
        worker.shutdown();
        // Multiple asynchronous tasks are generated when the ClientWorker instance is running. We recommend that you wait for 30 seconds to ensure that all running tasks exit after shutdown.
        Thread.sleep(30 * 1000);
    }
}

```

3. Create a file named SampleLogHubProcessor.java.

```

import com.aliyun.openservices.log.common.FastLog;
import com.aliyun.openservices.log.common.FastLogContent;
import com.aliyun.openservices.log.common.FastLogGroup;
import com.aliyun.openservices.log.common.FastLogTag;
import com.aliyun.openservices.log.common.LogGroupData;
import com.aliyun.openservices.loghub.client.ILogHubCheckPointTracker;
import com.aliyun.openservices.loghub.client.exceptions.LogHubCheckPointException;
import com.aliyun.openservices.loghub.client.interfaces.ILogHubProcessor;
import com.aliyun.openservices.loghub.client.interfaces.ILogHubProcessorFactory;
import java.util.List;
public class SampleLogHubProcessor implements ILogHubProcessor {
    private int shardId;
    // Record the last persistent checkpoint time.
    private long mLastCheckTime = 0;
    public void initialize(int shardId) {
        this.shardId = shardId;
    }
    // The main logic of data consumption. All exceptions must be captured and cannot be thrown.
    public String process(List<LogGroupData> logGroups,
        ILogHubCheckPointTracker checkPointTracker) {

```

```

// Display the retrieved data.
for (LogGroupData logGroup : logGroups) {
    FastLogGroup flg = logGroup.GetFastLogGroup();
    System.out.println(String.format("\tcategory\t:\t%s\n\tsource\t:\t%s\n\ttopic\t:\t%s\n\tmachineUUID\t:\t%s\n",
        flg.getCategory(), flg.getSource(), flg.getTopic(), flg.getMachineUUID());
    System.out.println("Tags");
    for (int tagIdx = 0; tagIdx < flg.getLogTagsCount(); ++tagIdx) {
        FastLogTag logtag = flg.getLogTags(tagIdx);
        System.out.println(String.format("\t%s\t:\t%s", logtag.getKey(), logtag.getValue()));
    }
    for (int lIdx = 0; lIdx < flg.getLogsCount(); ++lIdx) {
        FastLog log = flg.getLogs(lIdx);
        System.out.println("-----\nLog: " + lIdx + ", time: " + log.getTime() + ", GetContentCount: " + log.getContentCount());
        for (int cIdx = 0; cIdx < log.getContentsCount(); ++cIdx) {
            FastLogContent content = log.getContents(cIdx);
            System.out.println(content.getKey() + "\t:\t" + content.getValue());
        }
    }
    long curTime = System.currentTimeMillis();
    // Write checkpoints to the server every 30 seconds. If a ClientWorker instance does not respond within 30 seconds,
    // a new ClientWorker instance consumes data starting from the last checkpoint. A small amount of duplicate data may exist.
    if (curTime - mLastCheckTime > 30 * 1000) {
        try {
            // If the parameter is set to true, checkpoints are synchronized to the server immediately. If the parameter is set to false, checkpoints are locally cached. The default synchronization interval of checkpoints is 60 seconds.
            checkPointTracker.saveCheckPoint(true);
        } catch (LogHubCheckPointException e) {
            e.printStackTrace();
        }
        mLastCheckTime = curTime;
    }
    return null;
}
// The ClientWorker instance calls this function upon exit. You can perform a cleanup.
public void shutdown(ILogHubCheckPointTracker checkPointTracker) {
    // Save consumption breakpoints to the server.
    try {
        checkPointTracker.saveCheckPoint(true);
    } catch (LogHubCheckPointException e) {
        e.printStackTrace();
    }
}
}
class SampleLogHubProcessorFactory implements ILogHubProcessorFactory {
    public ILogHubProcessor generatorProcessor() {
        // Generate a consumption instance.
        return new SampleLogHubProcessor();
    }
}
}

```

 **Note** Run the preceding code to display all data in a Logstore. If you want multiple consumers to consume data from the same Logstore, you can modify the code based on the comments. You can use the same consumer group name and different consumer names to start new consumption processes.

Limits and troubleshooting

You can create a maximum of 10 consumer groups for each Logstore. The `ConsumerGroupQuotaExceed` error is reported if the number of consumer groups exceeds 10.

We recommend that you configure Log4j for the consumer program to throw error messages within consumer groups. This improves the troubleshooting efficiency. If you save the `log4j.properties` file to the resources directory and run the program, the following error message appears:

```
[WARN ] 2018-03-14 12:01:52,747 method:com.aliyun.openservices.loghub.client.LogHubConsumer.sampleLogError(LogHubConsumer.java:159)
com.aliyun.openservices.log.exception.LogException: Invalid loggroup count, (0,1000]
```

The following example is a typical `log4j.properties` configuration file:

```
log4j.rootLogger = info,stdout
log4j.appender.stdout = org.apache.log4j.ConsoleAppender
log4j.appender.stdout.Target = System.out
log4j.appender.stdout.layout = org.apache.log4j.PatternLayout
log4j.appender.stdout.layout.ConversionPattern = [%-5p] %d{yyyy-MM-dd HH:mm:ss,SSS} method:%l%n%m%n
```

Advanced operations

The preceding code is suitable for log consumption in common scenarios. This section describes how to perform advanced operations in other scenarios.

- Consume data that is logged from a certain time point

`LogHubConfig` in the preceding code has two constructors:

```
// The value of the consumerStartTimeInSeconds parameter is a UNIX timestamp representing the number of seconds
that have elapsed since 00:00:00 on January 1, 1970, 00:00:00 UTC.
public LogHubConfig(String consumerGroupName,
    String consumerName,
    String loghubEndPoint,
    String project, String logStore,
    String accessId, String accessKey,
    int consumerStartTimeInSeconds);

// The position parameter is an enumeration variable. LogHubConfig.ConsumePosition.BEGIN_CURSOR indicates that
the consumption starts from the earliest data. LogHubConfig.ConsumePosition.END_CURSOR indicates that the consumption
starts from the latest data.
public LogHubConfig(String consumerGroupName,
    String consumerName,
    String loghubEndPoint,
    String project, String logStore,
    String accessId, String accessKey,
    ConsumePosition position);
```

You can use different constructors based on your requirements. However, if a checkpoint is stored on the server, you must start data consumption from this checkpoint.

- Reset a checkpoint

In scenarios such as data padding or repeated computing, you may need to set the consumption position to a time point for a consumer group. Then data consumption is started from the consumption position. To set the consumption position, you can use either one of the following two methods:

- o Delete the consumer group.
 - a. Stop the consumption processes.
 - b. Delete the consumer group from the Log Service console.
 - c. Modify the code to specify the start time point for data consumption.
 - d. Restart the consumption processes.
- o Use an SDK to reset the start time point of data consumption for the consumer group.
 - a. Stop the consumption processes.
 - b. Use an SDK to modify the checkpoint.
 - c. Restart the consumption processes.

```
public static void updateCheckpoint() throws Exception {
    Client client = new Client(host, accessId, accessKey);
    long timestamp = Timestamp.valueOf("2017-11-15 00:00:00").getTime() / 1000;
    ListShardResponse response = client.ListShard(new ListShardRequest(project, logStore));
    for (Shard shard : response.GetShards()) {
        int shardId = shard.GetShardId();
        String cursor = client.GetCursor(project, logStore, shardId, timestamp).GetCursor();
        client.UpdateCheckPoint(project, logStore, consumerGroup, shardId, cursor);
    }
}
```

Use a RAM user to access consumer groups

Before you use a RAM user to access consumer groups, you must grant relevant permissions to the RAM user. For more information, see [Grant permissions to a RAM role](#).

The following table lists the actions you can perform as a RAM user.

| Action | Resource |
|-----------------------------------|---|
| log:GetCursorOrData | acs:log: \${regionName}: \${projectOwnerAliUid}: project / \${projectName}/logstore/ \${logstoreName} |
| log:CreateConsumerGroup | acs:log: \${regionName}: \${projectOwnerAliUid}: project / \${projectName}/logstore/ \${logstoreName}/ consumergroup/* |
| log:ListConsumerGroup | acs:log: \${regionName}: \${projectOwnerAliUid}: project / \${projectName}/logstore/ \${logstoreName}/ consumergroup/* |
| log:ConsumerGroupUpdateCheckPoint | acs:log: \${regionName}: \${projectOwnerAliUid}: project / \${projectName}/logstore/ \${logstoreName}/ consumergroup/ \${consumerGroupName} |
| log:ConsumerGroupHeartBeat | acs:log: \${regionName}: \${projectOwnerAliUid}: project / \${projectName}/logstore/ \${logstoreName}/ consumergroup/ \${consumerGroupName} |
| log:UpdateConsumerGroup | acs:log: \${regionName}: \${projectOwnerAliUid}: project / \${projectName}/logstore/ \${logstoreName}/ consumergroup/ \${consumerGroupName} |
| log:GetConsumerGroupCheckPoint | acs:log: \${regionName}: \${projectOwnerAliUid}: project / \${projectName}/logstore/ \${logstoreName}/ consumergroup/ \${consumerGroupName} |

For example, a project named project-test resides in the China (Hangzhou) region. The ID of the Apsara Stack tenant account to which the project belongs is 1234567. The name of the Logstore to be consumed is logstore-test and the consumer group name is consumergroup-test. To allow a RAM user to access the consumer group, you must grant the following permissions to the RAM user:

```

{
  "Version": "1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "log:GetCursorOrData"
      ],
      "Resource": "acs:log:cn-hangzhou:1234567:project/project-test/logstore/logstore-test"
    },
    {
      "Effect": "Allow",
      "Action": [
        "log:CreateConsumerGroup",
        "log:ListConsumerGroup"
      ],
      "Resource": "acs:log:cn-hangzhou:1234567:project/project-test/logstore/logstore-test/consumergroup/*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "log:ConsumerGroupUpdateCheckPoint",
        "log:ConsumerGroupHeartBeat",
        "log:UpdateConsumerGroup",
        "log:GetConsumerGroupCheckPoint"
      ],
      "Resource": "acs:log:cn-hangzhou:1234567:project/project-test/logstore/logstore-test/consumergroup/consumergroup-test"
    }
  ]
}

```

23.6.3.2. View the status of a consumer group

This topic describes how to use the console, API, and SDK to view the status of a consumer group. Log data consumption by using consumer groups is an advanced real-time data consumption method provided by Log Service. Automatic load balancing is implemented among multiple consumption instances. Spark Streaming and Storm use consumer groups as the basic mode to consume log data.

View the consumption progress in the console

1. [Log on to the Log Service console](#).
2. In the Projects section, click the target project.
3. Find the target Logstore, and choose  > **Data Consumption**.
4. Click the consumer group whose data consumption progress you want to view. The data consumption progress of each shard in the Logstore is displayed.

Use an API or SDK to view the data consumption progress

The following example uses the SDK for Java to describe how to call API operations to view the data consumption progress.

```

package test;
import java.util.ArrayList;
import com.aliyun.openservices.log.Client;
import com.aliyun.openservices.log.common.Consts.CursorMode;
import com.aliyun.openservices.log.common.ConsumerGroup;

```

```

import com.aliyun.openservices.log.common.ConsumerGroup;
import com.aliyun.openservices.log.common.ConsumerGroupShardCheckPoint;
import com.aliyun.openservices.log.exception.LogException;
public class ConsumerGroupTest {
    static String endpoint = "";
    static String project = "";
    static String logstore = "";
    static String accessKeyId = "";
    static String accessKey = "";
    public static void main(String[] args) throws LogException {
        Client client = new Client(endpoint, accessKeyId, accessKey);
        //Query all consumer groups of this Logstore. If no consumer group exists, the length of consumerGroups is 0.
        ArrayList<ConsumerGroup> consumerGroups;
        try{
            consumerGroups = client.ListConsumerGroup(project, logstore).GetConsumerGroups();
        }
        catch(LogException e){
            if(e.GetErrorCode() == "LogStoreNotExist")
                System.out.println("this logstore does not have any consumer group");
            else{
                //internal server error branch
            }
            return;
        }
        for(ConsumerGroup c: consumerGroups){
            //Print consumer group properties, including the name, heartbeat timeout, and whether data is consumed in order
            System.out.println("Name: " + c.getConsumerGroupName());
            System.out.println("Heartbeat timeout: " + c.getTimeout());
            System.out.println("Consumption in order: " + c.isInOrder());
            for(ConsumerGroupShardCheckPoint cp: client.GetCheckPoint(project, logstore, c.getConsumerGroupName()).GetCheckPoints()){
                System.out.println("shard: " + cp.getShard());
                //Format the returned time. The time is a long integer that is accurate to milliseconds.
                System.out.println("The last time when data was consumed: " + cp.getUpdateTime());
                System.out.println("Consumer name: " + cp.getConsumer());
                String consumerPrg = "";
                if(cp.getCheckPoint().isEmpty())
                    consumerPrg = "Consumption not started";
                else{
                    //The UNIX timestamp. Unit: seconds. Format the output value of the timestamp.
                    try{
                        int prg = client.GetPrevCursorTime(project, logstore, cp.getShard(), cp.getCheckPoint()).GetCursorTime();
                        consumerPrg = "" + prg;
                    }
                    catch(LogException e){
                        if(e.GetErrorCode() == "InvalidCursor")
                            consumerPrg = "Invalid. The previous consumption time has exceeded the retention period of the data in the Logstore.";
                        else{
                            //internal server error
                            throw e;
                        }
                    }
                }
                System.out.println("Consumption progress: " + consumerPrg);
                String endCursor = client.GetCursor(project, logstore, cp.getShard(), CursorMode.END).GetCursor();
                int endPrg = 0;
                try{
                    endPrg = client.GetPrevCursorTime(project, logstore, cp.getShard(), endCursor).GetCursorTime();
                }
            }
        }
    }
}

```


- If the data volume in a shard exceeds the processing capacity of a single spout, you can split the shard to reduce its data volume.
- LogHub spouts and bolts must use the `ack` method to check whether log data is successfully sent from spouts to bolts and whether the data is successfully processed by the bolts.

Examples

- Use the following code to create spouts and construct a topology:

```

public static void main( String[] args )
{
    String mode = "Local"; // Specify to use the local test mode.
    String conumser_group_name = ""; // Specify a unique consumer group name for each topology. The name cannot be an empty string. It must be 3 to 63 characters in length and can contain lowercase letters, digits, hyphens (-), and underscores (_). It must start and end with a lowercase letter or digit.
    String project = ""; // Specify the project in Log Service.
    String logstore = ""; // Specify the Logstore in Log Service.
    String endpoint = ""; // Specify the endpoint of Log Service.
    String access_id = ""; // Specify the AccessKey ID of the user.
    String access_key = "";
    // Configure a LogHub Storm spout.
    LogHubSpoutConfig config = new LogHubSpoutConfig(conumser_group_name,
        endpoint, project, logstore, access_id,
        access_key, LogHubCursorPosition.END_CURSOR);
    TopologyBuilder builder = new TopologyBuilder();
    // Create a LogHub Storm spout.
    LogHubSpout spout = new LogHubSpout(config);
    // You can create the same number of spouts as that of shards in a Logstore in actual business scenarios.
    builder.setSpout("spout", spout, 1);
    builder.setBolt("exclaim", new SampleBolt()).shuffleGrouping("spout");
    Config conf = new Config();
    conf.setDebug(false);
    conf.setMaxSpoutPending(1);
    // Configure the serialization method of LogGroupData by using the LogGroupDataSerializSerializer class if Kryo is used to serialize and deserialize data.
    Config.registerSerialization(conf, LogGroupData.class, LogGroupDataSerializSerializer.class);
    if (mode.equals("Local")) {
        logger.info("Local mode...");
        LocalCluster cluster = new LocalCluster();
        cluster.submitTopology("test-jstorm-spout", conf, builder.createTopology());
        try {
            Thread.sleep(6000 * 1000); //waiting for several minutes
        } catch (InterruptedException e) {
            // TODO Auto-generated catch block
            e.printStackTrace();
        }
        cluster.killTopology("test-jstorm-spout");
        cluster.shutdown();
    } else if (mode.equals("Remote")) {
        logger.info("Remote mode...");
        conf.setNumWorkers(2);
        try {
            StormSubmitter.submitTopology("stt-jstorm-spout-4", conf, builder.createTopology());
        } catch (AlreadyAliveException e) {
            // TODO Auto-generated catch block
            e.printStackTrace();
        } catch (InvalidTopologyException e) {
            // TODO Auto-generated catch block
            e.printStackTrace();
        }
    } else {
        logger.error("invalid mode: " + mode);
    }
}
}

```

- Use the following example code of bolts to consume log data and display the content of each log entry:

```

public class SampleBolt extends BaseRichBolt {
    private static final long serialVersionUID = 4752656887774402264L;
    private static final Logger logger = Logger.getLogger(BaseBasicBolt.class);
    private OutputCollector mCollector;
    @Override
    public void prepare(@SuppressWarnings("rawtypes") Map stormConf, TopologyContext context,
        OutputCollector collector) {
        mCollector = collector;
    }
    @Override
    public void execute(Tuple tuple) {
        String shardId = (String) tuple
            .getValueByField(LogHubSpout.FIELD_SHARD_ID);
        @SuppressWarnings("unchecked")
        List<LogGroupData> logGroupDatas = (ArrayList<LogGroupData>) tuple.getValueByField(LogHubSpout.FIELD_LOGGROUPS);
        for (LogGroupData groupData : logGroupDatas) {
            // Each log group consists of one or more log entries.
            LogGroup logGroup = groupData.getLogGroup();
            for (Log log : logGroup.getLogsList()) {
                StringBuilder sb = new StringBuilder();
                // Each log entry has a time field and other key-value pairs.
                int log_time = log.getTime();
                sb.append("LogTime:").append(log_time);
                for (Content content : log.getContentsList()) {
                    sb.append("\t").append(content.getKey()).append(":")
                        .append(content.getValue());
                }
                logger.info(sb.toString());
            }
        }
        // Spouts must use the ack method to indicate whether data has been successfully sent to bolts.
        // In addition, bolts must use the ack method to indicate whether data is successfully processed by the bolts.
        mCollector.ack(tuple);
    }
    @Override
    public void declareOutputFields(OutputFieldsDeclarer declarer) {
        //do nothing
    }
}

```

Maven

Use the following code to add Maven dependencies for Storm versions earlier than Storm 1.0 (for example, Storm 0.9.6):

```

<dependency>
  <groupId>com.aliyun.openservices</groupId>
  <artifactId>loghub-storm-spout</artifactId>
  <version>0.6.6</version>
</dependency>

```

Use the following code to add Maven dependencies for Storm 1.0 and later versions:

```
<dependency>
<groupId>com.aliyun.openservices</groupId>
<artifactId>loghub-storm-1.0-spout</artifactId>
<version>0.1.3</version>
</dependency>
```

23.6.5. Use Flume to consume log data

This topic describes how to use Flume to consume log data. You can use the aliyun-log-flume plug-in to connect LogHub of Log Service to Flume, and then write and consume log data.

The aliyun-log-flume plug-in connects LogHub to Flume. When LogHub is connected to Flume, you can connect Log Service to other systems such as HDFS and Kafka through Flume. The aliyun-log-flume plug-in provides the Sink and Source methods to connect Log Service to Flume.

- Sink: uses Flume to read data from other data sources and then writes data to LogHub.
- Source: uses Flume to consume LogHub data and then writes data to other systems.

LogHub Sink

You can use the Sink method to transmit data from other data sources to LogHub through Flume. Data can be parsed into the following two formats:

- SIMPLE: writes a Flume event to LogHub as a field.
- DELIMITED: delimits a Flume event with delimiters, parses an event into fields based on the configured column names, and then writes the fields to LogHub.

The following table lists the parameters you can configure when you use the Sink method to read data.

| Parameter | Required | Description |
|---------------|----------|--|
| type | Yes | Valid value: com.aliyun.loghub.flume.sink.LoghubSink. |
| endpoint | Yes | The endpoint of Log Service. |
| project | Yes | The name of the project. |
| logstore | Yes | The name of the Logstore. |
| accessKeyId | Yes | The AccessKey ID of your Apsara Stack tenant account. |
| accessKey | Yes | The AccessKey secret of your Apsara Stack tenant account. |
| batchSize | No | The number of log entries that are written to LogHub at a time. Default value: 1000. |
| maxBufferSize | No | The maximum size of the queue in the buffer. Default value: 1000. |

| Parameter | Required | Description |
|---------------|----------|--|
| serializer | No | The serialization format of log data. Valid values: <ul style="list-style-type: none"> DELIMITED: Data is parsed into the DELIMITED format. If you set this parameter to DELIMITED, you must set the columns parameter. SIMPLE: Data is parsed into the SIMPLE format. This is the default value. Custom serializer: Data is parsed into a custom serialization format. If you set this parameter to a custom serializer, you must specify the full name of the class. |
| columns | No | The configured column names. You must set this parameter if you set the serializer parameter to DELIMITED . Separate multiple columns with commas (,). Ensure that the columns are sorted in the same order as those of the log data. |
| separatorChar | No | The delimiter, which must be a single character. You can set this parameter if you set the serializer parameter to DELIMITED . Default value: <code>,</code> . |
| quoteChar | No | The quote character. You can set this parameter if you set the serializer parameter to DELIMITED . Default value: <code>"</code> . |
| escapeChar | No | The escape character. You can set this parameter if you set the serializer parameter to DELIMITED . Default value: <code>\</code> . |
| useRecordTime | No | Specifies whether to use the value of the timestamp field as the time when log data is written to Log Service. The default value false indicates that the current time is used. |

Loghub Source

You can use the Source method to ship data from LogHub to other data systems through Flume. Data can be output in the following two formats:

- **DELIMITED**: writes delimited log data to Flume.
- **JSON**: writes JSON-formatted log data to Flume.

The following table lists the parameters you can configure when you use the Source method to read data.

| Parameter | Required | Description |
|-------------|----------|---|
| type | Yes | Valid value: <code>com.aliyun.loghub.flume.source.LoghubSource</code> . |
| endpoint | Yes | The endpoint of Log Service. |
| project | Yes | The name of the project. |
| logstore | Yes | The name of the Logstore. |
| accessKeyId | Yes | The AccessKey ID of your Apsara Stack tenant account. |

| Parameter | Required | Description |
|---------------------|----------|--|
| accessKey | Yes | The AccessKey secret of your Apsara Stack tenant account. |
| heartbeatIntervalMs | No | The heartbeat interval between the Flume client and LogHub. Unit: milliseconds. Default value: 30000. |
| fetchIntervalMs | No | The interval for pulling data from LogHub. Unit: milliseconds. Default value: 100. |
| fetchInOrder | No | Specifies whether to consume log data in the order that log data was generated. Default value: false. |
| batchSize | No | The number of log entries that are read at a time. Default value: 100. |
| consumerGroup | No | The name of the consumer group that reads data. The name is randomly generated. |
| initialPosition | No | The start point from which data is read. Valid values: begin, end, and timestamp. Default value: begin.  Note If a checkpoint exists on the server, the checkpoint is used. |
| timestamp | No | The Unix timestamp. You must set this parameter if you set the initialPosition parameter to timestamp . Unix timestamp. |
| deserializer | Yes | The deserialization format of log data. Valid values: <ul style="list-style-type: none"> DELIMITED: Data is parsed into the DELIMITED format. This is the default value. If you set this parameter to DELIMITED, you must set the columns parameter. JSON: Data is parsed into the JSON format. Custom deserializer: Data is parsed into a custom deserialization format. If you set this parameter to a custom deserializer, you must specify the full name of the class. |
| columns | No | The configured column names. You must set this parameter if you set the deserializer parameter to DELIMITED . Separate multiple columns with commas (,). Ensure that the columns are sorted in the same order as those of the log data. |
| separatorChar | No | The delimiter, which must be a single character. You can set this parameter if you set the deserializer parameter to DELIMITED . Default value: <code>,</code> . |
| quoteChar | No | The quote character. You can set this parameter if you set the deserializer parameter to DELIMITED . Default value: <code>""</code> . |
| escapeChar | No | The escape character. You can set this parameter if you set the deserializer parameter to DELIMITED . Default value: <code>""</code> . |

| Parameter | Required | Description |
|-----------------|----------|--|
| appendTimestamp | No | Specifies whether to append the timestamp as a field to the end of each log entry. You can set this parameter if you set the deserializer parameter to DELIMITED . Default value: false. |
| sourceAsField | No | Specifies whether to add the log source as a field named <code>__source__</code> . You can set this parameter if you set the deserializer parameter to JSON . Default value: false. |
| tagAsField | No | Specifies whether to add the log tags as a field with the field name <code>__tag__: {tag names}</code> . You can set this parameter if you set the deserializer parameter to JSON . Default value: false. |
| timeAsField | No | Specifies whether to add the log time as a field named <code>__time__</code> . You can set this parameter if you set the deserializer parameter to JSON . Default value: false. |
| useRecordTime | No | Specifies whether to use the value of the timestamp field as the time when log data is read from Log Service. The default value false indicates that the current time is used. Default value: false. |

23.6.6. Use open source Flink to consume log data

Log Service provides the `flink-log-connector` plug-in to connect with Flink. This topic describes how to integrate the `flink-log-connector` plug-in with Flink to consume log data.

Prerequisites

- An `AccessKey` pair, a Log Service project, and a Logstore are created.
- If you log on to Log Service with a RAM user, relevant permissions to access a Logstore are granted to a RAM user. For more information, see [Grant permissions to a RAM role](#).

Context

The `flink-log-collector` plug-in includes `flink-log-consumer` and `flink-log-producer`.

- The `flink-log-consumer` plug-in reads data from Log Service. This plug-in supports the exactly-once semantics and load balancing among shards.
- The `flink-log-producer` plug-in writes data into Log Service. When you use the `flink-log-producer` plug-in, you must add the following Maven dependencies to a project:

```
<dependency>
  <groupId>com.aliyun.openservices</groupId>
  <artifactId>flink-log-connector</artifactId>
  <version>0.1.13</version>
</dependency>
<dependency>
  <groupId>com.google.protobuf</groupId>
  <artifactId>protobuf-java</artifactId>
  <version>2.5.0</version>
</dependency>
```

Log Consumer

The flink-log-consumer plug-in can consume the log data of a Logstore in Log Service. The exactly-once semantics is achieved during log consumption. The flink-log-consumer plug-in detects the change of the number of shards in a Logstore. This increases efficiency.

Each Flink subtask consumes data of some shards in a Logstore. If shards in a Logstore are split or merged, the shards consumed by the subtask also change.

When you use the flink-log-consumer plug-in to consume data from Log Service, you can call the following API operations:

- **GetCursorOrData**

You can call this operation to pull log data from a shard. Frequent API requests may exceed the read speed and IOPS limits of the shard. You can specify the `ConfigConstants.LOG_FETCH_DATA_INTERVAL_MILLIS` and `ConfigConstants.LOG_MAX_NUMBER_PER_FETCH` parameters to control the interval of API requests and number of log entries pulled in each request.

