ALIBABA CLOUD

# Alibaba Cloud

## Apsara Stack Enterprise

### Security Whitepaper

Product Version: 2006, Internal: V3.12.0

Document Version: 20201020

Alibaba Cloud

# 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:<br>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:<br>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:<br>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:<br>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**. |
| `Courier font` | Courier font is used for commands | Run the `cd /d C:/window` command to enter the Windows system folder. |
| *Italic* | Italic formatting is used for parameters and variables. | `bae log list --instanceid`<br>*Instance_ID* |
| [] or [a\|b] | This format is used for an optional value, where only one item can be selected. | `ipconfig [-all\|-t]` |
| {} or {a\|b} | This format is used for a required value, where only one item can be selected. | `switch {active\|stand}` |

# Table of Contents

1.Overview 17
2.Attribution of security power and sharing of security respon... 18
2.1. Attribution of security power 18
2.2. Sharing of security responsibilities 18
2.2.1. Security responsibilities of Alibaba Cloud 18
2.2.2. Security responsibilities of users 18
3.Security compliance 20
3.1. Overview 20
3.2. Security compliance 21
3.3. Alibaba Apsara Stack classified protection 2.0 complianc... 22
4.Apsara Stack security architecture 23
4.1. Apsara Stack security architecture 23
4.2. Cloud platform security 24
4.2.1. Apsara Stack infrastructure security 24
4.2.2. Apsara Stack system security 25
4.2.2.1. System security of physical host 25
4.2.2.2. Virtualization system security 25
4.2.3. Distributed Apsara system security 26
4.2.3.1. Security of distributed file system 26
4.2.3.2. Security of remote process call module 27
4.2.3.3. Security of job scheduling module 27
4.2.3.4. Security of basic service module 27
4.2.4. Network security 27
4.2.4.1. Basic network security 27
4.2.4.2. Network device security 28
4.2.5. Application security 28

4.2.5.1. Secure Product Lifecycle (SPLC) 28
4.2.5.2. Web Application Firewall (on the platform side) 30
4.2.6. Data security 30
4.2.6.1. Data security system 30
4.2.6.2. Data ownership 30
4.2.6.3. Multi-copy redundancy storage 30
4.2.6.4. Full-stack encryption 30
4.2.6.5. Residual data cleanup 30
4.2.6.6. Operations data security 30
4.2.7. Account system security 31
4.2.7.1. Overview 31
4.2.7.2. Super administrator 31
4.2.7.3. Apsara Stack account 31
4.2.7.4. Identity credential 31
4.2.7.5. RAM 32
4.2.8. Operations security 33
4.2.8.1. Overview 33
4.2.8.2. OAM permission and authorization 33
4.2.8.3. Apsara Infrastructure Management Framework p... 34
4.3. Security operation service (on the platform side) 35
4.4. Cloud user (tenant) security 35
4.4.1. Network security 35
4.4.1.1. Virtual Private Cloud 35
4.4.1.2. Distributed firewall 36
4.4.1.3. Server Load Balancer 36
4.4.1.4. Apsara Stack Security - Traffic Security Monitoring 36
4.4.1.5. Apsara Stack Security - DDoS Traffic Scrubbing 36
4.4.1.6. Cloud Firewall 37

4.4.2. Host security 37
4.4.2.1. Operating system of ECS instances 37
4.4.2.2. Image reinforcement 37
4.4.2.3. Image snapshot 37
4.4.2.4. Apsara Stack Security - Server Guard 37
4.4.3. Application security 38
4.4.3.1. Code security 38
4.4.3.2. Apsara Stack Security - Web Application Firewall 38
4.4.4. Data security 38
4.4.4.1. ApsaraDB 39
4.4.4.2. MaxCompute 40
4.4.4.3. Sensitive Data Discovery and Protection of Apsar... 40
4.4.5. Security management 41
4.4.5.1. Apsara Stack Security - Threat Detection Service 41
4.4.6. Security operation service (on the tenant side) 41
4.4.7. Best security practices 41
5.Security of Apsara Stack products 43
5.1. Elastic Compute Service (ECS) 43
5.1.1. Platform security 43
5.1.1.1. Security isolation 43
5.1.1.2. Authentication 44
5.1.1.2.1. Identity authentication 44
5.1.1.2.2. Access control 44
5.1.1.3. Data security 45
5.1.1.3.1. Overview 45
5.1.1.3.2. Triplicate storage technology 45
5.1.1.3.3. ECS disk encryption 46
5.1.1.4. Encrypted transmission 47