```
configProps.put(ConfigConstants.LOG_FETCH_DATA_INTERVAL_MILLIS, "100");
configProps.put(ConfigConstants.LOG_MAX_NUMBER_PER_FETCH, "100");
```

- **ListShards**

You can call this operation to view all shards in a Logstore and the status of each shard. If the shards are frequently split and merged, you can adjust the call interval to detect the changes in the number of shards.

```
// Call the ListShards operation once every 30 seconds.
configProps.put(ConfigConstants.LOG_SHARDS_DISCOVERY_INTERVAL_MILLIS, "30000");
```

- **CreateConsumerGroup**

You can call this operation to create a consumer group to synchronize checkpoints. This operation can be called only when consumption progress monitoring is enabled.

- **ConsumerGroupUpdateCheckPoint**

You can call this operation to synchronize snapshots of Flink to a consumer group.

The following table lists the Apsara Stack resources required for RAM users to call the preceding API operations.

| API | Alibaba Resource Name (ARN) |
|-------------------------------|--|
| GetCursorOrData | <code>acs:log:\${regionName}:\${projectOwnerAliUid}:project/\${projectName}/logstore/\${logstoreName}</code> |
| ListShards | <code>acs:log:\${regionName}:\${projectOwnerAliUid}:project/\${projectName}/logstore/\${logstoreName}</code> |
| CreateConsumerGroup | <code>acs:log:\${regionName}:\${projectOwnerAliUid}:project/\${projectName}/logstore/\${logstoreName}/consumergroup/*</code> |
| ConsumerGroupUpdateCheckPoint | <code>acs:log:\${regionName}:\${projectOwnerAliUid}:project/\${projectName}/logstore/\${logstoreName}/consumergroup/\${consumerGroupName}</code> |

For more information, see [Grant permissions to a RAM role](#).

1. Configure startup parameters.

The following example shows how to consume log data. The `java.util.Properties` class is used as the configuration tool. You can find all constants to be configured in the `ConfigConstants` class.

```

Properties configProps = new Properties();
// Specify the endpoint of Log Service.
configProps.put(ConfigConstants.LOG_ENDPOINT, "cn-hangzhou.log.aliyuncs.com");
// Specify the AccessKey pair.
configProps.put(ConfigConstants.LOG_ACCESSKEYID, "");
configProps.put(ConfigConstants.LOG_ACCESSKEY, "");
// Specify the project.
configProps.put(ConfigConstants.LOG_PROJECT, "ali-cn-hangzhou-sls-admin");
// Specify the Logstore.
configProps.put(ConfigConstants.LOG_LOGSTORE, "sls_consumergroup_log");
// Specify the start position to consume logs.
configProps.put(ConfigConstants.LOG_CONSUMER_BEGIN_POSITION, Consts.LOG_END_CURSOR);
// Specify the data deserialization method.
RawLogGroupListDeserializer deserializer = new RawLogGroupListDeserializer();
final StreamExecutionEnvironment env = StreamExecutionEnvironment.getExecutionEnvironment();
DataStream<RawLogGroupList> logTestStream = env.addSource(
    new FlinkLogConsumer<RawLogGroupList>(deserializer, configProps));

```

Note The number of subtasks in the Flink Streaming is independent of the number of shards in a Logstore. If the number of shards is greater than that of subtasks, each subtask consumes multiple shards exactly once. If the number of shards is less than that of subtasks, some subtasks are idle until new shards are generated.

2. Specify the start consumption position.

The `flink-log-consumer` plug-in enables you to specify the start consumption position of log data in a shard. By specifying the `ConfigConstants.LOG_CONSUMER_BEGIN_POSITION` parameter, you can start data consumption from the earliest, latest, or a specific time point. The `flink-log-connector` plug-in also allows a consumer group to resume consumption from a specific position. Valid values:

- `Consts.LOG_BEGIN_CURSOR`: The consumption starts from the earliest data.
- `Consts.LOG_END_CURSOR`: The consumption starts from the latest data.
- `Consts.LOG_FROM_CHECKPOINT`: The consumption starts from a checkpoint that is stored in a specific consumer group. You can use the `ConfigConstants.LOG_CONSUMERGROUP` parameter to specify the consumer group.
- `UnixTimestamp`: a string of the `INTEGER` data type. The timestamp is the number of seconds that have elapsed since 00:00:00 January 1, 1970. The value indicates that data in a shard is consumed from this time point.

You can use the following code to specify a start consumption position:

```

configProps.put(ConfigConstants.LOG_CONSUMER_BEGIN_POSITION, Consts.LOG_BEGIN_CURSOR);
configProps.put(ConfigConstants.LOG_CONSUMER_BEGIN_POSITION, Consts.LOG_END_CURSOR);
configProps.put(ConfigConstants.LOG_CONSUMER_BEGIN_POSITION, "1512439000");
configProps.put(ConfigConstants.LOG_CONSUMER_BEGIN_POSITION, Consts.LOG_FROM_CHECKPOINT);

```

Note If you have configured consumption resumption from a state backend of Flink when you start a Flink job, the `flink-log-connector` plug-in uses checkpoints stored in the state backend.

3. (Optional) Configure consumption progress monitoring.

The `flink-log-consumer` plug-in enables you to configure consumption progress monitoring. Consumption progress indicates the real-time consumption position of each shard. These positions are indicated by timestamps. For more information, see [View the status of a consumer group](#).

```

configProps.put(ConfigConstants.LOG_CONSUMERGROUP, "your consumer group name");

```

Note This configuration item is optional. If you specify this configuration item and no consumer group exists, the flink-log-consumer plug-in creates a consumer group. Snapshots in the flink-log-consumer plug-in are automatically synchronized to the consumer group of Log Service, and you can view the consumption progress of the flink-log-consumer plug-in in the Log Service console.

4. Configure consumption resumption and the exactly-once semantics.

If the checkpointing feature of Flink is enabled, the flink-log-consumer plug-in periodically stores the consumption progress of each shard. When a job fails, Flink restores the flink-log-consumer plug-in and starts to consume data from the latest checkpoint.

While you configure the checkpointing period, the maximum amount of data to be re-consumed when a failure occurs is defined. You can use the following code to configure the checkpointing period:

```
final StreamExecutionEnvironment env = StreamExecutionEnvironment.getExecutionEnvironment();
// Configure the exactly-once semantics.
env.getCheckpointConfig().setCheckpointingMode(CheckpointingMode.EXACTLY_ONCE);
// Store checkpoints every five seconds.
env.enableCheckpointing(5000);
```

For more information about the Flink checkpoints, see [Checkpoints](#) in the Flink documentation.

Log Producer

The flink-log-producer plug-in writes data into Log Service.

Note The flink-log-producer plug-in supports the Flink at-least-once semantics. If a job fails, data written into Log Service may be duplicated but never lost.

When you use the flink-log-producer plug-in to writes data to Log Service, you can call the following API operations:

- PostLogStoreLogs
- ListShards

The following table lists the Apsara Stack resources required for RAM users to call the preceding API operations.

| API | ARN |
|------------------|--|
| PostLogStoreLogs | <code>acs:log:\${regionName}:\${projectOwnerAliUid}:project/\${projectName}/logstore/\${logstoreName}</code> |
| ListShards | <code>acs:log:\${regionName}:\${projectOwnerAliUid}:project/\${projectName}/logstore/\${logstoreName}</code> |

For more information about RAM users and how to authorize RAM users, see [Grant permissions to a RAM role](#).

1. Initialize the flink-log-producer plug-in.

- i. Configure startup parameters for the flink-log-producer plug-in.

The initialization process for the flink-log-producer plug-in is similar to that for the flink-log-consumer plug-in. The following code shows the available parameters that you can configure for the flink-log-producer plug-in. You can use the default values of these parameters. You can also specify the parameters based on your needs.

```
// The number of I/O threads used to send data. Default value: 8.
ConfigConstants.LOG_SENDER_IO_THREAD_COUNT
// The time it takes to send the data after log data is cached. Default value: 3000.
ConfigConstants.LOG_PACKAGE_TIMEOUT_MILLIS
// The number of logs in the cached package. Default value: 4096.
ConfigConstants.LOG_LOGS_COUNT_PER_PACKAGE
// The size of the cached package. Default value: 3 Mbits.
ConfigConstants.LOG_LOGS_BYTES_PER_PACKAGE
// The total memory size that the job can use. Default value: 100 Mbits.
ConfigConstants.LOG_MEM_POOL_BYTES
```

 **Note** These parameters are optional. You can use their default values.

- ii. Reload LogSerializationSchema to define the method of serializing data into RawLogGroup.

To use the hash key to specify the shard for data writes, you can use the LogPartitioner method to generate the hash key for the data.

Example:

```
FlinkLogProducer<String> logProducer = new FlinkLogProducer<String>(new SimpleLogSerializer(), configProps);
logProducer.setCustomPartitioner(new LogPartitioner<String>() {
    // Generate a 32-bit hash value.
    public String getHashKey(String element) {
        try {
            MessageDigest md = MessageDigest.getInstance("MD5");
            md.update(element.getBytes());
            String hash = new BigInteger(1, md.digest()).toString(16);
            while(hash.length() < 32) hash = "0" + hash;
            return hash;
        } catch (NoSuchAlgorithmException e) {
        }
        return "0000000000000000000000000000000000000000000000000000000000000000";
    }
});
```

 **Note** The LogPartitioner method is optional. If you do not configure this method, data is randomly written into a shard.

2. Run the following code and write the generated strings to Log Service.

```
// Serialize data into the format of raw log groups.
class SimpleLogSerializer implements LogSerializationSchema<String> {
    public RawLogGroup serialize(String element) {
        RawLogGroup rlg = new RawLogGroup();
        RawLog rl = new RawLog();
        rl.setTime((int)(System.currentTimeMillis() / 1000));
        rl.addContent("message", element);
        rlg.addLog(rl);
        return rlg;
    }
}

public class ProducerSample {
    public static String sEndpoint = "cn-hangzhou.log.aliyuncs.com";
    public static String sAccessKeyId = "";
    public static String sAccessKey = "";
    public static String sProject = "ali-cn-hangzhou-sls-admin";
    public static String sLogstore = "test-flink-producer";
    private static final Logger LOG = LoggerFactory.getLogger(ConsumerSample.class);
    public static void main(String[] args) throws Exception {
        final ParameterTool params = ParameterTool.fromArgs(args);
        final StreamExecutionEnvironment env = StreamExecutionEnvironment.getExecutionEnvironment();
        env.getConfig().setGlobalJobParameters(params);
        env.setParallelism(3);
        DataStream<String> simpleStringStream = env.addSource(new EventsGenerator());
        Properties configProps = new Properties();
        // Specify the endpoint of Log Service.
        configProps.put(ConfigConstants.LOG_ENDPOINT, sEndpoint);
        // Specify the AccessKey pair to access Log Service.
        configProps.put(ConfigConstants.LOG_ACCESSKEYID, sAccessKeyId);
        configProps.put(ConfigConstants.LOG_ACCESSKEY, sAccessKey);
        // Specify the project to which logs are written.
        configProps.put(ConfigConstants.LOG_PROJECT, sProject);
        // Specify the Logstore to which logs are written.
        configProps.put(ConfigConstants.LOG_LOGSTORE, sLogstore);
        FlinkLogProducer<String> logProducer = new FlinkLogProducer<String>(new SimpleLogSerializer(), configProps);
    };
    simpleStringStream.addSink(logProducer);
    env.execute("flink log producer");
}

// Simulate log generation.
public static class EventsGenerator implements SourceFunction<String> {
    private boolean running = true;
    @Override
    public void run(SourceContext<String> ctx) throws Exception {
        long seq = 0;
        while (running) {
            Thread.sleep(10);
            ctx.collect((seq++) + "-" + RandomStringUtils.randomAlphabetic(12));
        }
    }
    @Override
    public void cancel() {
        running = false;
    }
}
}
```

23.6.7. Use Logstash to consume log data

Log Service allows you to use SDKs developed in various languages and various stream computing systems to consume log data. In addition, Log Service allows you to use Logstash to consume log data. You can configure the Logstash input plug-in to read log data from Log Service and write the data to other systems, such as Kafka and Hadoop Distributed File System (HDFS).

Features

- **Distributed collaborative consumption:** Multiple servers can be configured to consume log data from a Logstore at the same time.
- **High performance:** If you use consumer groups written in Java to consume log data, the consumption speed of a CPU core can reach up to 20 MB/s.
- **High reliability:** Log Service saves the consumption progress. This mechanism enables automatic resumption of log consumption from the last checkpoint after a consumption exception is solved.
- **Automatic load balancing:** Shards are automatically allocated based on the number of consumers in a consumer group. The shards are reallocated if consumers join or leave the consumer group.

23.6.8. Use Spark Streaming to consume log data

This topic describes how to use Spark Streaming to consume log data. After logs are collected to Log Service, you can use the Spark SDK provided by Log Service to process log data in Spark Streaming.

The Spark SDK supports two consumption modes: Receiver and Direct.

The Maven dependency is as follows:

```
<dependency>
<groupId>com.aliyun.emr</groupId>
<artifactId>emr-logservice_2.11</artifactId>
<version>1.7.2</version>
</dependency>
```

Receiver mode

In the Receiver mode, a consumer group consumes data from Log Service and temporarily stores the data in Spark Executor. After a Spark Streaming job is started, it reads and processes data from Spark Executor. For more information, see [Use consumer groups to consume log data](#). Each log entry is returned as a JSON string. The consumer group periodically saves checkpoints to Log Service. You do not need to update checkpoints.

- Example code

```

import org.apache.spark.storage.StorageLevel
import org.apache.spark.streaming.aliyun.logservice.LoghubUtils
import org.apache.spark.streaming.{ Milliseconds, StreamingContext}
import org.apache.spark.SparkConf
object TestLoghub {
  def main(args: Array[String]): Unit = {
    if (args.length < 7) {
      System.err.println(
        """Usage: TestLoghub <project> <logstore> <loghub group name> <endpoint>
        |   <access key id> <access key secret> <batch interval seconds>
        """
        .stripMargin)
      System.exit(1)
    }
    val project = args(0)
    val logstore = args(1)
    val consumerGroup = args(2)
    val endpoint = args(3)
    val accessKeyId = args(4)
    val accessKeySecret = args(5)
    val batchInterval = Milliseconds(args(6).toInt * 1000)
    def functionToCreateContext(): StreamingContext = {
      val conf = new SparkConf().setAppName("Test Loghub")
      val ssc = new StreamingContext(conf, batchInterval)
      val loghubStream = LoghubUtils.createStream(
        ssc,
        project,
        logstore,
        consumerGroup,
        endpoint,
        accessKeyId,
        accessKeySecret,
        StorageLevel.MEMORY_AND_DISK)
      loghubStream.checkpoint(batchInterval * 2).foreachRDD(rdd =>
        rdd.map(bytes => new String(bytes)).top(10).foreach(println)
      )
      ssc.checkpoint("hdfs://tmp/spark/streaming") // set checkpoint directory
      ssc
    }
    val ssc = StreamingContext.getOrCreate("hdfs://tmp/spark/streaming", functionToCreateContext _)
    ssc.start()
    ssc.awaitTermination()
  }
}

```

- Parameter description

| Parameter | Type | Description |
|---------------|--------|--|
| project | String | The project in Log Service. |
| logstore | String | The Logstore in Log Service. |
| consumerGroup | String | The name of the consumer group. |
| endpoint | String | The endpoint of the region to which the project belongs. |
| accessKeyId | String | The AccessKey ID used to access Log Service. |

| Parameter | Type | Description |
|-----------------|--------|--|
| accessKeySecret | String | The AccessKey secret used to access Log Service. |

- **Notes**

In the Receiver mode, data loss may occur in some cases. To avoid data loss, you can turn on the Write-Ahead Logs switch, which is supported in Spark 1.2 and later versions. For more information, visit [Spark Streaming Programming Guide](#).

Direct mode

In the Direct mode, no consumer group is required. API operations are called to request data from Log Service. Compared with the Receiver mode, the Direct mode has the following benefits:

- Simplified parallelism. The number of Spark partitions is the same as the number of shards in a Logstore. You can split shards to improve the parallelism of tasks.
- Increased efficiency. You no longer need to turn on the Write-Ahead Logs switch to prevent data loss.
- Exactly-once semantics. Data is read directly from Log Service. Checkpoints are submitted after the task is successful. In some cases, data may be repeatedly consumed if the task is not ended due to unexpected exit of Spark.

You must configure the ZooKeeper service when you use the Direct mode. You must set a checkpoint directory in ZooKeeper to store intermediate state data. If you want to re-consume data after restarting a task, you must delete the directory from ZooKeeper and change the name of the consumer group.

- Example code

```

import com.aliyun.openservices.loghub.client.config.LogHubCursorPosition
import org.apache.spark.SparkConf
import org.apache.spark.streaming.{ Milliseconds, StreamingContext}
import org.apache.spark.streaming.aliyun.logservice.{ CanCommitOffsets, LoghubUtils}
object TestDirectLoghub {
  def main(args: Array[String]): Unit = {
    if (args.length < 7) {
      System.err.println(
        """Usage: TestDirectLoghub <project> <logstore> <loghub group name> <endpoint>
        | <access key id> <access key secret> <batch interval seconds> <zookeeper host:port=localhost:2181>
        """
      ).stripMargin
      System.exit(1)
    }
    val project = args(0)
    val logstore = args(1)
    val consumerGroup = args(2)
    val endpoint = args(3)
    val accessKeyId = args(4)
    val accessKeySecret = args(5)
    val batchSize = Milliseconds(args(6).toInt * 1000)
    val zkAddress = if (args.length >= 8) args(7) else "localhost:2181"
    def functionToCreateContext(): StreamingContext = {
      val conf = new SparkConf().setAppName("Test Direct Loghub")
      val ssc = new StreamingContext(conf, batchSize)
      val zkParas = Map("zookeeper.connect" -> zkAddress,
        "enable.auto.commit" -> "false")
      val loghubStream = LoghubUtils.createDirectStream(
        ssc,
        project,
        logStore,
        consumerGroup,
        accessKeyId,
        accessKeySecret,
        endpoint,
        zkParas,
        LogHubCursorPosition.END_CURSOR)
      loghubStream.checkpoint(batchSize).foreachRDD(rdd => {
        println(s"count by key: ${rdd.map(s => {
          s.sorted
          (s.length, s)
        }).countByKey().size}")
        loghubStream.asInstanceOf[CanCommitOffsets].commitAsync()
      })
      ssc.checkpoint("hdfs://tmp/spark/streaming") // set checkpoint directory
      ssc
    }
    val ssc = StreamingContext.getOrCreate("hdfs://tmp/spark/streaming", functionToCreateContext _)
    ssc.start()
    ssc.awaitTermination()
  }
}

```

- Parameter description

| Parameter | Type | Description |
|-----------|--------|------------------------------|
| project | String | The project in Log Service. |
| logstore | String | The Logstore in Log Service. |

| Parameter | Type | Description |
|-----------------|--------|---|
| consumerGroup | String | The name of the consumer group (only used to save consumption checkpoints). |
| endpoint | String | The endpoint of the region to which the project belongs. |
| accessKeyld | String | The AccessKey ID used to access Log Service. |
| accessKeySecret | String | The AccessKey secret used to access Log Service. |
| zkAddress | String | The endpoint of ZooKeeper. |

- Parameter settings

In the Direct mode, you must specify the number of log entries that are consumed in each shard in each batch. Otherwise, the data reading process cannot be ended. You can throttle the transmission rate of a single batch by setting the two parameters listed in the following table.

| Parameter | Description | Default value |
|--|---|---------------|
| spark.loghub.batchGet.step | The number of log groups returned for a single request. | 100 |
| spark.streaming.loghub.maxRatePerShard | The maximum number of log entries that are read in a shard. | 10,000 |

The number of log entries processed in each batch is calculated as follows: $\text{number of shards} \times \max(\text{spark.loghub.batchGet.step} \times n \times \text{number of log entries in a log group}, \text{spark.streaming.loghub.maxRatePerShard} \times \text{duration})$.

- o n : the number of requests required to increase the returned rows to $\text{spark.streaming.loghub.maxRatePerShard} \times \text{duration}$.
- o duration: the interval between batch processing. Unit: milliseconds.

If you need to combine multiple data streams, the number of shards refers to the total number of shards in all Logstores.

- o Example

For example, the number of shards is 100. Each log group contains 50 log entries. Batches are processed at an interval of two seconds. If you want to process 20,000 log entries in each batch, use the following configurations:

- `spark.loghub.batchGet.step: 4`
- `spark.streaming.loghub.maxRatePerShard: 100`

If each log group contains 60 log entries and you want to process 20,000 log entries in each batch, 24,000 log entries will be processed based on the preceding configurations ($60 \times 4 \times 100 = 24,000$).

- o Accurate transmission rate throttling

A smaller `spark.loghub.batchGet.step` value increases the accuracy of throttling and the number of requests. We recommend that you count the average number of log entries contained in a log group and then set the preceding two parameters.

23.6.9. Use Realtime Compute to consume log data

You can use Realtime Compute (Blink) to create a schema for Log Service data and consume the data.

Log Service stores streaming data. Therefore, Realtime Compute can use the streaming data as input data. In Log Service, each log entry contains multiple fields and each field is a key-value pair. The following example is a sample log entry:

```
__source__: 11.85.123.199
__tag__:__receive_time__: 1562125591
__topic__: test-topic
a: 1234
b: 0
c: hello
```

You can use the following data definition language (DDL) statement to create a table in Realtime Compute:

```
create table sls_stream(
  a int,
  b int,
  c varchar
) with (
  type ='sls',
  endPoint ='<your endpoint>',
  accessId ='<your access key id>',
  accessKey ='<your access key>',
  startTime = '2017-07-05 00:00:00',
  project ='ali-cloud-streamtest',
  logStore ='stream-test',
  consumerGroup ='consumerGroupTest1'
);
```

Attribute fields

Realtime Compute can extract fields from log content. In addition, Realtime Compute can extract three attribute fields and custom tag fields, such as `__receive_time__`. The following table lists the three attribute fields.

Attribute fields

| Field name | Description |
|----------------------------|---|
| <code>__source__</code> | The source of the log entry. |
| <code>__topic__</code> | The topic of the log entry. |
| <code>__timestamp__</code> | The time when the log entry is generated. |

To extract the preceding fields, you must add HEADERS in the DDL statement. For example:

```

create table sls_stream(
  __timestamp__ bigint HEADER,
  __receive_time__ bigint HEADER
  b int,
  c varchar
) with (
  type='sls',
  endPoint='<your endpoint>',
  accessId='<your access key id>',
  accessKey='<your access key>',
  startTime='2017-07-05 00:00:00',
  project='ali-cloud-streamtest',
  logStore='stream-test',
  consumerGroup='consumerGroupTest1'
);

```

Parameters in the WITH clause

The following table describes the parameters in the WITH clause.

| Parameter | Required | Description |
|------------------------|----------|---|
| endPoint | Yes | The endpoint of Log Service. |
| accessId | Yes | The AccessKey ID of the Apsara Stack tenant account or RAM user that is used to access Log Service. |
| accessKey | Yes | The AccessKey secret of the Apsara Stack tenant account or RAM user that is used to access Log Service. |
| project | Yes | The name of the project in Log Service. |
| logStore | Yes | The name of the Logstore in Log Service. |
| consumerGroup | No | The name of the consumer group. |
| startTime | No | The time when Realtime Compute starts to consume data. |
| heartBeatIntervalMills | No | The heartbeat interval of the client that consumes log data. Unit: seconds. Default value: 10. |
| maxRetryTimes | No | The maximum number of retries to read data. Default value: 5. |
| batchGetSize | No | The number of log groups that are read at a time. Default value: 10. If the version of Blink is 1.4.2 or later, the default value is 100 and the maximum value is 1000. |

| Parameter | Required | Description |
|------------------|----------|---|
| columnErrorDebug | No | Specifies whether to enable debugging. If debugging is enabled, log entries that fail to be parsed are displayed. Default value: <code>false</code> . |

Field type mapping

All log fields in Log Service are of the string type. The following table lists the mapping between the type of Log Service fields and the type of Realtime Compute fields. We recommend that you declare the mapping in a data definition language (DDL) statement.

| Data type of Log Service | Data type of Realtime Compute |
|--------------------------|-------------------------------|
| STRING | VARCHAR |

If you specify another data type to convert Log Service data, an automatic conversion attempt is performed. For example, you can specify *BIGINT* as the data type to convert the string "1000" and specify *timestamp* as the data type to convert the string "2018-01-12 12:00:00".

Note

- Blink versions earlier than 2.2.0 do not support shard scaling. If you split or merge shards when a job is reading data from a Logstore, the job fails and cannot continue. In this case, you must restart the job.
- No versions of Blink allow you to delete or recreate a Logstore whose log data is being consumed.
- For Blink version 1.6.0 and earlier, if you specify *consumerGroup* to consume log data from a Logstore that contains a large number of shards, the read performance may be affected.
- When you create a schema, Log Service data cannot be converted to data of the map type.
- Fields whose values are empty are set to *null*.
- Unordered field conversions are supported. However, we recommend that you convert the fields in the same order as the fields in the schema.
- The *batchGetSize* parameter specifies the number of log groups that are obtained based on the *logGroup* parameter. If the size of each log entry and the value of the *batchGetSize* parameter are both large, garbage collection (GC) of data in the memory may frequently occur.

Precautions

- If no new data is written to a shard, the overall latency of a job increases. In this case, you need to modify the number of concurrent tasks in the job to the number of shards from which data is read and written.
- We recommend that you set the number of concurrent tasks in a job to the same as the number of shards. Otherwise, data may be filtered out when the job reads historical data from two shards at significantly different speeds.
- To extract fields in tags such as `__tag__:__hostname__` and `__tag__:__path__`, you can delete the `__tag__:` prefix and use the method of extracting attribute fields.

 **Note** This type of data cannot be extracted during debugging. We recommend that you use the local debugging method and the print method to display data in logs.

23.7. Data shipping

23.7.1. Ship logs to OSS

23.7.1.1. Overview

Log Service provides the data shipping feature. You can use this feature to ship logs to Object Storage Service (OSS) in real time by using the Log Service console. This topic describes the benefits and scenarios of the data shipping feature.

In the Log Service console, you can ship logs to other Apsara Stack services. Then, you can store or consume the log data by using other systems such as E-MapReduce. After you enable the log shipping feature, Log Service ships the collected logs to the specified cloud service at regular intervals.

Scenarios

The data shipping feature can be used to connect Log Service with data warehouses.

Benefits

The data shipping feature of Log Service has the following benefits:

- **Ease of use**
You only need to complete a few settings in the Log Service console before you can ship logs from Logstores to other Apsara Stack services such as OSS.
- **High efficiency**
Log Service stores logs that are collected from multiple servers. This improves efficiency when you ship log data to Apsara Stack services such as OSS.
- **Effective management**
You can ship logs from different projects or Logstores to different OSS buckets. This way, you can efficiently manage the logs by log type or log source.

Log shipping destinations

For information about how to ship logs to OSS, see [Ship log data from Log Service to OSS](#).

23.7.1.2. Ship log data from Log Service to OSS

You can use Log Service to collect log data and ship the log data to Object Storage Service (OSS) for storage and analysis. This topic describes how to ship log data from Log Service to OSS.

Prerequisites

- Log data is collected. For more information, see [Log collection methods](#).
- OSS is activated. A bucket is created in the region where the Log Service project resides. For more information, see the **Create bucket**s section in the *Service User Guide - Object Storage Service (OSS)*.
- A Resource Access Management (RAM) role is created for the level-1 organization. For more information, see [Obtain the ARN of a RAM role](#).

Context

Log Service can automatically ship log data from a Logstore to an OSS bucket.

- You can set a custom retention period for the log data in the OSS bucket. Permanent retention is supported.
- You can use data processing platforms such as E-MapReduce and Data Lake Analytics (DLA) or use custom programs to consume log data from the OSS bucket.

Procedure

1. [Log on to the Log Service console](#).
2. In the Projects section, click the project from which you want to ship log data to OSS.

3. On the Logstores tab, click the > icon on the left of the specific Logstore and choose **Data Transformation > Export > Object Storage Service(OSS)**.
4. On the OSS Shipper page, click **Enable**.
5. In the **OSS LogShipper** pane, configure the shipping rules.

The following table describes the required parameters.

| Parameter | Description |
|---------------------------------------|---|
| OSS Shipper Name | The name of the shipping rule. The name can contain only lowercase letters, digits, hyphens (-), and underscores (_). The name must start and end with a lowercase letter or digit and must be 2 to 128 characters in length. |
| OSS Bucket | The name of the OSS bucket to which you want to ship log data.  Notice You must specify the name of an existing OSS bucket. The specified OSS bucket must reside in the same region as the Log Service project. |
| OSS Prefix | The directory to which log data is shipped in the OSS bucket. |
| Partition format | The partition format of the bucket directory for the shipping task. The directory is automatically generated based on the time when the shipping task is created. The default format is %Y/%m/%d/%H/%M. The partition format cannot start with a forward slash (/). For information about partition format examples, see Partition format . For more information about parameters, see strptime API . |
| (Resource Access Management) RAM Role | The Alibaba Cloud Resource Name (ARN) of the RAM role. The RAM role is the identity that the OSS bucket owner creates for access control. Example: acs:ram::45643:role/aliyunlogdefaultrole. For information about how to obtain the ARN, see Obtain the ARN of a RAM role . |
| Shipping Size | The maximum size of raw log data that can be shipped to the OSS bucket in a shipping task. Valid values: 5 to 256. Unit: MB. If the size of shipped data exceeds the specified value, a new shipping task is automatically created. |
| Compress | Specifies whether to compress log data that is shipped to OSS. Valid values: <ul style="list-style-type: none"> ◦ No Compress: The log data that is shipped to OSS is not compressed. ◦ Compress (snappy): The snappy utility is used to compress the log data that is shipped to OSS. This way, the log data occupies less storage space of the OSS bucket. |
| Storage Format | The storage format of the log data that is shipped to OSS. Valid values: JSON, CSV, and Parquet. For more information, see Storage Formats . |
| Ship Tags | Specifies whether to ship log tags. |
| Shipping Time | The time period during which a shipping task runs. Valid values: 300 to 900. Default value: 300. Unit: seconds. If the specified time period expires, another shipping task is created. |

6. Click **OK**.

Note

- After you configure a shipping rule, multiple shipping tasks can concurrently run. If the size of the data shipped from a shard reaches the specified threshold or the specified time period expires, another task is created.
- After you create a shipping task, you can check whether the shipping rule satisfies your business requirements based on the task status and the data shipped to OSS.

View OSS data

After log data is shipped to OSS, you can access the log data in the OSS console, or by using the OSS API, an SDK, or another method. For more information, see the **Objects > Search for objects** section of the *Service User Guide - Object Storage Service(OSS)*.

The following script shows a sample OSS directory:

```
oss://OSS-BUCKET/OSS-PREFIX/PARTITION-FORMAT_RANDOM-ID
```

OSS-BUCKET is the name of the OSS bucket. OSS-PREFIX is the prefix of the directory in the OSS bucket. PARTITION-FORMAT is the partition format of the directory for a shipping task. The partition format is calculated based on the time when the shipping task is created. For more information, see [strptime API](#). RANDOM-ID is the unique identifier of the shipping task.

Note The directory in the OSS bucket is created based on the time when the shipping task is created. For example, the shipping task is created at 00:00:00 on June 23, 2016 to ship data to OSS. The data is written to Log Service after 23:55:00 on June 22, 2016. The shipping interval is 5 minutes. To retrieve all logs shipped on June 22, 2016, you must check all objects in the *2016/06/22* directory. You must also check the *2016/06/23/00/* directory for the objects that are generated in the first 10 minutes of June 23, 2020.

Partition format

For each shipping task, log data is written to a directory of an OSS bucket. The directory is in the *oss://OSS-BUCKET/OSS-PREFIX/PARTITION-FORMAT_RANDOM-ID* format. A partition format is obtained by formatting the time when a shipping task is created. The following table describes the partition formats and directories that are obtained when a shipping task is created at 19:50:43 on January 20, 2017.

| OSS Bucket | OSS Prefix | Partition format | OSS directory |
|-------------|----------------------|----------------------------------|--|
| test-bucket | test-table | %Y/%m/%d/%H/%M | <i>oss://test-bucket/test-table/2017/01/20/19/50_1484913043351525351_2850008</i> |
| test-bucket | log_ship_oss_example | year=%Y/mon=%m/day=%d/log_%H%M%s | <i>oss://test-bucket/log_ship_oss_example/year=2017/mon=01/day=20/log_195043_1484913043351525351_2850008.parquet</i> |
| test-bucket | log_ship_oss_example | ds=%Y%m%d/%H | <i>oss://test-bucket/log_ship_oss_example/ds=20170120/19_1484913043351525351_2850008.snappy</i> |

| OSS Bucket | OSS Prefix | Partition format | OSS directory |
|-------------|----------------------|------------------|--|
| test-bucket | log_ship_oss_example | %Y%m%d/ | <code>oss://test-bucket/log_ship_oss_example/20170120/_1484913043351525351_2850008</code> <div style="border: 1px solid #ccc; background-color: #e0f2f1; padding: 5px; margin-top: 10px;"> <p> Note This format may prevent platforms such as Hive from parsing the log data in the OSS bucket. We recommend that you do not use this format.</p> </div> |
| test-bucket | log_ship_oss_example | %Y%m%d%H | <code>oss://test-bucket/log_ship_oss_example/2017012019_1484913043351525351_2850008</code> |

You can use Hive, MaxCompute, or Data Lake Analytics (DLA) to analyze OSS data. In this case, if you want to use partition information, you can set PARTITION-FORMAT in the key=value format. For example, you can set the partition format to `oss://test-bucket/log_ship_oss_example/year=2017/mon=01/day=20/log_195043_1484913043351525351_2850008.parquet`. In this example, year, mon, and day are specified as three partition keys.

What to do next

After shipping tasks are created based on a shipping rule, you can modify the shipping rule. You can also disable the data shipping feature, view the statuses and error messages of the tasks, and retry failed tasks on the **OSS Shipper** page of a Logstore.

- **Modify the shipping rule.**
Click **Settings** to modify the shipping rule. For information about the parameters, see [Procedure](#).
- **Disable the data shipping feature.**
Click **Disable**. The data in the Logstore is no longer shipped to OSS.
- **View the statuses and error messages of the tasks.**
You can view the log shipping tasks of the last two days and their statuses.
 - **Statuses of a shipping task**

| Status | Equivalent |
|-----------|---|
| Succeeded | The shipping task has succeeded. |
| Running | The shipping task is running. Check whether the task succeeds later. |
| Failed | The shipping task has failed. If the task cannot be restarted due to external causes, troubleshoot the failure based on the error message and retry the task. |

o Error messages

If a shipping task fails, an error message is returned for the task.

| Error message | Error cause | Solution |
|---------------------|--|---|
| Unauthorized | The error message returned because the AliyunLogDefaultRole role does not have the required permissions. | To fix the error, check the following configurations: <ul style="list-style-type: none"> ▪ Check whether the AliyunLogDefaultRole role is created by the OSS bucket owner. ▪ Check whether the specified ID of the Alibaba Cloud account in the permission policy is valid. ▪ Check whether the AliyunLogDefaultRole role is granted the write permissions on the OSS bucket. ▪ Check whether the ARN of the AliyunLogDefaultRole role that you entered in the RAM Role field is valid. |
| ConfigNotExist | The error message returned because the task does not exist. | Check whether the data shipping feature is disabled. If the feature is disabled, enable the feature, configure a shipping rule, and then retry the shipping task. |
| InvalidOssBucket | The error message returned because the specified OSS bucket does not exist. | To fix the error, check the following configurations: <ul style="list-style-type: none"> ▪ Check whether the OSS bucket resides in the same region as the Log Service project. ▪ Check whether the specified bucket name is valid. |
| InternalServerError | The error message returned because an internal error has occurred in Log Service. | Retry the failed shipping task. |

o Retry a shipping task

By default, if a shipping task fails, Log Service retries the task based on the retry policy. You can also manually retry the task. By default, Log Service retries all tasks of the last two days. The minimum interval between two consecutive retries is 15 minutes. If a task fails for the first time, Log Service retries the task 15 minutes later. If the task fails for the second time, Log Service retries the task 30 minutes later. If the task fails for the third time, Log Service retries the task 60 minutes later. A similar method is used for subsequent attempts.

To immediately retry a failed task, you can click **Retry All Failed Tasks** or **Retry** on the right of the task. You can also use the Log Service API or an SDK to retry a task.

23.7.1.3. Obtain the ARN of a RAM role

When you use a RAM user to ship data from Log Service to Object Storage Service (OSS), you must first create a Resource Access Management (RAM) role and specify the ARN of the RAM role. This topic describes how to create a RAM role and obtain the ARN of a RAM role.

Procedure

1. [Log on to the Log Service console](#).
2. In the top navigation bar, click **Configurations**.

3. On the **Service-Linked Roles** page, click **Create RAM Role**.
4. In the **Organization Name** drop-down list, select the organization that you created. In the **Service Name** drop-down list, select **Log Service**, and click **OK**.
5. On the **RAM Service Role** page, enter `AliyunLogDefaultRole` in the **Role Name** search box and click **Search**.
6. Obtain the ARN of the RAM role.

In the search results, the value in the **role identifier** column is the ARN of the RAM role.

23.7.1.4. Storage Formats

Different storage formats are supported when Log Service ships logs to OSS, including JSON, CSV, and Parquet. This topic describes the field details of the formats.

JSON format

You can set the storage format for the data that is shipped to OSS. The following table shows how to set the **storage format** to **JSON**. For more information, see [Configure a data shipping rule](#).

| Compression type | File extension | Example file address | Description |
|------------------|----------------|---|---|
| Uncompressed | N/A | oss://oss-shipper-shenzhen/ecs_test/2016/01/26/20/54_1453812893059571256_937 | <p>You can download the raw JSON object to the local host and open each object as a text file. The following example is the content of a sample file:</p> <pre> {"__time__":1453809242,"__topic__":",","__source__":"10.170.**.***","ip":"10.200.**.**","time":"26/Jan/2016:19:54:02+0800","url":"POST/PutData?Category=YunOsAccountOpLog&AccessKeyId=<yourAccessKeyId>&Date=Fri%2C%2028%20Jun%202013%2006%3A53%3A30%20GMT&Topic=raw&Signature=<yourSignature> HTTP/1.1","status":"200","user-agent":"aliyun-sdk-java"} </pre> |
| snappy | .snappy | oss://oss-shipper-shenzhen/ecs_test/2016/01/26/20/54_1453812893059571256_937.snappy | JSON objects are compressed by using Snappy. For more information, see Decompression tools for Snappy-compressed files . |

CSV-format

You can set the storage format for the data that is shipped to OSS. The following table shows how to set the **storage format** to CSV. For more information, see [Configure a data shipping rule](#).

The following table describes the parameters. For more information, see [Common Format and MIME Type for Comma-Separated Values \(CSV\) Files](#) and [PostgreSQL 9.4.26 Documentation](#).

| Parameter | Description |
|------------------|--|
| CSV Fields | <p>The names of the log fields that you want to ship to OSS. You can view log fields on the Raw Logs tab of a Logstore and enter the names of the fields that you want to ship to OSS in the Key Name column. The log fields that you can ship to OSS include the fields in the log content and the reserved fields such as <code>__time__</code>, <code>__topic__</code>, and <code>__source__</code>.</p> <div style="border: 1px solid #ccc; background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p> Note The keys that you enter in the CSV Fields section must be unique.</p> </div> |
| Delimiter | You can use commas (,), vertical bars (), spaces, or tabs to delimit fields. |
| Escape Character | If a field contains a delimiter, you must use an escape character to enclose the field. This ensures that the field is not delimited. |
| Invalid Fields | If a key that you specify in the CSV Fields section does not exist, enter the value of the key in the Invalid Fields field. |
| Shipped Fields | If you turn on the Shipped Fields switch, field names are written in a CSV file. |

The following table lists the directories in OSS buckets that store the data shipped from Log Service.

| Compression type | File extension | Example | Description |
|------------------|----------------|---|---|
| No | .csv | oss://oss-shipper-shenzhen/ecs_test/2016/01/26/20154_1453812893059571256_937.csv | You can download the raw JSON object to the local host and open the object as a text file. |
| snappy | .snappy.csv | oss://oss-shipper-shenzhen/ecs_test/2016/01/26/20154_1453812893059571256_937.snappy.csv | Decompression tools for Snappy compressed files For more information, see Decompression tools for Snappy-compressed files . |

Parquet-format

You can set the storage format for the data that is shipped to OSS. The following figure shows how to set the **storage format** to Parquet. For more information, see [Configure a data shipping rule](#).

The following table describes the related parameters.

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

| Parameter | Description |
|-----------|---|
| Key Name | <p>The name of the log field that you want to ship to OSS. You can view log fields on the Raw Logs tab of a Logstore. You can also enter the names of the fields that you want to ship to OSS in the Key Name column. When the fields are shipped to OSS, they are stored in the Parquet format in the order that the field names are entered. The names of the fields are the column names in OSS. The log fields that you can ship to OSS include the fields in the log content and the reserved fields such as <code>__time__</code>, <code>__topic__</code>, and <code>__source__</code>. The value of a field in the Parquet format is null in the following two scenarios:</p> <ul style="list-style-type: none"> The field does not exist in logs. The value of the field fails to be converted from the string type to a non-string type, for example, double or Int64. <div style="background-color: #e6f2ff; padding: 5px; border: 1px solid #d9e1f2;"> <p> Note The keys that you enter in the Parquet Keys field must be unique.</p> </div> |
| Type | <p>The Parquet storage format supports six data types: string, Boolean, Int32, Int64, float, and double.</p> <p>Log fields are converted from the string type to a data type that the Parquet storage format supports. If the data type of a log field fails to be converted, the value of the log field is null.</p> |

The following table lists the directories in OSS buckets that store data shipped from Log Service.

| Compression type | File extension | Example | Description |
|------------------|-----------------|---|---|
| Uncompressed | .parquet | oss://oss-shipper-shenzhen/ecs_test/2016/01/26/20154_1453812893059571256_937.parquet | You can download the OSS buckets to the local host and use the <code>parquet-tools</code> utility to open the objects. For more information about the <code>parquet-tools</code> utility, visit parquet-tools . |
| Snappy | .snappy.parquet | oss://oss-shipper-shenzhen/ecs_test/2016/01/26/20154_1453812893059571256_937.snappy.parquet | You can download the OSS buckets to the local host and use the <code>parquet-tools</code> utility to open the objects. For more information about the <code>parquet-tools</code> utility, visit parquet-tools . |

23.7.1.5. Decompress Snappy compressed files

When you ship data from Log Service to Object Storage Service (OSS), you can use Snappy to compress OSS objects. After the data is shipped to OSS, you can decompress OSS objects by using the C++ library, Java library, Python library, and decompression tool for Linux.

Use the C++ library to decompress OSS objects

Download the C++ library from the [snappy](#) page and use the `Snappy.Uncompress` method to decompress Snappy compressed OSS objects.

Use the Java library to decompress OSS objects

Download the Java library from the [xerial snappy-java](#) page and use the `Snappy.Uncompress` or `Snappy.SnappyInputStream` method to decompress Snappy compressed OSS objects. The `SnappyFramedInputStream` method is not supported.

 **Note** If you use Java Library 1.1.2.1, some Snappy compressed OSS objects may fail to be decompressed. For more information, see [Bad handling of the MAGIC HEADER](#). To fix this issue, you can use Java Library 1.1.2.6 or later.

```
<dependency>
<groupId>org.xerial.snappy</groupId>
<artifactId>snappy-java</artifactId>
<version>1.0.4.1</version>
<type>jar</type>
<scope>compile</scope>
</dependency>
```

- `Snappy.Uncompress`

```
String fileName = "C:\\Downloads\\36_1474212963188600684_4451886.snappy";
RandomAccessFile randomFile = new RandomAccessFile(fileName, "r");
int fileLength = (int) randomFile.length();
randomFile.seek(0);
byte[] bytes = new byte[fileLength];
int byteread = randomFile.read(bytes);
System.out.println(fileLength);
System.out.println(byteread);
byte[] uncompressed = Snappy.uncompress(bytes);
String result = new String(uncompressed, "UTF-8");
System.out.println(result);
```

- `Snappy.SnappyInputStream`

```
String fileName = "C:\\Downloads\\36_1474212963188600684_4451886.snappy";
SnappyInputStream sis = new SnappyInputStream(new FileInputStream(fileName));
byte[] buffer = new byte[4096];
int len = 0;
while ((len = sis.read(buffer)) != -1) {
    System.out.println(new String(buffer, 0, len));
}
```

Use the Python Library to decompress OSS objects

1. Download and install the [Python library](#).
2. Run the decompression script.

The following example is a sample decompression script:

```
import snappy
compressed = open('/tmp/temp.snappy').read()
snappy.uncompress(compressed)
```

 **Note** The following two commands cannot be used to decompress Snappy compressed OSS objects. These commands can be used only in Hadoop mode (`hadoop_stream_decompress`) or streaming mode (`stream_decompress`).

```
python -m snappy -c uncompressed_file compressed_file.snappy
python -m snappy -d compressed_file.snappy uncompressed_file
```

Use decompression tools for Linux to decompress OSS buckets

Log Service allows you to decompress Snappy compressed files by using the decompression tool for Linux. Click [snappy_tool](#) to download the tool. Replace `03_1453457006548078722_44148.snappy` and `03_1453457006548078722_44148` in the following code with the values specific to your environment and then run the following code:

```
./snappy_tool 03_1453457006548078722_44148.snappy 03_1453457006548078722_44148
compressed.size: 2217186
snappy::Uncompress return: 1
uncompressed.size: 25223660
```

23.8. RAM

23.8.1. Overview

Resource Access Management (RAM) is a resource access control service provided by Apsara Stack.

You can use RAM to manage users, including employees, systems, and applications. You can also use RAM to grant users permissions to access resources.

RAM provides the following features:

- RAM Role

To authorize a cloud service in a level-1 organization to use other resources in the organization, you must create a RAM role. This role specifies the operations that the cloud service can perform on the resources.

Only system administrators and level-1 organization administrators can create RAM roles.

- User group

You can create multiple RAM users for an organization and grant the users different permissions on the same cloud resources in the organization.

You can create RAM user groups to classify and authorize RAM users within your Apsara Stack tenant account. This simplifies the management of RAM users and their permissions.

You can create RAM policies to grant permissions to different user groups.

23.8.2. Create a RAM role

To authorize a cloud service in a level-1 organization to use other resources in the organization, you must create a RAM role. This role specifies the operations that the cloud service can perform on the resources.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an administrator.
For more information, see [Log on to the Log Service console](#).
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.

4. In the upper-right corner of the page, click **Create RAM Role**.
5. On the **Roles - Create RAM Role** page that appears, set the **Role Name** and **Description** parameters.
6. Click **Create**.

23.8.3. Create a user

This topic describes how an administrator creates a user and assigns a role to the user. The role varies based on the cloud resources that the user needs to access.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an administrator.
For more information, see [Log on to the Log Service console](#).
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Users**.
4. On the **Users** page, click **Create**.
5. In the dialog box that appears, set the parameters based on your requirements.

| Parameter | Description |
|--------------------------------|--|
| Username | The username. |
| Display Name | The display name of the user. |
| Roles | The role that you want to assign to the user. |
| Organization | The organization to which the user belongs. |
| Logon Policy | The logon policy that restricts the logon time and IP address of the user. If you do not specify this parameter, the default policy is attached to the created user. |
| Mobile Number | The mobile phone number of the user. If you need to send text messages about the usage and requests for resources to the mobile phone number, make sure that the specified mobile phone number is correct. |
| Landline Number | Optional. The landline number of the user. |
| Email | The email address of the user. If you need to send emails about the usage and requests for resources to the email address, make sure that the specified email address is correct. |
| DingTalk Key | Optional. The DingTalk key. |
| Notify User by SMS | Specifies whether to send text messages about the usage and requests for resources to the specified mobile phone number. |
| Notify User by Email | Specifies whether to send emails about the usage and requests for resources to the specified email address. |
| Notify User by DingTalk | Specifies whether to send messages about the usage and requests for resources to the specified DingTalk user. |

6. Click **OK**.

23.8.4. Create a RAM user group

This topic describes how to create a RAM user group in an organization and grant permissions to RAM users in the RAM user group.

Prerequisites

An organization is created.

Context

The relationships between RAM user groups and RAM users:

- A RAM user group can contain zero or more RAM users.
- A RAM user does not need to belong to a RAM user group.
- You can add a RAM user to multiple RAM user groups.

The relationships between RAM user groups and organizations:

- A RAM user group belongs to only one organization.
- You can create multiple RAM user groups in an organization.

The relationships between RAM user groups and RAM roles:

- Only one RAM role can be assigned to each RAM user group.
- A RAM role can be assigned to multiple RAM user groups.
- When a RAM role is assigned to a RAM user group, the permissions that the RAM role has are automatically granted to RAM users in the RAM user group.

The relationships between RAM user groups and resource sets:

- You can add zero or more user groups to a resource set.
- A user group can be added to multiple resource sets.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an administrator.
For more information, see [Log on to the Log Service console](#).
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **User Groups**.
4. In the upper-right corner of the page, click **Create User Group**.
5. In the dialog box that appears, set the **User Group Name** and **Organization** parameters.
6. Click **OK**.

23.8.5. Add a RAM user to a RAM user group

This topic describes how to add a RAM user to a RAM user group.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an administrator.
For more information, see [Log on to the Log Service console](#).
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **User Groups**.
4. Find the user group to which you want to add a RAM user, and click **Add User** in the **Actions** column.
5. In the dialog box that appears, select a RAM user from the left pane, and click the right arrow to move the RAM user to the right pane.
6. Click **OK**.

23.8.6. Create a permission policy

If you want to use a cloud service to access the resources of other cloud services, you must create a permission policy for a RAM role. Then, the policy is automatically attached to the RAM user group to which the RAM role is assigned.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an administrator.
For more information, see [Log on to the Log Service console](#).
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. In the list of role names, find the RAM role for which you want to create a permission policy and choose **More > Modify**.
5. Click **Permissions**.
6. Click **Add Permission Policy**.
7. In the **Add Permission Policy** dialog box, enter the policy information.

Add Permission Policy

*Policy Name:
Enter a policy name 0/15

Description:
Enter 0 to 100 characters 0/100

*Policy Details:
1 | The details of the specified policy must be 2,048 characters in length, and follow the JSON format

OK Cancel

For more information about how to specify the policy information, see [Use custom policies to grant RAM user the required permissions](#).

23.8.7. Grant a RAM user the permissions to manage a project

This topic describes how to grant a Resource Access Management (RAM) user the permissions to manage a specified project.

Prerequisites

- A RAM user is created. For more information, see [Create a user](#).
- A resource set is created and the RAM user is added to the resource set. For more information, see [Enterprise Center > Resource Sets](#) in *Apsara Uni-manager Management Console User Guide*.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an administrator.
For more information, see [Log on to the Log Service console](#).
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Data Permissions**.
4. Click the resource set that you want to manage, for example, ResourceSet(appstreaming).
5. In the **Product Type** section, click **Log Service**.
6. Find the instance that you want to manage and click **Authorize** in the Actions column.
In this example, the instance is a Log Service project whose name is test-project.
7. Grant the RAM user the permissions to manage the project.
If you turn on the Action switch of the **Update Project** permission, the RAM user can modify the project that is selected in Step .

23.8.8. Grant permissions to a RAM role

This topic describes how to grant permissions to a RAM role. After a RAM role is granted permissions, the RAM users in the associated RAM user group inherit the permissions.

Procedure

1. Log on to the Apsara Uni-manager Management Console as an administrator.
For more information, see [Log on to the Log Service console](#).
2. In the top navigation bar, click **Enterprise**.
3. In the left-side navigation pane of the **Enterprise** page, click **Roles**.
4. In the list of role names, find the RAM role for which you want to create a permission policy and choose **More > Modify**.
5. On the **Permissions** tab, click **Select Existing Permission Policy**.
6. In the **Select Existing Permission Policy** dialog box, select a permission policy and click **OK**.
If no policies are available, create a policy. For more information, see [Create a permission policy](#).

23.8.9. Use custom policies to grant RAM user the required permissions

This topic describes how to use custom Resource Access Management (RAM) policies to grant RAM users the required permissions. You can grant permissions to the RAM users under your Apsara Stack tenant account.

Context

For data security, we recommend that you follow the principle of least privilege (PoLP) when you grant permissions to RAM users. You must grant the read-only permission on the project list to RAM users. Otherwise, the RAM users cannot view the projects in the project list.

Use the RAM console to grant permissions to a RAM user

- The read-only permission on projects

For example, you can use your Apsara Stack tenant account to grant RAM users the following permissions:

- The permission to view the list of projects that belong to the Apsara Stack tenant account
- The permission to read specific projects

Use the following policy:

```
{
  "Version": "1",
  "Statement": [
    {
      "Action": ["log:ListProject"],
      "Resource": ["acs:log:*:*:project/*"],
      "Effect": "Allow"
    },
    {
      "Action": [
        "log:Get*",
        "log:List*"
      ],
      "Resource": "acs:log:*:*:project/<The name of the project>/*",
      "Effect": "Allow"
    }
  ]
}
```

- The permission to read a Logstore, save searches, and use saved searches.

For example, you can use your Apsara Stack tenant account to grant RAM users the following permissions:

- The permission to view the project list of the Apsara Stack tenant account
- The permission to read a Logstore, save searches, and use saved searches

Use the following policy:

```

{
  "Version": "1",
  "Statement": [
    {
      "Action": [
        "log:ListProject"
      ],
      "Resource": "acs:log:*:*:project/*",
      "Effect": "Allow"
    },
    {
      "Action": [
        "log:List*"
      ],
      "Resource": "acs:log:*:*:project/<The name of the project>/logstore/*",
      "Effect": "Allow"
    },
    {
      "Action": [
        "log:Get*",
        "log:List*"
      ],
      "Resource": [
        "acs:log:*:*:project/<The name of the project>/logstore/<The name of the Logstore>"
      ],
      "Effect": "Allow"
    },
    {
      "Action": [
        "log:List*"
      ],
      "Resource": [
        "acs:log:*:*:project/<The name of the project>/dashboard",
        "acs:log:*:*:project/<The name of the project>/dashboard*"
      ],
      "Effect": "Allow"
    },
    {
      "Action": [
        "log:Get*",
        "log:List*",
        "log:Create*"
      ],
      "Resource": [
        "acs:log:*:*:project/<The name of the project>/savedsearch",
        "acs:log:*:*:project/<The name of the project>/savedsearch*"
      ],
      "Effect": "Allow"
    }
  ]
}

```

 **Note** In the policy, a value of the Resource attribute that does not end with an asterisk (*) indicates the exact resource. A value that ends with an asterisk (*) indicates all resources that match the value.

- The permission to read a Logstore and view all saved searches and dashboards in a project

For example, you can use your Apsara Stack tenant account to grant a RAM user the following permissions:

- The permission to view the project list of the Apsara Stack tenant account
- The permission to read a Logstore and view all saved searches and dashboards in a project

Use the following policy:

```
{
  "Version": "1",
  "Statement": [
    {
      "Action": [
        "log:ListProject"
      ],
      "Resource": "acs:log:*:*:project/*",
      "Effect": "Allow"
    },
    {
      "Action": [
        "log:List*"
      ],
      "Resource": "acs:log:*:*:project/<The name of the project>/logstore/*",
      "Effect": "Allow"
    },
    {
      "Action": [
        "log:Get*",
        "log:List*"
      ],
      "Resource": [
        "acs:log:*:*:project/<The name of the project>/logstore/<The name of the Logstore>"
      ],
      "Effect": "Allow"
    },
    {
      "Action": [
        "log:Get*",
        "log:List*"
      ],
      "Resource": [
        "acs:log:*:*:project/<The name of the project>/dashboard",
        "acs:log:*:*:project/<The name of the project>/dashboard/*"
      ],
      "Effect": "Allow"
    },
    {
      "Action": [
        "log:Get*",
        "log:List*"
      ],
      "Resource": [
        "acs:log:*:*:project/<The name of the project>/savedsearch",
        "acs:log:*:*:project/<The name of the project>/savedsearch/*"
      ],
      "Effect": "Allow"
    }
  ]
}
```

Grant RAM users the permissions that are required to call Log Service operations

- The permission to write data to a project

To grant RAM users only the permission to write data to a project, use the following policy:

```
{
  "Version": "1",
  "Statement": [
    {
      "Action": [
        "log:Post*"
      ],
      "Resource": "acs:log:*:*:project/<The name of the project>/*",
      "Effect": "Allow"
    }
  ]
}
```

- The permission to consume data of a project

To grant RAM users only the permission to consume data of a project, use the following policy:

```
{
  "Version": "1",
  "Statement": [
    {
      "Action": [
        "log:ListShards",
        "log:GetCursorOrData",
        "log:GetConsumerGroupCheckPoint",
        "log:UpdateConsumerGroup",
        "log:ConsumerGroupHeartBeat",
        "log:ConsumerGroupUpdateCheckPoint",
        "log:ListConsumerGroup",
        "log:CreateConsumerGroup"
      ],
      "Resource": "acs:log:*:*:project/<The name of the project>/*",
      "Effect": "Allow"
    }
  ]
}
```

- The permission to consume data of a Logstore

To grant RAM users only the permission to consume data of a Logstore, use the following policy:

```
{
  "Version": "1",
  "Statement": [
    {
      "Action": [
        "log:ListShards",
        "log:GetCursorOrData",
        "log:GetConsumerGroupCheckPoint",
        "log:UpdateConsumerGroup",
        "log:ConsumerGroupHeartBeat",
        "log:ConsumerGroupUpdateCheckPoint",
        "log:ListConsumerGroup",
        "log:CreateConsumerGroup"
      ],
      "Resource": [
        "acs:log:*:*:project/<The name of the project>/logstore/<The name of the Logstore>",
        "acs:log:*:*:project/<the name of the project>/logstore/<the name of the Logstore>/*"
      ],
      "Effect": "Allow"
    }
  ]
}
```

23.9. FAQ

23.9.1. Log collection

23.9.1.1. How do I troubleshoot Logtail collection errors?

If the Logtail preview page is blank or "No Data" is displayed on the query page, perform the following steps:

Procedure

1. Check whether Log Service receives heartbeats from the server group.

You can view the Logtail heartbeat status in the Log Service console. For more information, see [View the status of a server group](#).

If the heartbeat status is OK, go to the next step. If the heartbeat status is FAIL, proceed with further troubleshooting. For more information, see [What can I do if no heartbeat packet is received from a Logtail client?](#).

2. Check whether the Logtail configuration is created.

If the heartbeat status of Logtail is OK, check whether the Logtail configuration is created. Make sure that the path and name of monitored logs match the files that are stored on the server. The path can be a full path or a path that includes wildcards.

3. Make sure that the Logtail configuration is applied to the server group.

Check whether the target Logtail configuration is applied to the server group. For more information, see [Manage server group configurations](#).

4. Check collection errors.

If Logtail is configured correctly, check whether new logs are generated in real time. Logtail only collects incremental log data. Logtail does not read log files in which no log is generated. If a log file is updated but the updates cannot be queried in Log Service, you can diagnose the problem as follows:

- o View logs of the Logtail client

The client logs include key INFO logs, all WARNING logs, and all ERROR logs. To view complete error information in real time, check the following client logs:

- Linux: `/usr/local/ilogtail/ilogtail.LOG`.
 - Linux: `/usr/local/ilogtail/ilogtail_plugin.LOG`. The file contains the logs such as HTTP logs, MySQL binary logs, and MySQL query results.
 - 64-bit Windows: `C:\Program Files (x86)\Alibaba\Logtail\logtail_*.log`.
 - 32-bit Windows: `C:\Program Files\Alibaba\Logtail\logtail_*.log`.
- Check whether the log volume exceeds the limit.

To collect large volumes of logs, you may need to modify the Logtail startup parameters for higher log collection throughput. For more information, see [Set Logtail startup parameters](#).

23.9.1.2. What can I do if Log Service does not receive heartbeats from a Logtail client?

If Log Service does not receive heartbeats from a Logtail client, perform the steps that are described in the topic to troubleshoot the problem.

Context

After Logtail is installed on a server, the Logtail client sends heartbeats to Log Service. If the status page of the machine group shows that Log Service does not receive heartbeats from a Logtail client, it indicates that the Logtail client is not installed or disconnected from the server.

Step 1: Check whether Logtail is installed

Use the following method to check whether Logtail is installed:

- On a Linux server, run the following command:

```
sudo /etc/init.d/ilogtaild status
```

If the command returns `ilogtail is running`, it indicates that Logtail is installed. The following script shows an example command and response:

```
[root@*****~]# sudo /etc/init.d/ilogtaild status
ilogtail is running
```

- On a Windows server:
 - i. Press Win+R. In the Run dialog box, enter `services.msc` and click **OK**.
 - ii. In the **Services** window, check the status of the `LogtailDaemon` and `LogtailWorker` services. If the services are in the Running state, it indicates that Logtail is installed.

If Logtail is not installed, install it. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#). Ensure that the Log Service endpoint in the Logtail installation command corresponds to the region to which the Log Service project belongs. If Logtail is running, go to the next step.

Step 2: Check the Log Service endpoint in the Logtail installation command

When you install Logtail, you must specify a [Log Service endpoint](#) based on the region to which the Log Service project belongs. If the endpoint is incorrect or the Logtail installation command is invalid, Log Service cannot receive heartbeats from the Logtail client.

You can view the Log Service endpoint and installation method in the Logtail configuration file named `ilogtail_config.json`. The file is stored in the following path:

- Linux: `/usr/local/ilogtail/ilogtail_config.json`

- 64-bit Windows: `C:\Program Files (x86)\Alibaba\Logtail\logtail_config.json`
- 32-bit Windows: `C:\Program Files\Alibaba\Logtail\logtail_config.json`

In the Logtail configuration file, check the value of the `config_server_address` parameter. This parameter specifies the Log Service endpoint. Then, check whether the Logtail client can connect to Log Service based on the endpoint. For example, if the endpoint that is recorded in the Logtail configuration file is `logtail.cn-qingdao-env25-d01.sls-pub.inter.env25.shuguang.com`, you can run the following command to check the connection:

- Linux:

```
curl logtail.cn-qingdao-env25-d01.sls-pub.inter.env25.shuguang.com
```

- Windows:

```
telnet logtail.cn-qingdao-env25-d01.sls-pub.inter.env25.shuguang.com 80
```

If the Log Service endpoint in the Logtail installation command is incorrect, re-install Logtail. For more information, see [Install Logtail in Linux](#) or [Install Logtail in Windows](#).

If the Log Service endpoint in the Logtail installation command is correct, go to the next step.

Step 3: Check the server IP addresses in the machine group

The server IP address that is obtained by a Logtail client must be configured in the machine group. Otherwise, Log Service cannot receive heartbeats or collect logs from the Logtail client. Logtail uses the following methods to obtain the IP address of a server:

- If the server is not bound with a hostname, Logtail obtains the IP address of the first network interface controller (NIC) card of the server.
- If the server is bound with a hostname, Logtail obtains the IP address that corresponds to the hostname. You can view the hostname and IP address in the `/etc/hosts` file.

 **Note** You can run the `hostname` command to query the hostname.

Perform the following steps to check whether the server IP address that is obtained by the Logtail client is configured in the machine group.

1. Check the server IP address that is obtained by Logtail.

The `ip` field in the `app_info.json` file indicates the server IP address that is obtained by Logtail. The file is stored in the following path:

- Linux: `/usr/local/logtail/app_info.json`
- 64-bit Windows: `C:\Program Files (x86)\Alibaba\Logtail\app_info.json`
- 32-bit Windows: `C:\Program Files\Alibaba\Logtail\app_info.json`

 **Note**

- If the `ip` field in the `app_info.json` file is empty, Logtail cannot work. In this case, you must configure an IP address for the server and restart Logtail.
- The `app_info.json` file is used only to record information. If you modify the IP address in the file, the server IP address obtained by Logtail is not updated.

2. Check the server IP addresses in the machine group.

Log on to the Log Service console. In the **Projects** section, click the project to which the machine group belongs. In the left-side navigation pane, click the **Machine Groups** icon. In the **Machine Groups** pane, click the machine group. In the **Machine Group Status** section of the **Machine Group Settings** page, check the server IP addresses.

If no server IP address in the machine group is the same as the IP address that is obtained by Logtail, perform the following step to modify the IP address configurations in the Log Service console:

- If a server IP address in the machine group is incorrect, change the IP address to the IP address that is obtained by Logtail. Then, check the heartbeat status 1 minute after you save the change.
- If you have modified the IP address of the server where Logtail is installed (for example, the `/etc/hosts` file is modified), restart Logtail. After Logtail obtains the new server IP address, set a server IP address in the machine group to the value of the `ip` field in the `app_info.json` file.

You can use the following method to restart Logtail:

- On a Linux server, run the following commands:

```
sudo /etc/init.d/ilogtaild stop
sudo /etc/init.d/ilogtaild start
```

- On a Windows server:

Press Win+R. In the Run dialog box, enter `services.msc` and click OK. In the **Services** window, find and restart the `LogtailWorker` service.

23.9.1.3. How do I query the local log collection statuses?

You can use Logtail to query the health status of Logtail and log collection statuses. The statuses help you troubleshoot log collection issues and customize status monitoring for log collection.

Instructions

After a Logtail client that supports the status query feature is installed, you can query the local log collection statuses by running commands on the client. For more information about how to install Logtail, see [Install Logtail in Linux](#).

You can run the `/etc/init.d/ilogtaild -h` command on the client to check whether a client supports the feature of querying the local log collection status. If the command output includes the `logtail insight, version` keyword, it indicates that the client supports the status query feature.

```
/etc/init.d/ilogtaild -h
Usage: ./ilogtaild { start | stop (graceful, flush data and save checkpoints) | force-stop | status | -h for help}$
logtail insight, version : 0.1.0
command list :
  status all [index]
  get logtail running status
  status active [--logstore | --logfile] index [project] [logstore]
  list all active logstore | logfile. if use --logfile, please add project and logstore. default --logstore
  status logstore [--format=line | json] index project logstore
  get logstore status with line or json style. default --format=line
  status logfile [--format=line | json] index project logstore fileFullPath
  get log file status with line or json style. default --format=line
  status history beginIndex endIndex project logstore [fileFullPath]
  query logstore | logfile history status.
index: from 1 to 60. in all, it means last $(index) minutes; in active/logstore/logfile/history, it means last $(index)*10 minutes
```

The following table describes the commands that are supported by Logtail:

| Command | Function | Maximum time range that can be queried | Time window |
|---------|--------------------------------|--|-------------|
| all | Queries the status of Logtail. | Last 60 minutes | 1 minute |

| Command | Function | Maximum time range that can be queried | Time window |
|----------|---|--|-------------|
| active | Queries the active Logstores that are collecting logs and the active log files from which logs are being collected. | Last 600 minutes | 10 minutes |
| logstore | Queries the collection status of a Logstore. | Last 600 minutes | 10 minutes |
| logfile | Queries the collection status of a log file. | Last 600 minutes | 10 minutes |
| history | Queries the collection status of a Logstore or log file in the query time window. | Last 600 minutes | 10 minutes |

Note

- The `index` parameter in the preceding commands indicates the index of the time window. Valid values: 1 to 60. The index of the latest time window is 1 and the time window ends at the current system time. If you specify a 1-minute time window, the status in the past interval of `(index, index-1]` minutes is returned. If you specify a 10-minute time window, the status in the past interval of `(10*index, 10*(index-1)]` minutes is returned.
- All commands in the preceding table is the subcommands of the status command.

Command all

Command syntax

```
/etc/init.d/ilogtaild status all [ index ]
```

Note The all command is used to query the status of Logtail. The index parameter is optional. Default value: 1.

Examples

```
/etc/init.d/ilogtaild status all 1
ok
/etc/init.d/ilogtaild status all 10
busy
```

Response

| Status | Description | Priority | Troubleshooting |
|--------|---|----------|------------------------|
| ok | Logtail is running as expected. | N/A | No action is required. |
| busy | The collection speed is high, and Logtail is running as expected. | N/A | No action is required. |

| Status | Description | Priority | Troubleshooting |
|----------------|---|----------|---|
| many_log_files | A large number of log files are being collected by Logtail. | Low | You can check the Logtail configuration for log files that do not need to be collected. |
| process_block | The process of log parsing is blocked. | Low | You can check whether a large number of logs are generated in a short time. If you use the all command for multiple times and the returned value is always process_block, you can modify the limit of CPU usage or the limit of concurrent packet sending . |
| send_block | The process of sending log packets is blocked. | High | You can check whether a large number of logs are generated in a short time and the network connection is stable. If you use the all command for multiple times and the returned value is always send_block, you can modify the limit of CPU usage or the limit of concurrent packet sending . |

Command active

Command syntax

```
/etc/init.d/ilogtaild status active [--logstore] index
/etc/init.d/ilogtaild status active --logfile index project-name logstore-name
```

Note

- You can use the `active [--logstore] index` command to query all active Logstores. The `--logstore` parameter is optional.
- The command `active --logfile index project-name logstore-name` is used to query all active log files in the Logstore of a project.
- The active command is used to query log files. We recommend that you query active Logstores before querying active log files in the Logstores.

Examples

```
/etc/init.d/ilogtaild status active 1
sls-zc-test : release-test
sls-zc-test : release-test-ant-rpc-3
sls-zc-test : release-test-same-regex-3
/etc/init.d/ilogtaild status active --logfile 1 sls-zc-test release-test
/disk2/test/normal/access.log
```

Response

- If you run the `active --logstore index` command, the names of the active Logstores are returned in the following format: `project-name : logstore-name`. If you run the command `active --logfile index project-name logstore-name`, the paths of active log files are returned.
- The status of the inactive Logstores or inactive log files in the query time window is not returned.

Command logstore

Command syntax