5.1.1.5. ARP spoofing prevention 47
5.1.2. Tenant security 47
5.1.2.1. Log audit 47
5.1.2.2. Secure images 47
5.1.2.3. Block storage 48
5.2. Container Service for Kubernetes 48
5.2.1. Platform security 48
5.2.1.1. Security isolation 48
5.2.1.2. Account authentication 49
5.2.1.3. Link security 49
5.2.2. Tenant security 49
5.2.2.1. Application security 49
5.2.2.2. Host security 51
5.3. Auto Scaling (ESS) 52
5.3.1. Platform security 52
5.3.1.1. Security isolation 52
5.3.1.2. Authentication 52
5.3.1.2.1. Authentication 52
5.3.1.2.2. Access control 52
5.3.2. Tenant security 53
5.3.2.1. Log audit 53
5.4. Resource Orchestration Service (ROS) 53
5.4.1. Platform security 53
5.4.1.1. Data security 53
5.4.1.2. Authentication 53
5.4.2. Tenant security 53
5.4.2.1. Log audit 53
5.5. Object Storage Service (OSS) 53

5.5.1. Platform security 53
5.5.1.1. Security isolation 54
5.5.1.2. Authentication and access control 54
5.5.1.2.1. Authentication 54
5.5.1.2.2. Access control 54
5.5.1.2.3. Support for RAM and STS 55
5.5.1.3. Data security 55
5.5.1.4. Data encryption 55
5.5.1.4.1. Server-side encryption 56
5.5.1.4.2. Client-side encryption 56
5.5.2. Tenant security 56
5.5.2.1. Key management 56
5.5.2.2. Log audit 56
5.5.2.3. Configure hotlink protection 56
5.6. Apsara File Storage NAS 57
5.6.1. Platform security 57
5.6.1.1. Security isolation 57
5.6.1.2. Authentication 57
5.6.1.3. Data security 59
5.6.2. Tenant security 60
5.6.2.1. Log audit 60
5.6.2.2. Directory-level ACLs 60
5.7. Tablestore 61
5.7.1. Platform security 61
5.7.1.1. Security isolation 61
5.7.1.2. Authentication 62
5.7.1.3. Data security 62
5.7.2. Tenant security 63

5.7.2.1. Key management 63
5.8. ApsaraDB for RDS 63
5.8.1. Platform security 63
5.8.1.1. Secure isolation 63
5.8.1.2. Authentication 63
5.8.1.3. Data security 64
5.8.1.4. Data encryption 64
5.8.1.5. DDoS attack prevention 64
5.8.2. Tenant security 65
5.8.2.1. Log audit 65
5.8.2.2. IP address whitelist 65
5.8.2.3. Software update 66
5.9. Cloud Native Distributed Database PolarDB-X 66
5.9.1. Platform security 66
5.9.1.1. Security isolation 66
5.9.1.2. Authentication 66
5.9.2. Tenant security 67
5.9.2.1. IP address whitelist 67
5.9.2.2. Protection against high-risk SQL operations 67
5.9.2.3. Slow SQL audit 67
5.9.2.4. Performance monitoring 68
5.10. AnalyticDB for MySQL 68
5.10.1. Platform security 68
5.10.1.1. Security isolation 68
5.10.1.2. Authentication 68
5.10.1.3. Data security 70
5.10.2. Tenant security 70
5.10.2.1. Log audit 70

5.11. AnalyticDB for PostgreSQL 71
5.11.1. Platform security 71
5.11.1.1. Security isolation 71
5.11.1.2. Authentication 71
5.11.1.3. Primary and secondary nodes 72
5.11.2. Tenant security 72
5.11.2.1. Database account 72
5.11.2.2. IP address whitelists 72
5.11.2.3. SQL audit 72
5.11.2.4. Backup and restoration 72
5.11.2.5. Software upgrade 72
5.12. KVStore for Redis 73
5.12.1. Platform security protections 73
5.12.1.1. Security isolation 73
5.12.1.2. Authentication 73
5.12.1.3. Transmission encryption 74
5.12.2. Tenant security protections 74
5.12.2.1. Database account 74
5.12.2.2. IP address whitelist 74
5.12.2.3. Backup and recovery 74
5.12.2.4. Software upgrade 74
5.13. ApsaraDB for MongoDB 75
5.13.1. Platform security 75
5.13.1.1. Isolation 75
5.13.1.2. Authentication 75
5.13.1.3. Data security 76
5.13.1.4. Data encryption 76
5.13.1.5. Anti-DDoS 76