```
/etc/init.d/ilogtailed status logstore [--format={line|json}] index project-name logstore-name
```

 Note

- The logstore command is used to query the collection status of the specified project and Logstore in the LINE or JSON format.
- The default value of the `--format=` parameter is `--format=line`, which indicates that the status is returned in the LINE format. Noted that the `--format=` parameter is placed after the `logstore` parameter.
- If the Logstore specified in the preceding command does not exist or is not active in the query time window, an empty response in LINE format or the `null` value in the JSON format is returned.

Examples

```

/etc/init.d/ilogtaild status logstore 1 sls-zc-test release-test-same
time_begin_readable : 17-08-29 10:56:11
time_end_readable : 17-08-29 11:06:11
time_begin : 1503975371
time_end : 1503975971
project : sls-zc-test
logstore : release-test-same
status : ok
config : ##1.0##sls-zc-test$same
read_bytes : 65033430
parse_success_lines : 230615
parse_fail_lines : 0
last_read_time : 1503975970
read_count : 687
avg_delay_bytes : 0
max_unsend_time : 0
min_unsend_time : 0
max_send_success_time : 1503975968
send_queue_size : 0
send_network_error_count : 0
send_network_quota_count : 0
send_network_discard_count : 0
send_success_count : 302
send_block_flag : false
sender_valid_flag : true
/etc/init.d/ilogtaild status logstore --format=json 1 sls-zc-test release-test-same
{
  "avg_delay_bytes" : 0,
  "config" : "##1.0##sls-zc-test$same",
  "last_read_time" : 1503975970,
  "logstore" : "release-test-same",
  "max_send_success_time" : 1503975968,
  "max_unsend_time" : 0,
  "min_unsend_time" : 0,
  "parse_fail_lines" : 0,
  "parse_success_lines" : 230615,
  "project" : "sls-zc-test",
  "read_bytes" : 65033430,
  "read_count" : 687,
  "send_block_flag" : false,
  "send_network_discard_count" : 0,
  "send_network_error_count" : 0,
  "send_network_quota_count" : 0,
  "send_queue_size" : 0,
  "send_success_count" : 302,
  "sender_valid_flag" : true,
  "status" : "ok",
  "time_begin" : 1503975371,
  "time_begin_readable" : "17-08-29 10:56:11",
  "time_end" : 1503975971,
  "time_end_readable" : "17-08-29 11:06:11"
}

```

Response

| Parameter | Description | Unit |
|-----------|-------------|------|
|-----------|-------------|------|

| Parameter | Description | Unit |
|---------------------|---|---------------------------|
| status | The status of the Logstore. For information about Logstore statuses and actions that are required to deal with each status, see the following table. | N/A |
| time_begin_readable | The time when logs become readable. | N/A |
| time_end_readable | The time when logs become unreadable. | N/A |
| time_begin | The time when statistics collection starts. | Unix timestamp in seconds |
| time_end | The time when statistics collection ends. | Unix timestamp in seconds |
| project | The name of the project. | N/A |
| logstore | The name of the Logstore. | N/A |
| config | The name of the Logtail configuration, which is globally unique. The format of the name is <code>##1.0## + project + \$ + config</code> . | N/A |
| read_bytes | The amount of the log data that is read in the query time window. | Byte |
| parse_success_lines | The number of the log lines that are parsed in the query time window. | Line |
| parse_fail_lines | The number of the log lines that fail to be parsed in the query time window. | Line |
| last_read_time | The last time when logs are read in the query time window. | Unix timestamp in seconds |
| read_count | The number of times that the log file is read in the query time window. | Times |
| avg_delay_bytes | The average of difference between the actual file size and the offset generated when reading log data each time in the query time window. | Byte |
| max_unsend_time | The maximum period of time for which an unsend packet waits in the sending queue. An unsend packet refers to a packet that has not been sent at the end of the query time window. If no packets exist in the queue, the value is 0. | Unix timestamp in seconds |

| Parameter | Description | Unit |
|----------------------------|---|---------------------------|
| min_unsend_time | The minimum period of time for which an unsend packet waits in the sending queue. Unsend packets refer to packets that have not been sent at the end of the query time window. If no packets exist in the queue, the value is 0. | Unix timestamp in seconds |
| max_send_success_time | The maximum period of time when a packet waited in the sending queue. | Unix timestamp in seconds |
| send_queue_size | The number of the unsend packets in the sending queue at the end of the query time window. | Number of packets |
| send_network_error_count | The number of the packets that cannot be sent due to network errors in the query time window. | Number of packets |
| send_network_quota_count | The number of the packets that cannot be sent due to quota limit in the query time window. | Number of packets |
| send_network_discard_count | The number of the packets that are discarded due to data errors or lack of permissions. | Number of packets |
| send_success_count | The number of the packets that are sent in the query time window. | Number of packets |
| send_block_flag | Indicates whether the sending queue is blocked at the end of the query time window. | N/A |
| sender_valid_flag | Indicates whether the sender flag of the Logstore is valid. The value true indicates that the sender flag is valid. The value false indicates that the sender flag is invalid and disabled because of a network error or quota error. | N/A |

Logstore statuses

| Status | Description | Troubleshooting |
|---------------|--|---|
| ok | Logtail is running as expected. | No action is required. |
| process_block | The process of log parsing is blocked. | You can check whether a large number of logs are generated in a short time. If you use the all command for multiple times and the returned value is always process_block, you can Set Logtail startup parameters modify the limit of CPU usage or of concurrent packet sending. |
| parse_fail | Logtail fails to parse logs. | You can check whether the format of logs is consistent with that you set in the Logtail configuration. |

| Status | Description | Troubleshooting |
|------------|--|---|
| send_block | The process of sending log packets is blocked. | You can check whether a large number of logs are generated in a short time and the network connection is stable. If you use the all command for multiple times and the returned value is always send_block, you can Set Logtail startup parameters modify the limit of CPU usage or of concurrent packet sending. |

Command logfile

Command syntax

```
/etc/init.d/ilogtail status logfile [--format={line|json}] index project-name logstore-name fileFullPath
```

Note

- The logfile command is used to query the collection status of the specified log files in the LINE or JSON format.
- The default value of the `--format=` parameter is `--format=line`, which indicates that the status is returned in the LINE format.
- If the log file specified in the command does not exist or is not active in the query time window, an empty response in the LINE format or the `null` value in the JSON format is returned.
- The `--format` parameter is placed after the `logfile` parameter.
- The value of the `filefullpath` parameter must be the full path of the log file.

Examples

```

/etc/init.d/ilogtaild status logfile 1 sls-zc-test release-test-same /disk2/test/normal/access.log
time_begin_readable : 17-08-29 11:16:11
time_end_readable : 17-08-29 11:26:11
time_begin : 1503976571
time_end : 1503977171
project : sls-zc-test
logstore : release-test-same
status : ok
config : ##1.0##sls-zc-test$same
file_path : /disk2/test/normal/access.log
file_dev : 64800
file_inode : 22544456
file_size_bytes : 17154060
file_offset_bytes : 17154060
read_bytes : 65033430
parse_success_lines : 230615
parse_fail_lines : 0
last_read_time : 1503977170
read_count : 667
avg_delay_bytes : 0
/etc/init.d/ilogtaild status logfile --format=json 1 sls-zc-test release-test-same /disk2/test/normal/access.log
{
  "avg_delay_bytes": 0,
  "config": "##1.0##sls-zc-test$same",
  "file_dev": 64800,
  "file_inode": 22544456,
  "file_path": "/disk2/test/normal/access.log",
  "file_size_bytes": 17154060,
  "last_read_time": 1503977170,
  "logstore": "release-test-same",
  "parse_fail_lines": 0,
  "parse_success_lines": 230615,
  "project": "sls-zc-test",
  "read_bytes": 65033430,
  "read_count": 667,
  "read_offset_bytes": 17154060,
  "status": "ok",
  "time_begin": 1503976571,
  "time_begin_readable": "17-08-29 11:16:11",
  "time_end": 1503977171,
  "time_end_readable": "17-08-29 11:26:11"
}

```

Response

| Parameter | Description | Unit |
|---------------------|---|------|
| status | The collection status of the log file in the query time window. For more information, see the status parameter in the Command logstore section. | N/A |
| time_begin_readable | The time when logs become readable. | N/A |
| time_end_readable | The time when logs become unreadable. | N/A |

| Parameter | Description | Unit |
|---------------------|---|---------------------------|
| time_begin | The time when statistics collection starts. | Unix timestamp in seconds |
| time_end | The time when statistics collection ends. | Unix timestamp in seconds |
| project | The name of the project. | N/A |
| logstore | The name of the Logstore. | N/A |
| file_path | The path of the log file. | N/A |
| file_dev | The ID of the device from which the log file is collected. | N/A |
| file_inode | The inode of the log file. | N/A |
| file_size_bytes | The size of the log file that is last scanned in the query time window. | Byte |
| read_offset_bytes | The parsing offset of the log file. | Byte |
| config | The name of the Logtail configuration, which is globally unique. The format of the name is <code>##1.0## + project + \$ + config</code> . | N/A |
| read_bytes | The amount of the log data that is read in the query time window. | Byte |
| parse_success_lines | The number of the log lines that are parsed in the query time window. | Line |
| parse_fail_lines | The number of the log lines that fail to be parsed in the query time window. | Line |
| last_read_time | The last time when logs are read in the query time window. | Unix timestamp in seconds |
| read_count | The number of times that the log file is read in the query time window. | Times |
| avg_delay_bytes | The average of difference between the actual file size and the offset generated when reading log data each time in the query time window. | Byte |

Command history

Command syntax

```
/etc/init.d/ilogtaild status history beginIndex endIndex project-name logstore-name [fileFullPath]
```

Note

- The history command is used to query the collection status of a Logstore or log file in the query time window.
- The `beginIndex` and `endIndex` parameters specify the start and end indexes of the range of time windows that you want to query. You must ensure that `beginIndex <= endIndex`.
- The `fileFullPath` parameter is optional. If you specify the path of a log file, the collection status of the log file is queried. If the path is not specified, the collection status of the Logstore is queried.

Examples

```
/etc/init.d/ilogtaild status history 1 3 sls-zc-test release-test-same /disk2/test/normal/access.log
begin_time status read parse_success parse_fail last_read_time read_count avg_delay device inode file_size read_offset
17-08-29 11:26:11 ok 62.12MB 231000 0 17-08-29 11:36:11 671 0B 64800 22544459 18.22MB 18.22MB
17-08-29 11:16:11 ok 62.02MB 230615 0 17-08-29 11:26:10 667 0B 64800 22544456 16.36MB 16.36MB
17-08-29 11:06:11 ok 62.12MB 231000 0 17-08-29 11:16:11 687 0B 64800 22544452 14.46MB 14.46MB
$ /etc/init.d/ilogtaild status history 2 5 sls-zc-test release-test-same
begin_time status read parse_success parse_fail last_read_time read_count avg_delay send_queue network_error quota_error discard_error send_success send_block send_valid max_unsend min_unsend max_send nd_success
17-08-29 11:16:11 ok 62.02MB 230615 0 17-08-29 11:26:10 667 0B 0 0 0 0 30
0 false true 70-01-01 08:00:00 70-01-01 08:00:00 17-08-29 11:26:08
17-08-29 11:06:11 ok 62.12MB 231000 0 17-08-29 11:16:11 687 0B 0 0 0 0 30
3 false true 70-01-01 08:00:00 70-01-01 08:00:00 17-08-29 11:16:10
17-08-29 10:56:11 ok 62.02MB 230615 0 17-08-29 11:06:10 687 0B 0 0 0 0 30
2 false true 70-01-01 08:00:00 70-01-01 08:00:00 17-08-29 11:06:08
17-08-29 10:46:11 ok 62.12MB 231000 0 17-08-29 10:56:11 692 0B 0 0 0 0 30
2 false true 70-01-01 08:00:00 70-01-01 08:00:00 17-08-29 10:56:10
```

Response

- The collection status of the Logstore or log file in each query time window is listed in a line.
- For more information about response parameters, see the Command `logstore` and Command `logfile` sections.

Response status codes**Success code**

If parameters that you specify in a command is valid (even if the queried Logstore or log file is not found), the code 0 is returned. The following section provides two examples:

```
/etc/init.d/ilogtaild status logfile --format=json 1 error-project error-logstore /no/this/file
null
echo $?
0
/etc/init.d/ilogtaild status all
ok
echo $?
0
```

Error codes

If a non-zero code is returned, it indicates that an error occurs. The following table describes the possible non-zero codes.

| Code | Description | Message | Troubleshooting |
|------|---|--|---|
| 10 | The command is invalid or required parameters in the command are not specified. | invalid param, use -h for help. | You can run the -h command for help. |
| 1 | The value of the index parameter is not in the range from 1 to 60. | invalid query interval | You can run the -h command for help. |
| 1 | The collection status in the specified query time window cannot be queried. | query fail, error: \$(error) . For more information, visit errno . | The startup time of Logtail is earlier than the query time window. Otherwise, you can submit a ticket for help. |
| 1 | The start time of querying is out of the query time window. | no match time interval, please check logtail status | You can check whether Logtail is running. If yes, you can submit a ticket for help. |
| 1 | No logs exist in the specified query time window. | invalid profile, maybe logtail restart | You can check whether Logtail is running. If yes, you can submit a ticket for help. |

Examples

```

/etc/init.d/ilogtail status nothiscmd
invalid param, use -h for help.
echo $?
10
/etc/init.d/ilogtail status/all 99
invalid query interval
echo $?
1
    
```

Scenarios

You can query the overall status of Logtail, and specific metrics by querying the collection status during collection. You can customize a mechanism to monitor the log collection status based on the queried information.

Monitor the status of Logtail

You can monitor the status of Logtail by using the `all` command.

For example, you can run the command every minute to query Logtail status. If the `process_block`, `send_block` or `send_error` value is returned for 5 consecutive minutes, an alert is triggered.

You can adjust the alert duration and monitoring scope based on the priorities of the collected log files.

Monitor the log collection status

You can monitor the log collection status of a Logstore by using the `logstore` command.

For example, you can run the `logstore` command every 10 minutes to query the collection status of the Logstore. If the value of the `avg_delay_bytes` parameter exceeds 1 MB (1024 × 1024 bytes) or the value of the `status` parameter is not `ok`, an alert is triggered.

You can adjust the alert threshold for the `avg_delay_bytes` metric based on the size of data that is generated during the log collection.

Check whether Logtail has finished collecting log files

You can check whether Logtail has finished collecting log files by using the `logfile` command.

After Logtail stops collecting log files, you can run the `logfile` command every 10 minutes to query the status of the log file. If the value of the `read_offset_bytes` parameter is the same as that of the `file_size_bytes` parameter, it indicates that the log file is collected.

Troubleshoot log collection issues

If log collection latency occurs on a server, you can use the `history` command to query the status history of log collection.

1. The value of the `send_block_flag` parameter is true. This indicates that the log collection is blocked because of unstable network connections.
 - If the value of the `send_network_quota_count` parameter is greater than 0, split shards in the Logstore. For more information, see [Split a shard](#).
 - If the value of the `send_network_error_count` parameter is greater than 0, check the network connections.
 - If no network error occurs, adjust the [limit of concurrent packet sending and data transfer speed](#) of Logtail.
2. The parameters related to packet sending are set to appropriate values. However, the value of the `avg_delay_bytes` parameter is large.
 - Use the value of the `read_bytes` parameter to calculate the average speed of parsing logs, and then determine whether a large amount of data is transferred during log collection based on the average speed.
 - Adjust the [resource usage limits](#) for Logtail.
3. The value of the `parse_fail_lines` parameter is greater than 0.

Check whether the regular expression can match all required log fields as expected.

23.9.1.4. How do I test a regular expression?

If you select the full regex mode when you configure Logtail to collect and parse text logs, you must specify a regular expression based on your sample log entries. This topic describes how to test a regular expression.

Context

To test a regular expression that you have specified in the Log Service console, you can click **Validate** in the console and check the results as follows:

- For the regular expression that matches the first line of logs, check whether the regular expression can match the expected number of log entries.
- For the fields extracted by the regular expression, check whether the value of each field meets your expectations.

If you want to validate more items and test a regular expression, you can use online tools such as [regex101.com](#) and [regextester.com](#). You can copy and paste the regular expression that is generated by Log Service to an online tool, and specify a sample log entry as the test string.

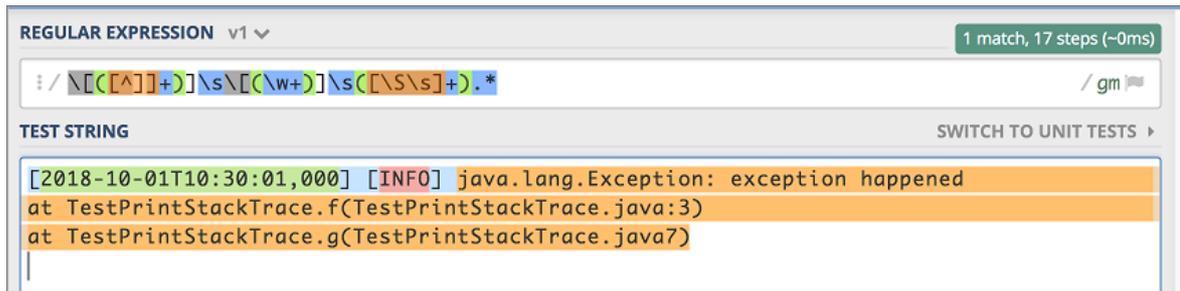
If you use the full regex mode, Log Service automatically generates a regular expression based on a sample log entry. However, the regular expression may fail to match the message field in multi-line log entries as expected. The following example describes how to use the [regex101.com](#) tool to test the regular expression.

Procedure

1. Copy the generated regular expression.
2. Visit the [regex101.com](#) website.
3. Paste the regular expression in the **REGULAR EXPRESSION** field.

On the right side of the page, you can view the explanation of the regular expression.

The following figure shows how the modified regular expression matches the sample log entry with two colons.



You can follow the preceding instructions to test your regular expression. After you validate the regular expression, you can apply it to a Logtail configuration.

23.9.1.5. How do I optimize regular expressions?

You can optimize regular expressions to improve the Logtail performance.

When you optimize regular expressions, we recommend that you follow these rules:

- Use precise characters

We recommend that you do not use the wildcard characters `.*` in a regular expression to match fields in log entries. Using wildcard characters may lead to mismatches and low matching performance. For example, if a field that you want to match only consists of letters, use `[A-Za-z]`.

- Use appropriate quantifiers

We recommend that you do not use plus signs (+) or asterisks (*). For example, you can use `\d` instead of `\d+` or `\d{1,3}` to match the IP address.

- Test and modify regular expressions

You can visit the regex101.com website to test and modify a regular expression to decrease the time required to match log entries.

23.9.1.6. How do I use the full regex mode to collect log entries in multiple formats?

The full regex mode requires that log entries to be collected be in the same format. Therefore, if you want to collect log entries that are in multiple formats, you must use the schema-on-write or schema-on-read solution.

Taking Java logs as an example, the following section lists the types of error log entry and normal log entry.

- Multi-line WARNING log entries
- Simple text INFO log entries
- Key-value DEBUG log entries

```
[2018-10-01T10:30:31,000] [WARNING] java.lang.Exception: another exception happened
at TestPrintStackTrace.f(TestPrintStackTrace.java:3)
at TestPrintStackTrace.g(TestPrintStackTrace.java:7)
at TestPrintStackTrace.main(TestPrintStackTrace.java:16)
[2018-10-01T10:30:32,000] [INFO] info something
[2018-10-01T10:30:33,000] [DEBUG] key:value key2:value2
```

To collect log entries of these types, you can use the following solutions:

- Schema-on-write: To extract log fields, you must apply multiple Logtail configurations with different regular expressions to a log file.

Note However, Logtail cannot apply multiple Logtail configurations directly to the same log file. Therefore, you must set up multiple symbolic links for the directory in which the log file resides. Each Logtail configuration applies to a symbolic link to collect log entries in a specific format.

- Schema-on-read: you can use a common regular expression to collect log entries in different formats.

For example, if you want to collect log entries in multiple formats, you can configure a regular expression that matches the time and log level fields as the first line, and specify the rest of the log entries as the log message. If you want to parse the message, create an index for the message, specify a regular expression to extract log messages, and then extract target fields.

Note We recommend that you use this solution only for scenarios in which tens of millions of log entries are collected, or fewer.

23.9.1.7. How do I set the time format for logs?

You must be familiar with the following rules before setting the time format for logs in Logtail configurations.

- The unit of the timestamp in Log Service is seconds. Therefore, you cannot set the unit as milliseconds or microseconds.
- You only need to set the time field. Other parameters are not required.

The following section lists commonly used formats:

```
Custom1 2017-12-11 15:05:07
%Y-%m-%d %H:%M:%S
Custom2 [2017-12-11 15:05:07.012]
[%Y-%m-%d %H:%M:%S
RFC822 02 Jan 06 15:04 MST
%d %b %y %H:%M
RFC822Z 02 Jan 06 15:04 -0700
%d %b %y %H:%M
RFC850 Monday, 02-Jan-06 15:04:05 MST
%A, %d-%b-%y %H:%M:%S
RFC1123 Mon, 02 Jan 2006 15:04:05 MST
%A, %d-%b-%y %H:%M:%S
RFC3339 2006-01-02T15:04:05Z07:00
%Y-%m-%dT%H:%M:%S
RFC3339Nano 2006-01-02T15:04:05.999999999Z07:00
%Y-%m-%dT%H:%M:%S
```

23.9.1.8. How do I configure non-printable characters in a sample log?

This topic describes how to configure non-printable characters in a sample log entry.

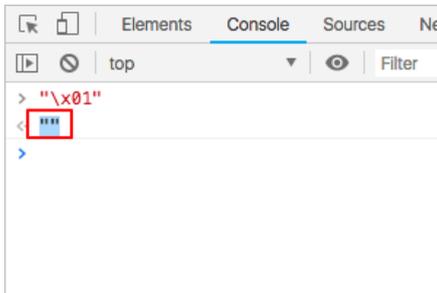
Context

Log Service allows you to specify a non-printable character as the delimiter or quote to collect logs. Non-printable characters are those whose decimal ASCII codes are in the range of 1 to 31 and 127. If you want to specify a non-printable character as the delimiter or quote, you must find the hexadecimal ASCII code of this character and enter this character in the following format: `0x the hexadecimal ASCII code of the non-printable character`. For example, a sample log entry is `123456780`. You can set `0x01` as the delimiter and `0x02` as the quote, and then enter a non-printable character `0x01` between the digits 5 and 6.

Procedure

1. [Log on to the Log Service console.](#)
2. Right-click the blank space on the browser and select **Inspect** from the shortcut menu.
3. Click the **Console** tab on the page that appears.
4. Enter `"\x01"` in the code editor and press Enter.
5. Copy the returned result.

A non-printable character is enclosed in quotation marks (").



6. Paste the returned result between the digits 5 and 6.



7. Delete the quotation marks between the digits 5 and 6.
Then, a non-printable character is configured in a sample log entry.



23.9.1.9. How do I troubleshoot errors during container log collection?

Perform the steps that are described in this topic to troubleshoot an error that occurs when you use Logtail to collect logs from common containers or containers in a Kubernetes cluster.

Related O&M operations

- [Log on to a Logtail container](#)
- [View the operational logs of Logtail](#)
- [Ignore the stdout logs of a Logtail container](#)
- [View the status of Log Service components in a Kubernetes cluster](#)
- [View the version number, IP address, and startup time of Logtail](#)

Troubleshoot an error if Log Service does not receive heartbeats from Logtail clients

Perform the following steps to check whether Logtail is installed:

1. In the machine group, count the number of the servers whose heartbeat status is OK.
 - i. [Log on to the Log Service console](#).
 - ii. Click the project to which the machine group belongs.
 - iii. In the left-side navigation pane, click **Machine Groups**.
 - iv. In the **Machine Groups** pane, click the name of the machine group.

In the **Machine Group Status** section, count the number of the servers whose heartbeat status is OK.

2. Count the number of worker nodes in the cluster.

Run the `kubectl get node | grep -v master` command to query the work nodes in the cluster. Count the number of the work nodes that are returned.

```
$kubectl get node | grep -v master
NAME                STATUS ROLES  AGE   VERSION
cn-hangzhou.i-bp17enxc2us3624wexh2 Ready <none> 238d v1.10.4
cn-hangzhou.i-bp1ad2b02jtqd1shi2ut Ready <none> 220d v1.10.4
```

3. Check whether the number of servers whose heartbeat status is OK in the machine group is equal to the number of worker nodes in the cluster. Troubleshoot the error based on the check result.
 - o The number of servers whose heartbeat status is OK is equal to the number of worker nodes. This means that the heartbeat status of all of the servers in the machine group is **Failed**.
 - If [Logtail is installed into a common container](#), check whether the values of the `your_region_name`, `your_aliyun_user_id`, and `your_machine_group_user_defined_id` parameters are correct. For information about how to set these parameters, see [Collect standard Docker logs](#).
 - If [Logtail is installed into a container in a Container Service for Kubernetes cluster](#), submit a ticket.
 - If [Logtail is installed into a container in a user-created Kubernetes cluster](#), check whether the values of the `your-project-suffix`, `regionId`, `aliuid`, `access-kev-id`, and `access-kev-secret` parameters are correct. If the value of a parameter is incorrect, run the `helm del --purge alibaba-log-controller` command to delete the installation package and re-install Logtail. For information about how to set these parameters, see [Collect Kubernetes logs](#).
 - o The number of servers whose heartbeat status is OK is less than the number of worker nodes.
 - a. Check whether you used a YAML file to manually deploy a DaemonSet.

Run the `kubectl get po -n kube-system -l k8s-app=logtail` command to perform the check. If the command returns pod information, it indicates that you manually deployed a DaemonSet by using a YAML file.
 - b. Download the latest version of the [Logtail DaemonSet template](#).
 - c. Set the `your_region_name`, `your_aliyun_user_id`, and `your_machine_group_name` parameters to the values that are specific to your environment.
 - d. Run the `kubectl apply -f ./logtail-daemonset.yaml` command to update the DaemonSet YAML file.

Submit a ticket if the error persists.

Troubleshoot an error if Log Service collects no logs from containers

If no log is displayed in the **Consumption Preview** pane or on the **Search & Analysis** page of a Logstore, it indicates that Log Service does not collect logs from the machine group of the Logstore. Check the status of the containers that correspond to the servers in the machine group. If the containers are working as expected, perform the following steps to troubleshoot the error:

1. [Check the heartbeat status of the servers in the machine group](#).
2. Check whether the parameter settings in the Logtail configuration files are correct.

Check whether the values of the `IncludeLabel`, `ExcludeLabel`, `IncludeEnv`, and `ExcludeEnv` parameters in the Logtail collection configuration files meet your requirements.

Note The `IncludeLabel` or `ExcludeLabel` parameter specifies whether to include or exclude the container images to which specified labels are attached. You can retrieve a list of container image labels by running the `docker inspect` command. The labels are not the labels that are defined by using Kubernetes. To check whether the parameter settings are correct in a Logtail configuration file, delete the `IncludeLabel`, `ExcludeLabel`, `IncludeEnv`, and `ExcludeEnv` parameters from the file. If Log Service can collect logs from the containers after the parameters are deleted, it indicates that the settings of the parameters are incorrect.

3. Check other items.

Log Service does not collect logs from containers in the following scenarios:

- Log files are not updated.
- The log files of a container are stored in locations that are neither the default storage nor a storage attached to the container.

Log on to a Logtail container

• Common container

- i. Run the `docker ps | grep logtail` command on the host to search for the Logtail container.
- ii. Run the `docker exec -it ***** bash` command to log on to the container.

```
$docker ps | grep logtail
223fbd3ed2a6e registry.cn-hangzhou.aliyuncs.com/log-service/logtail          "/usr/local/ilogta..." 8 days ago
Up 8 days          logtail-iba
$docker exec -it 223fbd3ed2a6e bash
```

• Container in a Kubernetes cluster

- i. Run the `kubectl get po -n kube-system | grep logtail` command to search for the pod where the Logtail container resides.
- ii. Run the `kubectl exec -it -n kube-system ***** bash` command to log on to the pod.

```
$kubectl get po -n kube-system | grep logtail
logtail-ds-g5wgd          1/1 Running 0    8d
logtail-ds-slpn8         1/1 Running 0    8d
$kubectl exec -it -n kube-system logtail-ds-g5wgd bash
```

View the operational logs of Logtail

The operational logs of Logtail are saved in the files named `ilogtail.LOG` and `logtail_plugin.LOG` under the `/usr/local/ilogtail/` directory of a Logtail container.

1. [Log on to a Logtail container.](#)
2. Open the `/usr/local/ilogtail/` directory.

```
cd /usr/local/ilogtail
```

3. View the `ilogtail.LOG` and `logtail_plugin.LOG` files.

```
cat ilogtail.LOG
cat logtail_plugin.LOG
```

Ignore the stdout logs of a Logtail container

The standard output (stdout) logs of a Logtail container are useless for troubleshooting. Ignore the following stdout logs:

```

start umount useless mount points, /shm$/merged$/mqueue$
umount: /logtail_host/var/lib/docker/overlay2/3fd0043af174cb0273c3c7869500f2b2db95d13b1e110172ef57fe840c82155/merged: must be superuser to unmount
umount: /logtail_host/var/lib/docker/overlay2/d5b10aa19399992755de1f85d25009528daa749c1bf8c16edff44beab6e69718/merged: must be superuser to unmount
umount: /logtail_host/var/lib/docker/overlay2/5c3125daddacedec29df72ad0c52fac800cd56c6e880dc4e8a640b1e16c22dbe/merged: must be superuser to unmount
.....
xargs: umount: exited with status 255; aborting
umount done
start logtail
ilogtail is running
logtail status:
ilogtail is running

```

View the status of Log Service components in a Kubernetes cluster

Run the `helm status alibaba-log-controller` command to view the status of Log Service components in a Kubernetes cluster.

View the version number, IP address, and startup time of Logtail

View the information in the `app_info.json` file under the `/usr/local/ilogtail/` directory of the Logtail container. For example, you can run the following command to view the content of the file:

```

kubect exec logtail-ds-gb92k -n kube-system cat /usr/local/ilogtail/app_info.json
{
  "UUID": "",
  "hostname": "logtail-gb92k",
  "instance_id": "0EBB2B0E-0A3B-11E8-B0CE-0A58AC140402_10.10.10_1517810940",
  "ip": "10.10.10.10",
  "logtail_version": "0.16.2",
  "os": "Linux; 3.10.0-693.2.2.el7.x86_64; #1 SMP Tue Sep 12 22:26:13 UTC 2017; x86_64",
  "update_time": "2018-02-05 06:09:01"
}

```

23.9.2. Log search and analysis

23.9.2.1. FAQ about log query

This topic describes the common issues that may occur when you query log data in the Log Service console. It also includes solutions to these issues.

How do I identify the source server from which Logtail collects logs during a query?

If a server group uses IP addresses as its identifier when logs are collected by using Logtail, servers in the server group are distinguished from one another by internal IP addresses. When querying logs, you can use the hostname and custom IP address to identify the source server from which logs are collected.

For example, you can use the following statement to count the times different hostnames appear in logs:

 **Note** You must enable the index feature for the Logstore and enable the statistics feature for the `__tag__:__hostname__` field in advance.

```
* | select "__tag__:__hostname__", count(1) as count group by "__tag__:__hostname__"
```

How do I query IP addresses in logs?

You can use the exact match method to query IP addresses in logs. For example, you can specify IP addresses to query log data that includes or excludes the specified IP addresses. However, you cannot use the partial match method to query log data related to specified IP addresses. This is because decimal points contained in an IP address are not default delimiters in Log Service. If you want to use the partial match method, you can configure the decimal point as a delimiter for indexes. For example, you can use the SDK to download data and then use regular expressions or the `string.indexOf` method in the code.

For example, you use the following statement to query projects that meet the specified conditions.

```
not ip:121.42.0 not status:200 not 360jk not DNSPod-Monitor not status:302 not jiankongbao  
not 301 and status:403
```

The retrieved log data still contains 121.42.0.x. An IP address such as 121.42.0.x is taken as a word in Log Service. To include or filter out 121.42.0.x in the query result, you must specify 121.42.0.x in the query statement. If you specify 121.42.0 in the query statement, you cannot retrieve log data that includes or excludes the keyword.

How do I query log data by using a keyword that contains a whitespace character?

If you use a keyword that contains a whitespace character to query log data, log data that contains the part of the keyword on the left or right of the whitespace character is retrieved. You can enclose the keyword that contains a whitespace character in double quotation marks (""). Then the entire enclosed content is regarded as a keyword to query the log data that you expect.

For example, you want to query log data that contains the keyword `POS version` from the following log data:

```
post():351]:&nbsp;device_id:&nbsp;BTAddr&nbsp;:&nbsp;B6:xF:xx:65:xx:A1&nbsp;IMEI&nbsp;:&nbsp;35847xx22xx81x9&  
&nbsp;WifiAddr&nbsp;:&nbsp;4c:xx:0e:xx:4e:xx&nbsp;|&nbsp;user_id:&nbsp;bb07263xxd2axx43xx9exxea26e39e5f&nbsp;P  
OS&nbsp;version:903
```

If you use `POS version` as the keyword, log data that contains `POS` and `version` is retrieved. This result does not meet your expectations. If you use `"POS version"` as the keyword, log data that contains the keyword `POS version` is retrieved.

How do I use two query conditions to query log data?

You can enter two query conditions at one time to query log data that you want.

For example, you want to query log data whose status field value is neither OK nor Unknown in a Logstore. You can use the `not OK not Unknown` statement to retrieve the expected result.

How can I query collected logs in Log Service?

You can use one of the following methods to query logs in Log Service:

- Use the Log Service console.
- Use the SDK.
- Use the Restful API.

23.9.2.2. What can I do if no log data is retrieved?

When you use the log search and analytics feature of Log Service to query data, you may not retrieve the data you want. In this case, you can troubleshoot the problem as follows:

Log collection failure

If log data fails to be collected by Log Service, the target log data cannot be queried. Check whether log data is available on the consumption preview page of the target Logstore.

If data is available, log data is collected by Log Service.

If data is unavailable, possible causes are as follows:

- The log source does not generate log data.

In this case, no logs can be sent to Log Service. Check your log source.

- Logtail has no heartbeat.

On the **Server Group Settings** page, check whether the relevant server has a heartbeat in the Server Group Status section. If it has no heartbeat, troubleshoot the problem. For more information, see [What can I do if no heartbeat packet is received from a Logtail client?](#)

- The monitoring file is not written in real time.

If the monitoring file is written in real time, you can open the `/usr/local/ilogtail/ilogtail.LOG` file to view the error message. Common error messages are as follows:

- parse delimiter log fail: The error message is returned because an error has occurred when Log Service collects logs in the delimiter mode.
- parse regex log fail: The error message is returned because an error has occurred when Log Service collects logs in the regex mode.

Delimiter setting errors

View the configured delimiters. Check whether you can use a keyword to query log data after the log content is split by using a delimiter. For example, the default delimiters `; = () [] { } ? @ & < > / : '` are used. If a log entry contains `abc"defg,hij`, it is split into `abc"defg` and `hij`. In this case, you cannot retrieve this log entry by searching for `abc`.

Fuzzy match is also supported. For more information, see [Query syntax](#).

Note

- To save your indexing cost, Log Service has optimized the index feature. If you configure an index for a field, full-text indexing is ineffective for the key of this field. For example, an index is configured for a log field whose key is `message`, and a whitespace character is used as a delimiter. To use a whitespace character as a delimiter, put it in the middle of delimiters that you have configured for an index. You can retrieve the log entries that contain `"message: this is a test message"` by searching for the key-value-pair-formatted keyword `message:this`. However, if you use the keyword `this` to query the log entries, you cannot retrieve the data because an index is configured for the `message` field and full-text indexing is ineffective.
- You can create indexes or modify existing indexes. However, new or modified indexes take effect only for new data.

You can click **Index Attributes** to check whether the configured delimiters meet the requirements.

Other reasons

If log data is available, modify the time range of the query and try again. Log Service allows you to preview log data in real time. Due to a maximum latency of one minute, we recommend that you query log data at least one minute after logs are generated.

23.9.2.3. What are the differences between log consumption and log search and analytics?

Both the log consumption and log search and analytics features provided by Log Service need to read log data. The log consumption feature provides log collection and distribution channels. In contrast, the log search and analytics feature allows you to query log data.

Both the log consumption and log search and analytics features need to read log data:

Log collection and consumption (LogHub): provides public channels for log collection and distribution. It reads and writes full data in first-in, first-out (FIFO) order, which is similar to Kafka.

- Each Logstore has one or more shards. Data is written to a random shard.
- You can read multiple log entries at a time from a specified shard based on the order in which the log entries were written to the shard.
- You can set the start position (cursor) to pull log entries from shards based on the time when Log Service receives the log entries.

Log search and analytics: enables you to set conditions to search and analyze large amounts of log data based on LogHub.

- This feature allows you to search for required data based on query conditions.
- This feature allows you to include a combination of Boolean keywords AND, NOT, and OR and SQL statements in a query.
- This feature is independent of shards.

The following table lists the differences between the log search and analytics feature and the LogHub feature.

| Feature | Log search and analytics (LogSearch) | LogHub |
|------------------------------------|--|---|
| Search by keyword | Supported. | Not supported. |
| Data read (a small amount of data) | Fast. | Fast. |
| Data read (full data) | Slow. LogSearch reads 100 log entries in 100 milliseconds. This method is not recommended. | Fast. LogHub reads 1 MB of log data in 10 milliseconds. This method is recommended. |
| Data read by topic | Yes. | No. Data is identified only by shard. |
| Data read by shard | No. Data in all shards of a Logstore is queried. | Yes. You need to specify a shard each time to read data. |
| Price | Relatively high. | Low. |
| Scenario | Monitoring, problem investigation, and analysis. | Full data processing scenarios, such as stream computing and batch processing. |

23.9.2.4. How do I resolve common errors returned in log data queries?

Common errors returned in log data queries are as follows:

line 1:44: Column 'my_key_field' cannot be resolved; please add the column in the index attribute

- Cause

The `my_key_field` key cannot be included in the query statement because it does not exist.

- Solution

In the upper-right corner of the Search & Analysis page, click Index Attributes to create an index for this field and enable the statistics feature for this field.

Column 'xxxxline' not in GROUP BY clause;please add the column in the index attribute

- Cause

You use the GROUP BY clause and include a non-aggregated field in a SELECT statement. However, this field is not specified in the GROUP BY clause. For example, the key1 field in the `select key1, avg(latency) group by key2` statement is not specified in the GROUP BY clause.

- Solution

An example correct statement is `select key1,avg(latency) group by key1,key2` .

sql query must follow search query,please read syntax doc

- Cause

You do not include a filtering condition in a query statement, for example, `select ip,count(*) group by ip` .

- Solution

An example correct statement is `*|select ip,count(*) group by ip` .

please read syntax document,and make sure all related fields are indexed. error after select .error detail:line 1:10: identifiers must not start with a digit; surround the identifier with double quotes

- Cause

The column name or variable name referenced in an SQL statement starts with a number and does not comply with the rules.

- Solution

Change the name so the name starts with a letter.

please read syntax document,and make sure all related fields are indexed. error after select .error detail:line 1:9: extraneous input " expecting

- Cause

Misspelled words exist in the query statement.

- Solution

Correct the misspelled words.

key (category) is not config as key value config,if symbol : is in your log,please wrap : with quotation mark "

- Cause

No index is configured for the category field. It cannot be used in a query statement.

- Solution

Configure an index for this field in the index attributes. For more information, see [Enable the index feature and configure indexes for a Logstore](#).

Query exceeded max memory size of 3GB

- Cause

The size of memory used by the current query exceeds 3 GB. The common cause is that a large number of values are still returned in the query result after you use the GROUP BY clause to remove duplicates.

- Solution

Reduce the number of keys specified in the GROUP BY clause.

ErrorType:ColumnNotExists.ErrorPosition,line:0,column:1.ErrorMessage:line 1:123: Column '__raw_log__' cannot be resolved; it seems __raw_log__ is wrapper by ";if __raw_log__ is a string ,not a key field, please use '__raw_log__'

- Cause

The `my_key_field` key cannot be included in the query statement because it does not exist.

- Solution

In the upper-right corner of the Search & Analysis page, click Index Attributes to create an index for this field and enable the statistics feature for this field.

23.9.2.5. Why data queries are inaccurate?

This topic describes the causes for inaccurate data queries. It also includes solutions to these issues.

When you search and analyze log data, the message **The results are inaccurate** may prompt in the console. This indicates that the returned result is inaccurate because some log data in a Logstore was not queried.

Possible causes include:

The time range for queries is excessive.

- Cause

The time range for a query is excessively wide, for example, three months or a year. In this case, Log Service cannot scan all log data generated within this time period for one query.

- Solution

Narrow down the query time range and perform multiple queries.

Query statements are complex.

- Cause

The query statement is exceedingly complex or contains multiple frequently used words. In this case, Log Service cannot scan all related log data or read the query results at one time.

- Solution

Narrow down the query scope and perform multiple queries.

The SQL computing needs to read an excessively large amount of data.

- Cause

The SQL computing needs to read an excessively large amount of data. In this case, query results are likely to become inaccurate. A maximum of 1 GB of data can be read from each shard. For example, if the SQL computing needs to read strings from multiple columns, which exceed the threshold data volume, inaccurate query results will be returned.

- Solution

Narrow down the query scope and perform multiple queries.

23.9.3. Alarm

23.9.3.1. FAQ about alerts

This topic describes the common issues that may occur when you configure alerts in the Log Service console. It also includes solutions to these issues.

How can I include the raw error log entries in the notification content?

- Issue

More than five error log entries were generated in the past five minutes, which triggered an alert. How can I include the raw error log entries in notifications that were sent when the alert was triggered?

- Solution

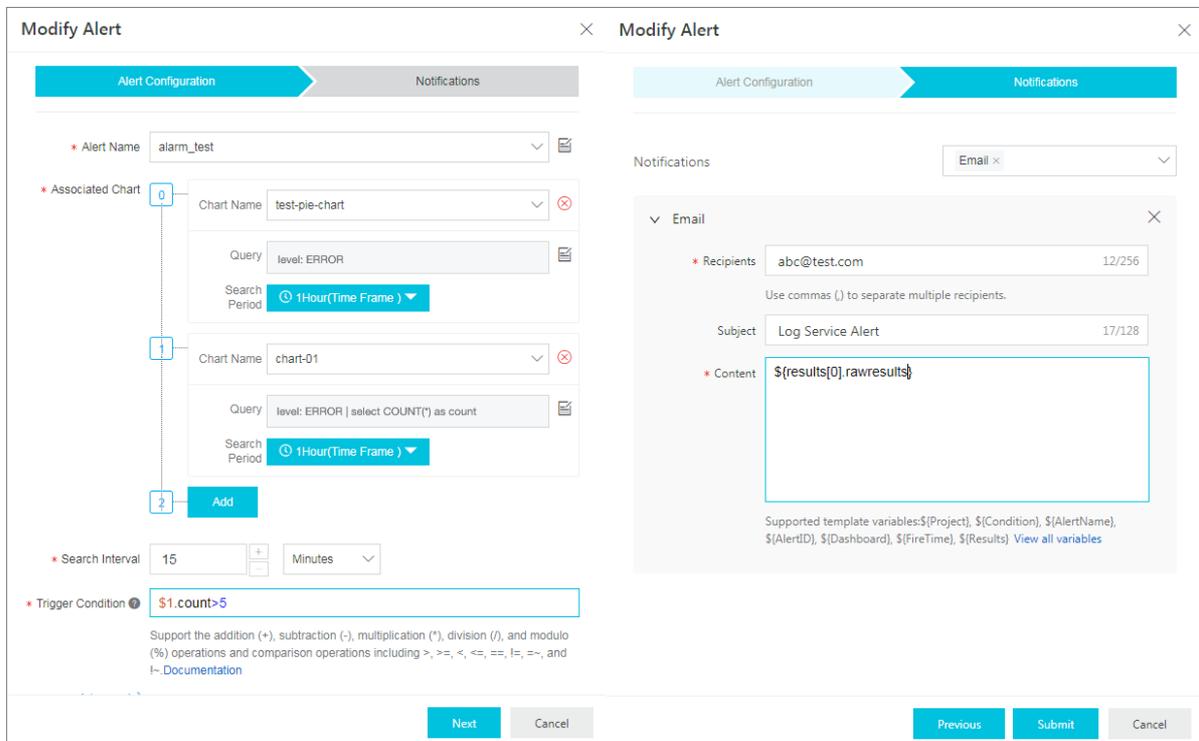
- Associated query statements

- Number 0: `level:ERROR`
- Number 1: `level:ERROR | select COUNT(*) as count`

- Trigger condition: `$.count > 5`

- Notification content: `${results[0].rawresults}`

- Configuration examples



24. Apsara Stack DNS

24.1. What is Apsara Stack DNS?

Apsara Stack DNS is a service that runs on Apsara Stack to resolve domain names. You can configure rules to map domain names to IP addresses. Apsara Stack DNS then distributes domain name requests from clients to cloud resources, business systems on your internal networks, or the business resources of Internet service providers (ISPs).

Apsara Stack DNS provides DNS resolution in VPCs. You can perform the following operations on your VPC by using Apsara Stack DNS:

- Access other ECS instances deployed in your VPC.
- Access cloud service instances provided by Apsara Stack.
- Access custom enterprise business systems.
- Access Internet services and businesses.
- Establish network connections between DNS and user-created DNS over a leased line.
- Manage internal domain names.
- Manage DNS records of internal domain names.
- Manage forwarding configurations.
- Manage recursive resolution configurations.

24.2. User roles and permissions

| Role | Permission |
|--|---|
| System administrator | A user of this role has read, write, and execute permissions on all level-1 organization resources, global resources, and system configurations. |
| Level-1 organization administrator | A user of this role has read, write, and execute permissions on level-1 organization resources to which the user belongs, but does not have permissions on level-1 organization resources, global resources, or system configurations to which other users belong. |
| Lower-level organization administrator | A user of this role does not have permissions on Apsara Stack DNS. The user does not have permissions on level-1 organization resources to which the user belongs, and does not have permissions on level-1 organization resources, global resources, or system configurations to which other users belong. |
| Resource user | A user of this role does not have permissions on Apsara Stack DNS. The user does not have permissions on level-1 organization resources to which the user belongs, and does not have permissions on level-1 organization resources, global resources, or system configurations to which other users belong. |
| Other roles | A user of this role does not have permissions on Apsara Stack DNS. The user does not have permissions on level-1 organization resources to which the user belongs, and does not have permissions on level-1 organization resources, global resources, or system configurations to which other users belong. |

24.3. Log on to the Apsara Stack DNS console

This topic describes how to log on to the Apsara Stack DNS console by using Google Chrome.

Prerequisites

- The URL of the Apsara Uni-manager Management Console is obtained from the deployment personnel before you log on to the Apsara Uni-manager Management Console.
- A browser is available. We recommend that you use the Google Chrome browser.

Procedure

1. In the address bar, enter the URL of the Apsara Uni-manager Management Console. Press the Enter key.
2. Enter your username and password.

Obtain the username and password that you can use to log on to the console from the operations administrator.

 **Note** When you log on to the Apsara Uni-manager Management Console for the first time, you must change the password of your username. Your password must meet complexity requirements. The password must be 8 to 20 characters in length and must contain at least two of the following character types:

- Uppercase or lowercase letters
- Digits
- Special characters, which include ! @ # \$ %

3. Click **Login**.
4. In the top navigation bar, choose **Products > Networking > Apsara Stack DNS**.

24.4. Internal DNS resolution management

Internal DNS resolution management allows you to manage global internal domain names, global forwarding configurations, and global recursive resolution configurations that you have created in Apsara Stack.

24.4.1. Global internal domain names

24.4.1.1. Overview

All the operations of this feature require administrator privileges.

24.4.1.2. View an internal domain name

Procedure

1. [Log on to the Apsara Stack DNS console](#).
2. In the left-side navigation pane, choose **Internal Domains > Global Internal Domains**.
3. In the **Domain Name** search box, enter the domain name that you want to view.
4. Click **Search**.

The search result is displayed.

24.4.1.3. Add a domain name

This topic describes how to add a domain name in the Apsara Uni-manager Management Console.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Internal Domain Names > Global Internal Domain Name**.
3. Click **Add Domain Name**.
4. In the dialog box that appears, enter **Global Internal Domain Name**.
5. Click **OK**.

24.4.1.4. Add a description for a domain name

This topic describes how to add a description for a domain name in the Apsara Uni-manager Management Console.

Context

You can add a description for a domain name to help you identify it. For example, you can add a hostname or internal system information to describe a domain name.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Internal Domain Names > Global Internal Domain Name**.
3. Find the domain name for which you want to add a description, click the  icon in the Actions column, and then select **Description**.
4. In the dialog box that appears, enter a description.
5. Click **OK**.

24.4.1.5. Delete a domain name

This topic describes how to delete a domain name in the Apsara Uni-manager Management Console.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Internal Domain Names > Global Internal Domain Name**.
3. Find the domain name that you want to delete, click the  icon in the **Actions** column, and then select **Delete**.
4. In the message that appears, click **OK**.

24.4.1.6. Delete multiple domain names

This topic describes how to delete unnecessary domain names at a time in the Apsara Uni-manager Management Console.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Internal Domain Names > Global Internal Domain Name**.
3. Select one or more domain names that you want to delete and click **Delete** in the upper-right corner.
4. In the message that appears, click **OK**.

24.4.1.7. Configure DNS records

This topic describes how to configure DNS records in the Apsara Uni-manager Management Console.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Internal Domain Names > Global Internal Domain Name**.
3. Find the domain name for which you want to configure DNS records, click the  icon in the Actions column, and then select **Configure DNS Records**.
4. On the **Configure DNS Records** page, click **Add DNS Record** in the upper-right corner.
5. Perform the following operations as needed:
 - Add a description for a DNS record
Select the DNS record for which you want to add a description, click  in the Actions column, and then select **Description** from the shortcut menu. In the dialog box that appears, enter a description and click **OK**.
 - Delete a DNS record
Select the DNS record that you want to delete, click  in the Actions column, and then select **Delete** from the shortcut menu. In the message that appears, click **OK**.
 - Modify a DNS record
Select the DNS record that you want to modify, click  in the Actions column, and then select **Modify** from the shortcut menu. In the dialog box that appears, set the required parameters and click **OK**.
 - Delete DNS records in batches
Select the DNS records that you want to delete and click **Delete** in the upper-right corner. In the message that appears, click **OK**.

24.4.1.8. View a resolution policy

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Internal Domains > Global Internal Domains**.
3. Find the target domain name, click the  icon in the Actions column, and then select **Configure DNS Records**.
4. On the page that appears, select the domain name for which you want to configure DNS records, and click **Weight** in the **Resolution Policy** column.
5. On the page that appears, view the details of **Resolution Policy**.

24.4.2. Global forwarding configurations

24.4.2.1. Global forwarding domain names

24.4.2.1.1. Overview

All operations of this feature require administrator privileges.

Apsara Stack DNS forwards specific domain names to other DNS servers for resolution.

Two forwarding modes are available: forward all requests without recursion and forward all requests with recursion.

- Forward all requests without recursion: Only a specified DNS server is used to resolve domain names. If the specified DNS server cannot resolve the domain names or the request times out, a message is returned to the DNS client to indicate that the query failed.
- Forward all requests with recursion: A specified DNS server is preferentially used to resolve domain names. If the specified DNS server cannot resolve the domain names, the local DNS is used instead.

24.4.2.1.2. View global forwarding domain names

This topic describes how to view global forwarding domain names in the Apsara Uni-manager Management Console. This operation requires administrator permissions.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Internal Domain Names > Forwarding Settings > Global Forwarding Domain Names**.
3. In the **Domain Name** search box, enter the domain name that you want to query and click **Search**.

24.4.2.1.3. Add a domain name

This topic describes how to add a domain name in the Apsara Uni-manager Management Console. This operation requires administrator permissions.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Internal Domain Names > Forwarding Settings > Global Forwarding Domain Names**.
3. Click **Add Domain Name**.
4. In the dialog box that appears, configure *Global Forwarding Domain*, *Forwarding Mode*, and *Forwarder IP Addresses*. Then, click **OK**.

24.4.2.1.4. Add a description for a domain name

This topic describes how to add a description for a domain name in the Apsara Uni-manager Management Console. This operation requires administrator permissions.

Context

You can add a description for a domain name to help you identify it. For example, you can describe a domain name by using a host name or internal system information.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Internal Domain Names > Forwarding Settings > Global Forwarding Domain Names**.

3. Select the domain name for which you want to add a description, click  in the Actions column, and then select **Description**.
4. In the dialog box that appears, enter a description and click **OK**.

24.4.2.1.5. Modify the forwarding configurations of a domain name

This topic describes how to modify the forwarding configurations of a domain name in the Apsara Uni-manager Management Console. This operation requires administrator permissions.

Procedure

1. [Log on to the Apsara Stack DNS console](#).
2. In the left-side navigation pane, choose **Internal Domain Names > Forwarding Settings > Global Forwarding Domain Names**.
3. Find the domain name whose forwarding configurations you want to modify, click the  icon in the Actions column, and then select **Modify**.
4. In the dialog box that appears, change the value of *Forwarding Mode* or *Forwarder IP Addresses*, and click **OK**.

24.4.2.1.6. Delete a domain name

This topic describes how to delete a domain name in the Apsara Uni-manager Management Console. This operation requires administrator permissions.

Procedure

1. [Log on to the Apsara Stack DNS console](#).
2. In the left-side navigation pane, choose **Internal Domain Names > Forwarding Settings > Global Forwarding Domain Names**.
3. Find the domain name that you want to delete, click the  icon in the Actions column, and then select **Delete**.
4. Click **OK**.

24.4.2.1.7. Delete multiple domain names

This topic describes how to delete multiple domain names at a time in the Apsara Uni-manager Management Console. This operation requires administrator permissions.

Procedure

1. [Log on to the Apsara Stack DNS console](#).
2. In the left-side navigation pane, choose **Internal Domain Domains > Forwarding Settings > Global Forwarding Domain Names**.
3. Select one or more domain names that you want to delete and click **Delete** in the upper-right corner.
4. In the message that appears, click **OK**.

24.4.2.2. Global default forwarding configurations

24.4.2.2.1. Enable default forwarding

This topic describes how to enable default forwarding in the Apsara Uni-manager Management Console. This operation requires administrator permissions.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Internal Domain Names > Forwarding Settings > Global Default Forwarding**.
3. Click the  icon in the Actions column and select **Enable**.
4. In the dialog box that appears, configure *Default Forwarding Mode* and *Forwarder IP Addresses*. Then, click **OK**. Make sure that **Enable Default Forwarding** is set to **ON**.

24.4.2.2.2. Modify default forwarding configurations

This topic describes how to modify default forwarding configurations in the Apsara Uni-manager Management Console. This operation requires administrator permissions.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Internal Domain Names > Forwarding Settings > Global Default Forwarding**.
3. Click the  icon in the Actions column and select **Modify**.
4. In the dialog box that appears, configure *Forwarding Mode* and *Forwarder IP Addresses*. Then, click **OK**.

24.4.2.2.3. Disable default forwarding

This topic describes how to disable default forwarding in the Apsara Uni-manager Management Console. This operation requires administrator permissions.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Internal Domain Names > Forwarding Settings > Global Default Forwarding**.
3. Click the  icon in the Actions column and select **Disable**.
4. In the message that appears, click **OK**.

24.4.3. Global recursive resolution

24.4.3.1. Enable global recursive resolution

Prerequisites

You have administrator permissions.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. Choose **Internal Domains** > **Global Recursive Resolution**.
3. Click the  icon in the Actions column and select **Enable**.
4. In the dialog box that appears, click **OK**.

24.4.3.2. Disable global recursive resolution

Prerequisites

You have administrator permissions.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. Choose **Internal Domains** > **Global Recursive Resolution**.
3. Click the  icon in the Actions column and select **Disable**.
4. In the dialog box that appears, click **OK**.

24.5. PrivateZone (DNS Standard Edition only)

The PrivateZone feature allows you to create VPC-specific tenant domain names. You can bind the domain names to VPCs as required to achieve tenant isolation.

24.5.1. Tenant internal domain name

24.5.1.1. View a domain name

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations** > **Tenant Internal Domains**.
3. In the **Domain Name** search box, enter the domain name that you want to view.
4. Click **Search**.

The search result is displayed.

24.5.1.2. Add a domain name

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations** > **Tenant Internal Domains**.
3. Click **Add Domain Name**.
4. In the dialog box that appears, set *Tenant Internal Domain Name*.
5. Click **OK**.

24.5.1.3. Bind an organization to a VPC

Tenant domain names are isolated based on VPCs. To ensure that the DNS forwarding configurations take effect, you must bind the organization of domain names to a VPC.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Internal Domains**.
3. Find the target domain name, click the  icon in the Actions column, and select **Associate VPCs**.
4. Select one or more VPCs from the list of VPCs to Select, click the right arrow to add them to the list of VPCs Selected, and then click **OK**.

24.5.1.4. Unbind a domain name from a VPC

This topic describes how to unbind a domain name from a VPC.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Internal Domains**.
3. Find the target domain name and click the number in the **VPCs Associated** column.
4. On the **VPCs Associated** page, find the target VPC, click the  icon in the **Actions** column, and then select **Disassociate**.
Make sure that the unbound VPC is no longer displayed on the **VPCs Associated** page.

24.5.1.5. Add a description for a domain name

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Internal Domains**.
3. Find the target domain name, click the  icon in the Actions column, and then select **Description**.
4. In the dialog box that appears, enter a description.
5. Click **OK**.

24.5.1.6. Delete a domain name

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Internal Domains**.
3. Find the target domain name, click the  icon in the Actions column, and then select **Delete**.
4. In the message that appears, click **OK**.

24.5.1.7. Delete multiple domain names

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Internal Domains**.
3. Select one or more domain names that you want to delete and click **Delete** in the upper-right corner.
4. In the message that appears, click **OK**.

24.5.1.8. Configure DNS records

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Internal Domains**.
3. Find the target domain name, click the  icon in the Actions column, and then select **Configure DNS Records**.
4. In the upper-right corner of the **Configure DNS Records** page, click **Add DNS Record**.
5. In the **Add DNS Record** dialog box, configure *Host*, *Type*, *TTL*, *Resolution Policy*, and *Record Set*. Then, click **OK**.

The following tables describe the types of DNS records.

- o A record

| Resolution policy | Formatting rule |
|-------------------|---|
| None | <p>You can enter up to 100 unique IPv4 addresses, each in a separate row.</p> <p>Make sure that the IPv4 addresses are valid.</p> <p>Example:</p> <ul style="list-style-type: none"> ▪ 192.168.1.1 ▪ 192.168.1.2 ▪ 192.168.1.3 |
| Weight | <p>You can enter up to 100 unique IPv4 addresses, each in a separate row.</p> <p>Format:</p> <ul style="list-style-type: none"> ▪ [IPv4 address] [Weight] (The IPv4 address and weight are separated with a space.) ▪ Make sure that the IPv4 addresses are valid. ▪ The weight value is an integer ranging from 0 to 999. A larger value indicates a greater weight. <p>Example:</p> <ul style="list-style-type: none"> ▪ 192.168.1.1 20 ▪ 192.168.1.1 30 ▪ 192.168.1.1 50 |

- o AAAA record

| Resolution policy | Formatting rule |
|-------------------|-----------------|
|-------------------|-----------------|

| Resolution policy | Formatting rule |
|-------------------|---|
| None | <p>You can enter up to 100 unique IPv6 addresses, each in a separate row.</p> <p>Make sure that the IPv6 addresses are valid.</p> <p>Example:</p> <ul style="list-style-type: none"> 2400:3200::6666 2400:3200::6688 2400:3200::8888 |
| Weight | <p>You can enter up to 100 unique IPv6 addresses, each in a separate row.</p> <p>Format:</p> <ul style="list-style-type: none"> [IPv6 address] [Weight] (The IPv6 address and weight are separated with a space.) Make sure that the IPv6 addresses are valid. The weight value is an integer ranging from 0 to 999. A larger value indicates a greater weight. <p>Example:</p> <ul style="list-style-type: none"> 2400:3200::6666 20 2400:3200::6688 20 2400:3200::8888 60 |

o CNAME record

| Resolution policy | Formatting rule |
|-------------------|---|
| None | <p>You can enter only one domain name.</p> <p>The domain name must be a fully qualified domain name (FQDN) that ends with a dot (.). It must be 1 to 255 characters in length.</p> <p>Example: www.example.com.</p> |
| Weight | <p>You can enter up to 100 unique domain names, each in a separate row.</p> <p>Format:</p> <ul style="list-style-type: none"> [Domain name] [Weight] (The domain name and weight are separated with a space.) The domain name must be an FQDN that ends with a period (.). It must be 1 to 255 characters in length. The weight value is an integer ranging from 0 to 999. A larger value indicates a greater weight. <p>Example:</p> <ul style="list-style-type: none"> www1.example.com. 20 www2.example.com. 20 www3.example.com. 60 |

o MX record

| Resolution policy | Formatting rule |
|-------------------|-----------------|
|-------------------|-----------------|

| Resolution policy | Formatting rule |
|-------------------|--|
| None | <p>You can enter 100 unique email server hostnames, each in a separate row.</p> <p>Format:</p> <ul style="list-style-type: none"> ▪ [Priority] [Email server hostname] (The priority and hostname are separated with a space.) ▪ The priority value is an integer ranging from 0 to 999. A smaller value indicates a higher priority. ▪ The email server hostname must be an FQDN that ends with a period (.). It must be 1 to 255 characters in length. <p>Example:</p> <ul style="list-style-type: none"> ▪ 10 mailserver1.example.com. ▪ 20 mailserver2.example.com. |

o TXT record

| Resolution policy | Formatting rule |
|-------------------|---|
| None | <p>You can enter up to 100 unique character strings, each in a separate row.</p> <p>A string must be 1 to 255 characters in length. No row can be left blank.</p> <p>Example: "v=spf1 ip4:192.168.0.1/16 ip6:2001::1/96 ~all"</p> |

o PTR record

| Resolution policy | Formatting rule |
|-------------------|--|
| None | <p>You can enter up to 100 unique domain names, each in a separate row.</p> <p>The DNS server address must be an FQDN that ends with a period (.). It must be 1 to 255 characters in length.</p> <p>Example:</p> <ul style="list-style-type: none"> ▪ www1.example.com. ▪ www2.example.com. ▪ www3.example.com. |

o SRV record

| Resolution policy | Formatting rule |
|-------------------|-----------------|
| | |

| Resolution policy | Formatting rule |
|-------------------|---|
| None | <p>You can enter up to 100 unique application server hostnames, each in a separate row.</p> <p>Format:</p> <ul style="list-style-type: none"> ▪ [Priority] [Weight] [Port number] [Application server hostname] (Every two items are separated with a space.) ▪ The priority value is an integer ranging from 0 to 999. A smaller value indicates a higher priority. ▪ The weight value is an integer ranging from 0 to 999. A larger value indicates a greater weight. ▪ The port number is an integer ranging from 0 to 65535. It indicates the TCP or UDP port used for network communications. ▪ The application server hostname must be an FQDN that ends with a period (.). It must be 1 to 255 characters in length. <p>Example:</p> <ul style="list-style-type: none"> ▪ 1 10 8080 www1.example.com. ▪ 2 20 8081 www2.example.com. |

o NAPTR record

| Resolution policy | Formatting rule |
|-------------------|---|
| None | <p>You can enter up to 100 unique NAPTR record values, each in a separate row.</p> <p>Format:</p> <ul style="list-style-type: none"> ▪ [Serial number] [Priority] [Flag] [Service information] [Regular expression] [Substitute domain name] (Every two items are separated with a space.) ▪ The serial number is an integer ranging from 0 to 999. A smaller value indicates a higher priority. ▪ The priority value is an integer ranging from 0 to 999. A smaller value indicates a higher priority. If two records have the same serial number, the one with a higher priority takes effect first. ▪ The flag value can be left blank or be a character from A to Z, a to z, or 0 to 9. It is not case-sensitive and must be enclosed in double quotation marks (""). ▪ The service information can be left blank or be a string of 1 to 32 characters. It must start with a letter and be enclosed in double quotation marks (""). ▪ The regular expression can be left blank or be a string of 1 to 255 characters enclosed in double quotation marks (""). ▪ The substitute domain name must be an FQDN that ends with a period (.). It must be 1 to 255 characters in length. <p>Example:</p> <ul style="list-style-type: none"> ▪ 100 50 "S" "Z3950+I2L+I2C" "" _z3950._tcp.example.com. ▪ 100 50 "S" "RCDS+I2C" "" _rcds._udp.example.com. ▪ 100 50 "S" "HTTP+I2L+I2C+I2R" "" _http._tcp.example.com. |

o CAA record

| Resolution policy | Formatting rule |
|-------------------|-----------------|
| | |

| Resolution policy | Formatting rule |
|-------------------|--|
| None | <p>You can enter up to 100 unique CAA records, each in a separate row.</p> <p>Format:</p> <ul style="list-style-type: none"> ▪ [Certificate authority flag] [Certificate property tag] [Authorization information] (Every two items are separated with a space.) ▪ The certification authority flag is an integer ranging from 0 to 255. ▪ The certificate property tag can be issue, issuewild, or iodef. ▪ The authorization information must be 1 to 255 characters in length and enclosed in double quotation marks (""). <p>Example:</p> <ul style="list-style-type: none"> ▪ 0 issue "caa.example.com" ▪ 0 issuewild ";" ▪ 0 iodef "mailto:example@example.com" |

o NS record

| Resolution policy | Formatting rule |
|-------------------|--|
| None | <p>You can enter up to 100 unique DNS server addresses, each in a separate row.</p> <p>The DNS server address must be an FQDN that ends with a period (.). It must be 1 to 255 characters in length. Wildcard domain names are not allowed.</p> <p>Example:</p> <ul style="list-style-type: none"> ▪ ns1.example.com. ▪ ns2.example.com. |

6. After you add DNS records, perform the following operations as required:

- o Add a description for a DNS record.

Find the target DNS record, click the  icon in the Actions column, and then select **Description**. In the dialog box that appears, enter a description and click **OK**.

- o Delete a DNS record.

Find the target DNS record, click the  icon in the Actions column, and then select **Delete**. In the message that appears, click **OK**.

- o Modify a DNS record.

Find the target DNS record, click the  icon in the Actions column, and then select **Modify**. In the dialog box that appears, modify the required parameters and click **OK**.

- o Delete multiple DNS records.

Select the DNS records that you want to modify and click **Delete** in the upper-right corner. In the message that appears, click **OK**.

24.5.1.9. View a resolution policy

This topic describes how to view the details of a resolution policy.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Internal Domains**.
3. Find the target domain name, click the  icon in the Actions column, and then select **Configure DNS Records**.
4. View the resolution policy in the DNS Records list.

24.5.2. Tenant forwarding configurations

24.5.2.1. Tenant forwarding domain names

24.5.2.1.1. View a tenant forwarding domain name

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Forwarding Domains**.
3. In the **Domain Name** search box, enter the domain name that you want to view.
4. Click **Search**.

The search result is displayed.

24.5.2.1.2. Add a tenant forwarding domain name

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Forwarding Domains**.
3. Click **Add Domain Name**.
4. In the dialog box that appears, configure parameters such as *Domain Name*, *Forwarding Mode*, and *Forwarder IP Addresses*.

| Parameter | Description |
|-------------|---|
| Domain Name | <p>The domain name, which must meet the following formatting rules:</p> <ul style="list-style-type: none"> ◦ The domain name must be 1 to 255 characters in length. This includes the period (.) at the end of the domain name. ◦ The domain name can contain multiple domain name segments that are separated with periods (.). A domain name segment must be 1 to 63 characters in length. It cannot contain consecutive periods (.) or be left blank. ◦ The domain name can only contain letters (a to z, A to Z), digits (0 to 9), hyphens (-), and underscores (_). ◦ The domain name must start with a letter, digit, or underscore (_) and end with a letter, digit, or period (.) ◦ The domain name is not case-sensitive. The system saves the domain name in lowercase letters. ◦ The period (.) at the end of the domain name is optional. The system adds a period (.) to the end of the domain name. |

| Parameter | Description |
|------------------------|---|
| Forwarding Mode | <p>For both domain name-based forwarding and default forwarding, the following two forwarding modes are supported:</p> <ul style="list-style-type: none"> Forward All Requests without Recursion: forwards DNS requests to the target DNS server. If the target DNS server cannot resolve the domain names, a message is returned to the DNS client indicating that the query failed. Forward All Requests with Recursion: A specified DNS server is preferentially used to resolve domain names. If the specified DNS server cannot resolve the domain names, the local DNS is used instead. If you enter internal IP addresses in the Forwarder IP Addresses field, unexpected results may occur during recursive resolution. For example, a domain name used for internal network services may be resolved to a public IP address. |
| Forwarder IP Addresses | <p>A list of destination IP addresses.</p> <div style="background-color: #e6f2ff; padding: 5px; border: 1px solid #d9e1f2;"> <p> Note Multiple IP addresses are separated with semicolons (;).</p> </div> |

5. Click OK.

24.5.2.1.3. Bind an organization to a VPC

Tenant domain names are isolated based on VPCs. You must bind the organization of domain names to a VPC before the DNS forwarding configurations can take effect.

Procedure

- Log on to the [Apsara Stack DNS console](#).
- In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Forwarding Domains**.
- Find the target domain name, click the  icon in the Actions column, and then select **Associate VPCs**.
- Select one or more VPCs from the list of VPCs to Select, click the right arrow to add them to the list of VPCs Selected, and then click OK.

24.5.2.1.4. Unbind a domain name from a VPC

This topic describes how to unbind a domain name from a VPC.

Procedure

- Log on to the [Apsara Stack DNS console](#).
- In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Forwarding Domains**.
- Find the target domain name and click the number in the **VPCs Associated** column.
- On the VPCs Associated page, find the target VPC, click the  icon in the Actions column, and then select **Disassociate**.
Make sure that the unbound VPC is no longer displayed on the VPCs Associated page.

24.5.2.1.5. Modify the forwarding configurations of a domain name

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Forwarding Domains.**
3. Find the target domain name, click the  icon in the Actions column, and then select **Modify**.
4. In the dialog box that appears, change the value of **Forwarding Mode** or **Forwarder IP Addresses**.
5. Click **OK**.

24.5.2.1.6. Add a description for a tenant forwarding domain name

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Forwarding Domains.**
3. Find the target domain name, click the  icon in the Actions column, and then select **Description**.
4. In the dialog box that appears, enter a description.
5. Click **OK**.

24.5.2.1.7. Delete a tenant forwarding domain name

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Forwarding Domains.**
3. Find the target domain name, click the  icon in the Actions column, and then select **Delete**.
4. In the message that appears, click **OK**.

24.5.2.1.8. Delete multiple tenant forwarding domain names

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Forwarding Domains.**
3. Select one or more domain names that you want to delete and click **Delete** in the upper-right corner.
4. In the message that appears, click **OK**.

24.5.2.2. Tenant default forwarding configurations

24.5.2.2.1. View default forwarding configurations

Prerequisites

You have the permissions of a system administrator or level-1 organization administrator.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Default Forwarding.**

24.5.2.2.2. Add a default forwarding configuration

Prerequisites

You have the permissions of a system administrator or level-1 organization administrator.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Default Forwarding.**
3. Click **Add Settings.**
4. In the dialog box that appears, configure parameters such as *Forwarding Mode* and *Forwarder IP Addresses.*

| Parameter | Description |
|------------------------|--|
| Forwarding Mode | <p>For both domain name-based forwarding and default forwarding, the following two forwarding modes are available:</p> <ul style="list-style-type: none">◦ Forward All Requests without Recursion: Only a specified DNS server is used to resolve domain names. If the specified DNS server cannot resolve the domain names, a message is returned to the DNS client to indicate that the query failed.◦ Forward All Requests with Recursion: A specified DNS server is preferentially used to resolve domain names. If the specified DNS server cannot resolve the domain names, a local DNS server is used instead. If you enter internal IP addresses in the Forwarder IP Addresses field, unexpected results may occur during recursive resolution. For example, a domain name used for internal network services may be resolved to a public IP address. |
| Forwarder IP Addresses | <p>A list of destination IP addresses.</p> <div style="border: 1px solid #add8e6; padding: 5px;"><p> Note Multiple IP addresses are separated with semicolons (;).</p></div> |

5. Click **OK.**

24.5.2.2.3. Bind an organization to a VPC

Tenant domain names are isolated based on VPCs. You must bind the organization of domain names to a VPC before the DNS forwarding configurations can take effect.

Prerequisites

You have the permissions of a system administrator or level-1 organization administrator.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Default Forwarding.**
3. Find the target organization, click the  icon in the Actions column, and then select **Associate VPCs.**
4. Select one or more VPCs from the list of VPCs to Select, click the right arrow to add them to the list of VPCs Selected, and then click OK.

24.5.2.2.4. Unbind a domain name from a VPC

This topic describes how to unbind a domain name from a VPC.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Default Forwarding.**
3. Find the target domain name and click the number in the **VPCs Associated** column.
4. On the VPCs Associated page, find the target VPC, click the  icon in the Actions column, and then select **Disassociate.**
Make sure that the unbound VPC is no longer displayed on the VPCs Associated page.

24.5.2.2.5. Modify a default forwarding configuration

Prerequisites

You have the permissions of a system administrator or level-1 organization administrator.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Default Forwarding.**
3. Find the target organization, click the  icon in the Actions column, and then select **Modify.**
4. In the dialog box that appears, change the value of **Forwarding Mode** or **Forwarder IP Addresses.**
5. Click OK.

24.5.2.2.6. Add a default forwarding configuration

Prerequisites

You have the permissions of a system administrator or level-1 organization administrator.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Default Forwarding.**
3. Find the target organization, click the  icon in the Actions column, and then select **Description.**
4. In the dialog box that appears, enter **Description.**
5. Click **OK.**

24.5.2.2.7. Delete a default forwarding configuration

Prerequisites

You have the permissions of a system administrator or level-1 organization administrator.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Default Forwarding.**
3. Find the target organization, click the  icon in the Actions column, and then select **Delete.**
4. In the dialog box that appears, click **OK.**

24.5.2.2.8. Delete multiple default forwarding configurations

Prerequisites

You have the permissions of a system administrator or level-1 organization administrator.

Procedure

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, choose **Forwarding Configurations > Tenant Forwarding Settings > Tenant Default Forwarding.**
3. Select one or more domain names that you want to delete and click **Delete** in the upper-right corner.
4. In the message that appears, click **OK.**

24.6. Internal Global Traffic Manager (internal GTM Standard Edition only)

Internal Global Traffic Manager (GTM) supports multi-cloud disaster recovery for domain names of customers. This feature manages traffic loads between multiple Apsara Stack networks.

24.6.1. Scheduling instance management

24.6.1.1. Scheduling Instance

The Scheduling Instance tab displays all existing scheduling instances. You can add, delete, modify, and configure scheduling instances on this tab. When you create a scheduling instance, you must associate an address pool and scheduling domain with the instance.

24.6.1.1.1. Create a scheduling instance

After you create a scheduling instance, you can associate the scheduling instance with a scheduling domain and address pool.

1. In the left-side navigation pane, choose **Internal Global Traffic Manager > Scheduling Instances > Scheduling Instance**.
2. Click **Create Scheduling Instance** in the upper-right corner of the instance list.
3. In the dialog box that appears, configure Scheduling Instance Name, CNAME Access Domain Name, and Global TTL. Then, click **OK**.

24.6.1.1.2. Modify a scheduling instance

1. In the left-side navigation pane, choose **Internal Global Traffic Manager > Scheduling Instances > Scheduling Instance**.
2. Find the instance that you want to modify and click **Modify** in the Actions column.
3. Modify the parameter settings as prompted and click **OK**.

24.6.1.1.3. Configure a scheduling instance

You can create, delete, modify, and query access policies of scheduling instances.

1. In the left-side navigation pane, click **Internal Global Traffic Manager**. On the **Scheduling Instances** tab of the page that appears, click the **Scheduling Instances** tab.
2. Find the scheduling instance whose access policies you want to view and click **Configure** in the **Actions** column.
3. On the Access Policy Configuration page, view information about all the access policies of the scheduling instance. The information includes **Access Policy Name**, **DNS Request Source**, **Address Type**, **Effective Address Pool**, and **Last Modified At**.
4. Click the closing angle bracket (>) next to an access policy to view the details, including information about the primary and secondary address pools.
5. View the setting of **Address Pool Switchover Policy**. The default value is **Automatic**. You can change the value to **Manual**. If Address Pool Switchover Policy is set to Automatic, the system automatically selects an available address pool. If Address Pool Switchover Policy is set to Manual, you must manually specify whether to use the primary address pool or secondary address pool.

Note

Whether an address pool is available is determined based on the number of normal addresses in the address pool and Min. Number of Available Addresses that you specified when you configure the access policy. If the number of normal addresses in the address pool is less than the value of Min. Number of Available Addresses, the address pool is considered unavailable. You can perform a health check to obtain the number of normal addresses in the address pool.

Processing logic for automatic switchover

| State of the primary address pool | State of the secondary address pool | Comparison between the numbers of normal addresses in the primary and secondary address pools | Effective address pool (list of available addresses) |
|-----------------------------------|-------------------------------------|---|---|
| Available | Available | - | Primary address pool (including the normal addresses that are intelligently returned and the addresses that are always online. Abnormal addresses that are intelligently returned are deleted or their weight values are set to 0.) |
| Available | Unavailable | - | Primary address pool (including the normal addresses that are intelligently returned and the addresses that are always online. Abnormal addresses that are intelligently returned are deleted or their weight values are set to 0.) |
| Unavailable | Available | - | Secondary address pool (including the normal addresses that are intelligently returned and the addresses that are always online. Abnormal addresses that are intelligently returned are deleted or their weight values are set to 0.) |
| Unavailable | Unavailable | Number of normal addresses in the primary address pool > Number of normal addresses in the secondary address pool | Primary address pool (including the normal addresses that are intelligently returned and the addresses that are always online. Abnormal addresses that are intelligently returned are deleted or their weight values are set to 0.) |
| Unavailable | Unavailable | Number of normal addresses in the primary address pool < Number of normal addresses in the secondary address pool | Secondary address pool (including the normal addresses that are intelligently returned and the addresses that are always online. Abnormal addresses that are intelligently returned are deleted or their weight values are set to 0.) |

| | | | |
|-------------|-------------|---|---|
| Unavailable | Unavailable | Number of normal addresses in the primary address pool = Number of normal addresses in the secondary address pool > 0 | Primary address pool (including the normal addresses that are intelligently returned and the addresses that are always online. Abnormal addresses that are intelligently returned are deleted or their weight values are set to 0.) |
| Unavailable | Unavailable | Number of normal addresses in the primary address pool = Number of normal addresses in the secondary address pool = 0 | If the DNS request source is a custom line, the system clears the DNS configurations of the line instead of selecting an address pool. The configurations of the custom line are deleted. If the DNS request source is the global default line, the system selects the primary address pool. The system returns all the configured addresses without considering their return mode. |

When the system compares the numbers of normal addresses between the primary and secondary address pools, normal addresses include normal addresses that are intelligently returned and all addresses that are always online (with the health status ignored). The abnormal addresses that are intelligently returned and all addresses that are always off line (with the health status ignored) are not normal addresses.

The following table describes the processing logic for address types if a line is selected in two access policies.

| Scenario | Address type for the effective address pool in two access policies | | Processing logic |
|-------------------------------------|--|-------------|---|
| Same DNS request source: Scenario 1 | IPv4 | IPv6 | IPv4 and IPv6 addresses take effect at the same time. |
| Same DNS request source: Scenario 2 | IPv4 | Domain name | Addresses of the domain name type take effect. |
| Same DNS request source: Scenario 3 | IPv6 | Domain name | Addresses of the domain name type take effect. |

24.6.1.1.3.1. Create an access policy for a scheduling instance

You can create multiple access policies to resolve different address pools based on different DNS request sources.

1. Log on to the Apsara Stack DNS console. In the left-side navigation pane, click Internal Global Traffic Manager. On the Scheduling Instances tab of the Global Traffic Management on the Internal Network page, find the scheduling instance for which you want to create an access policy and click Configure in the Actions column. On the page that appears, click **Create Access Policy**.
2. In the Create Access Policy dialog box, specify **Access Policy Name**, select items in the **DNS Request Source** section, and then configure parameters in the **Primary/Secondary Address Pool Configuration** section. In the DNS Request Source section, if you select Global default, you cannot select other items. The parameters on the Primary Address Pool tab must be configured, and the parameters on the Secondary Address Pool tab can be left empty.

3. In the **Primary/Secondary Address Pool Configuration** section, you can set **Address Type** to IPv4, IPv6, or Domain Name to select different types of address pools. You can also set **Min. Number of Available Addresses**. If the number of healthy addresses in an address pool is less than the value of this parameter, the address pool is determined to be unavailable. The value of Min. Number of Available Addresses must be an integer ranging from 1 to 100.

4. Click **OK**.

The following table describes the limits on address type conflicts between the primary and secondary address pools for two access policies that have the same DNS request source.

| Scenario | Address type of the primary address pool (access policy 1) | Address type of the secondary address pool (access policy 2) | Processing logic |
|-------------------------------------|--|--|--|
| Same DNS request source: Scenario 1 | IPv4 | IPv4 | The address pools are allowed to be added. |
| Same DNS request source: Scenario 2 | IPv4 | IPv6 | The address pools are not allowed to be added. |
| Same DNS request source: Scenario 3 | IPv4 | Domain name | The address pools are allowed to be added. |
| Same DNS request source: Scenario 4 | IPv6 | IPv4 | The address pools are not allowed to be added. |
| Same DNS request source: Scenario 5 | IPv6 | IPv6 | The address pools are allowed to be added. |
| Same DNS request source: Scenario 6 | IPv6 | Domain name | The address pools are allowed to be added. |
| Same DNS request source: Scenario 7 | Domain name | IPv6 | The address pools are allowed to be added. |
| Same DNS request source: Scenario 8 | Domain name | IPv4 | The address pools are allowed to be added. |
| Same DNS request source: Scenario 9 | Domain name | Domain name | The address pools are allowed to be added. |

The following tables describe the limits on address type conflicts between the primary address pools and those between the secondary address pools for two access policies that have the same DNS request source.

| Scenario | Address type of the primary address pool (access policy 1) | Address type of the secondary address pool (access policy 2) | Processing logic |
|-------------------------------------|--|--|--|
| Same DNS request source: Scenario 1 | IPv4 | IPv6 | The address pools are allowed to be added and they can coexist. |
| Same DNS request source: Scenario 2 | IPv4 | IPv4 | The address pools are not allowed to be added and they cannot coexist. |
| Same DNS request source: Scenario 3 | IPv4 | Domain name | The address pools are not allowed to be added and they cannot coexist. |

| | | | |
|--|-------------|-------------|--|
| Same DNS request source: Scenario 4 | IPv6 | IPv6 | The address pools are not allowed to be added and they cannot coexist. |
| Same DNS request source: Scenario 5 | Domain name | IPv6 | The address pools are not allowed to be added and they cannot coexist. |
| Same DNS request source: Scenario 6 | Domain name | Domain name | The address pools are allowed to be added and they can coexist. However, the following conditions must be met: (1) The two primary address pools are the same. (2) Both of the secondary address pools exist. In addition, one secondary address pool is of the IPv4 type and the other is of the IPv6 type. |

| Scenario | Address type of the secondary address pool (access policy 1) | Address type of the secondary address pool (access policy 2) | Processing logic |
|--|--|--|--|
| Same DNS request source: Scenario 1 | IPv4 | IPv6 | The address pools are allowed to be added and they can coexist. |
| Same DNS request source: Scenario 2 | IPv4 | IPv4 | The address pools are not allowed to be added and they cannot coexist. |
| Same DNS request source: Scenario 3 | IPv4 | Domain name | The address pools are not allowed to be added and they cannot coexist. |
| Same DNS request source: Scenario 4 | IPv6 | IPv6 | The address pools are not allowed to be added and they cannot coexist. |
| Same DNS request source: Scenario 5 | Domain name | IPv6 | The address pools are not allowed to be added and they cannot coexist. |
| Same DNS request source: Scenario 6 | Domain name | Domain name | The address pools are allowed to be added and they can coexist. However, the following conditions must be met: (1) The two secondary address pools are the same. (2) Both of the two primary address pools exist. In addition, one primary address pool is of the IPv4 type and the other is of the IPv6 type. |

24.6.1.1.3.2. Modify the access policy of a scheduling instance

Notice

If you do not change the primary or secondary address pool when you modify an access policy, no primary/secondary switchover is triggered. If you change one of the address pools, the system complies with the following processing rules:

1. Manual switchover mode: If the secondary address pool is deleted, the system forcibly switches services to the primary address pool. If the secondary address pool is not deleted, the effective address pool does not change.
2. Automatic switchover mode: The system determines the address pool that takes effect based on the status of the newly selected address pools. Exercise caution when you change the address pools.

1. On the Access Policy Configuration page, find the access policy that you want to modify and click **Modify** in the Actions column.
2. In the Modify Access Policy dialog box, modify **Access Policy Name**, select items in the **DNS Request Source** section, and then configure parameters in the **Primary/Secondary Address Pool Configuration** section. In the DNS Request Source section, if you select Global default, you cannot select other items. The parameters on the Primary Address Pool tab must be configured, and the parameters on the Secondary Address Pool tab can be left empty.
3. In the Primary/Secondary Address Pool Configuration section, you can set **Address Type** to IPv4, IPv6, or Domain Name to select address pools of different types. You can also set Min. Number of Available Addresses. If the number of healthy addresses in an address pool is less than the value of this parameter, the address pool is determined to be unavailable. The value of this parameter must be an integer ranging from 1 to 100.
4. Click **OK**.

24.6.1.1.3.3. Delete the access policy of a scheduling instance

1. On the Access Policy Configuration page, find the target access policy and click Delete in the Actions column.
2. In the dialog box that appears, click OK after you verify that the displayed information is correct.

24.6.1.1.4. Delete a scheduling instance

1. Log on to the Apsara Stack DNS console and choose Recursion Configurations > Scheduling Instances > Scheduling Instance.
2. Find the target instance and click Delete in the Actions column.
3. In the dialog box that appears, click OK.

Note: After you delete the instance, its configuration data is also deleted.

24.6.1.2. Address Pool

The Address Pool tab allows you to manage address pools. You can associate address pools with scheduling instances. The address pools are classified into three types: IPv4 address pools, IPv6 address pools, and domain name address pools. The load balancing policy of an address pool can be set to polling or weight.

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, click **Internal Global Traffic Manager**. On the **Scheduling Instance** tab of the Global Traffic Management on the Internal Network page, click the **Address Pool** tab.
3. Find the address pool whose information you want to view and click the closing angle bracket (>) next to the name of the address pool to view its detailed information. The information includes Address Pool ID, Address Pool Name, Address Type, Load Balancing Policy (Among Addresses), Created At, Last Modified At, Health Check, and Health Check Status.

24.6.1.2.1. Create an address pool

You can define a list of addresses that form an address pool, which can be associated with access policies of scheduling instances when you configure the access policies.

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, click **Internal Global Traffic Manager**. On the Scheduling Instance tab of the Global Traffic Management on the Internal Network page, click the Address Pool tab.
3. On the Address Pool tab, click **Create Address Pool**.
4. In the Create Address Pool dialog box, specify **Address Pool Name**, **Address Type**, and **Load Balancing Policy (Among Addresses)**, and add addresses one by one in the **Address List** section. You can also click **Batch Add** to add multiple addresses at a time. After you enter the required information, click **OK**.

| Parameter | Description |
|---|--|
| Address Pool Name | The name can contain a maximum of 20 characters. |
| Address Type | You can select IPv4, IPv6, or Domain Name from the drop-down list of this parameter. This configuration cannot be changed. |
| Load Balancing Policy (Among Addresses) | You can select Polling or Weight from the drop-down list of this parameter. This configuration cannot be changed. |
| Mode | Valid values: <ul style="list-style-type: none"> • Automatically Returned: The system determines whether the address is available based on the health check result of the address. • Always online: The system ignores the health check result of the address and sets the address to be always available. The health check task is still running. • Always Offline: The system ignores the health check result of the address and sets the address to be always unavailable. The health check task is still running. |

24.6.1.2.2. Modify the configurations of an address pool

 **Notice** After you modify the configurations of an address pool, the health check results of all addresses are reset to normal if health check is enabled. If the address pool has been associated with access policies and automatic switchover is enabled, a primary/secondary switchover may be triggered. Proceed with caution.

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, click **Internal Global Traffic Manager**. On the **Scheduling Instance** tab of the Global Traffic Management on the Internal Network page, click the **Address Pool** tab.
3. Find the address pool that you want to modify and click **Modify** in the Actions column.
4. In the Modify Address Pool dialog box, you can change only **Address Pool Name** and the addresses in **Address (One in Each Row)**.
5. Click **OK**.

24.6.1.2.3. Delete an address pool

1. Log on to the Apsara Stack DNS console and choose Recursion Configurations > Scheduling Instances > Address Pool.
2. Find the target address pool and click **Delete** in the Actions column.

3. In the dialog box that appears, click OK after you verify that the displayed information is correct.

24.6.1.2.4. Enable health check

You can enable health check to check the status of the addresses in an address pool. Only the addresses whose health check status is normal can be returned.

1. [Log on to the Apsara Stack DNS console.](#)
2. In the left-side navigation pane, click **Internal Global Traffic Manager**. On the **Scheduling Instance** tab of the Global Traffic Management on the Internal Network page, click the **Address Pool** tab.
3. Find the address pool for which you want to configure a health check, and click **Health Check** in the Actions column.
4. In the Health Check dialog box, turn on or turn off **Health Check Switch**, specify the parameters in the **Protocol Settings**, **Health Check Settings**, and **Node Settings** sections, and then click **OK**.

| Section | Parameter | Description | Protocol |
|-------------------|--|--|--------------------|
| Protocol Settings | Port | The port number that is used for health checks on the destination address. The value must be an integer in the range of 1 to 65535. This parameter cannot be empty. | HTTP/HTTPS/TCP/UDP |
| | Path | The HTTP or HTTPS path that is used for health checks on the destination address. This path is used to check whether the HTTP or HTTPS service of the destination address is normal. If the HTTP status code returned from this path is 2xx or 3xx, the HTTP or HTTPS service is normal. The system automatically adds a forward slash (/) before the path name. The path can be empty. The default value is /. The path name can contain a maximum of 255 characters. | HTTP/HTTPS |
| | Host Configuration | The host configuration that is used for health checks. If you do not specify this parameter, the primary domain name is used. | HTTP/HTTPS |
| | Returned Error Code Greater Than or Equal to | The minimum value of the HTTP status code when the health check result is abnormal. The HTTP status code returned must be greater than or equal to the value of this parameter. | HTTP/HTTPS |
| | ICMP Packages Sent | The number of ICMP packets sent each time an Internet Control Message Protocol (ICMP) health check is performed. | ICMP |
| | Packet Loss Rate | The threshold of the packet loss rate. The threshold is used to determine whether the result of an ICMP health check is abnormal. | ICMP |

| Section | Parameter | Description | Protocol |
|-----------------------|-------------------|---|----------|
| Health Check Settings | Check interval | The time interval at which health checks are performed on the destination address. | |
| | Timeout Duration | The timeout duration for which the system waits after an exception occurs during a health check. | |
| | Failure Threshold | The minimum number of consecutive health check failures when the status of the destination address is abnormal during a health check. | |
| Node Settings | Selected node | The nodes that initiate health checks on the destination address. | |

24.6.1.3. Scheduling Domain

The Scheduling Domain tab allows you to add, delete, and query scheduling domains.

You can log on to the Apsara Stack DNS console and choose Internal Global Traffic Manager > Scheduling Instances > Scheduling Domain to go to the scheduling domain list.

24.6.1.3.1. Create a scheduling domain

1. Log on to the Apsara Stack DNS console and choose Recursion Configurations > Scheduling Instances > Scheduling Domain. Then, click Create Scheduling Domain in the upper-right corner of the scheduling domain list.
2. In the dialog box that appears, enter the custom domain name and click OK.

24.6.1.3.2. Add a description for a scheduling domain

1. In the left-side navigation pane, choose Internal Global Traffic Manager > Scheduling Instances > Scheduling Domain.
2. Find the scheduling domain for which you want to add a description and click Edit in the Actions column.
3. In the dialog box that appears, add a description in the Edit field and click OK.

24.6.1.3.3. Delete a scheduling domain

1. In the left-side navigation pane, choose Internal Global Traffic Manager > Scheduling Instances > Scheduling Domain.
2. Find the scheduling domain that you want to delete and click Delete in the Actions column.
3. In the message that appears, click OK after you verify that the displayed information is correct.

24.6.1.4. View alert logs

Note

By default, the system saves alert logs of the last 90 days.

1. [Log on to the Apsara Stack DNS console.](#)

2. In the left-side navigation pane, click **Internal Global Traffic Manager**. On the **Scheduling Instance** tab of the Global Traffic Management on the Internal Network page, click the **Alert Logs** tab.
3. In the alert log list, view address pool information, such as the status of the address pool, health check results of addresses, and switchover between primary and secondary address pools. You can query alert logs of address pools by time or behavior, or by using a keyword.

24.6.2. Scheduling line management

24.6.2.1. IP Address Line Configuration

The IP Address Line Configuration tab allows you to define lines based on IP addresses. The lines are used to group request sources to achieve intelligent load balancing.

24.6.2.1.1. Add a line

1. In the left-side navigation pane, choose **Internal Global Traffic Manager > Scheduling Line Management > IP Address Line Configuration**.
2. Click **Add Line** in the upper-right corner of the line list.
3. In the dialog box that appears, configure the parameters as prompted and click **OK**.

24.6.2.1.2. Sort lines

1. In the left-side navigation pane, choose **Internal Global Traffic Manager > Scheduling Line Management > IP Address Line Configuration**.
2. Find the line whose sequence you want to change and click **Sort** in the Actions column.
3. Specify Sort Behavior as prompted and click **OK**.

24.6.2.1.3. Modify the configurations of a line

1. In the left-side navigation pane, choose **Internal Global Traffic Manager > Scheduling Line Management > IP Address Line Configuration**.
2. Find the line whose configurations you want to modify and click **Modify** in the Actions column.
3. Modify the configurations as prompted and click **OK**.

24.6.2.1.4. Delete a line

1. In the left-side navigation pane, choose **Internal Global Traffic Manager > Scheduling Line Management > IP Address Line Configuration**.
2. Find the line that you want to delete and click **Delete** in the Actions column.
3. In the message that appears, click **OK**.

24.6.3. Data synchronization management

Data synchronization management is used to synchronize Global Traffic Manager (GTM) data between clouds.

24.6.3.1. Synchronization cluster management

Synchronization clusters involve two operations: **Set Emergency Group** and **Merge GTM Control Domain**.

You can perform the following operations to go to the synchronization cluster management page:

1. Log on to the Apsara Stack DNS console.

2. In the left-side navigation pane, choose **Internal Global Traffic Manager**.
3. On the page that appears, click the **Data Synchronization** tab.

Set Emergency Group

You can select some service instances to form a cluster to provide services.

- Enable the emergency group feature: If the synchronization cluster is abnormal, click **Set Emergency Group**. In the Set Emergency Group dialog box, turn on Emergency Group Switch, select available service instances to form an emergency group, and then click **OK**.
- Disable the emergency group feature: If the synchronization cluster is restored to normal, click **Set Emergency Group**. In the Set Emergency Group dialog box, turn off the Emergency Group Switch and click **OK**.

Merge GTM Control Domain

In multi-cloud scenarios, you can click **Merge GTM Control Domain** and enter the IP address of the leader service instance of the merged Global Traffic Manager (GTM) control domain to form a large synchronization cluster.

View the status of the synchronization cluster

You can view the status of the synchronization cluster on the Synchronization Cluster Management tab.

View the service instances in the synchronization cluster

You can view the following information of the service instances in the current synchronization cluster:

Instance IP Address, Instance Role, Status, Latest Synchronization Log ID, IP Address, and Instance Description.

You can also perform the following operations to switch the role of a service instance in the synchronization cluster from follower to leader:

1. Find the service instance with the follower role and click **Switch Primary** in the Actions column.
2. In the message that appears, click **OK**.