5.13.2. Tenant security 77
5.13.2.1. Log audit 77
5.13.2.2. IP address whitelists 77
5.14. ApsaraDB for OceanBase 77
5.14.1. Platform security 77
5.14.1.1. Security isolation 77
5.14.1.2. Authentication 77
5.14.1.3. High-availability architecture 78
5.14.1.4. Compatibility 78
5.14.2. Tenant security 79
5.14.2.1. Database accounts 79
5.14.2.2. IP address whitelists 79
5.14.2.3. Log audit 79
5.14.2.4. Software upgrades 79
5.14.3. Notes on MySQL-related vulnerabilities in ApsaraDB... 79
5.15. Data Transmission Service (DTS) 80
5.15.1. Platform security 80
5.15.1.1. Security isolation 81
5.15.1.2. Authentication 81
5.15.1.3. Transmission security 81
5.15.1.4. Data security 81
5.16. Data Management (DMS) 81
5.16.1. Platform security 81
5.16.1.1. Security isolation 81
5.16.1.2. Authentication 81
5.16.1.3. Transmission security 82
5.16.2. Tenant security 82
5.16.2.1. ActionTrail 82

5.17. Server Load Balancer (SLB) 82
5.17.1. Platform security 82
5.17.1.1. Authentication 82
5.17.2. Tenant security 82
5.17.2.1. HTTPS 82
5.17.2.2. IP address whitelists 83
5.17.2.3. Log management 83
5.18. Virtual Private Cloud (VPC) 83
5.18.1. Platform security 83
5.18.1.1. Security isolation 83
5.18.1.2. Access control 83
5.18.2. Tenant security 84
5.18.2.1. Security groups 84
5.19. Log Service 84
5.19.1. Platform security 84
5.19.1.1. Security isolation 84
5.19.1.2. Authentication 84
5.19.1.3. Data security 85
5.19.1.4. Encrypted data transmission 85
5.19.2. Tenant security 86
5.19.2.1. Service monitoring 86
5.20. Key Management Service (KMS) 86
5.20.1. Platform security 86
5.20.1.1. Security isolation 86
5.20.1.2. Authentication 86
5.20.1.2.1. Identity authentication 86
5.20.1.2.2. Access control 87
5.20.1.2.3. RAM and STS support 87

5.20.1.3. Data security 87
5.20.1.4. Transmission encryption 87
5.21. Apsara Stack DNS 88
5.21.1. Tenant security 88
5.21.1.1. Tenant isolation 88
5.21.1.2. Network security hardening 88
5.21.1.3. Log audit 88
5.22. API Gateway 88
5.22.1. Platform security 88
5.22.1.1. Security isolation 88
5.22.1.2. Authentication 88
5.22.1.2.1. Authentication 88
5.22.1.2.2. API access control 89
5.22.1.2.3. RAM and STS support 89
5.22.1.3. Data security 89
5.22.1.4. Transmission encryption 89
5.22.2. Tenant security 89
5.22.2.1. Log audit 90
5.22.2.2. IP address-based access control 90
5.23. Enterprise Distributed Application Service (EDAS) 90
5.23.1. Platform security 90
5.23.1.1. Authentication 90
5.23.1.2. Transmission encryption 91
5.23.2. Tenant security 92
5.24. MaxCompute 93
5.24.1. Platform security 93
5.24.1.1. Security isolation 93
5.24.1.2. Authentication and authorization 94

5.24.1.3. Data security 99
5.24.1.4. KMS-based storage encryption 100
5.24.1.5. Transmission encryption 103
5.24.2. Tenant security 103
5.24.2.1. Cross-project resource sharing 103
5.24.2.2. Data protection mechanism (Project Protection) 106
5.24.2.3. Log audit 108
5.24.2.4. IP address whitelists 109
5.25. DataWorks 111
5.25.1. Permission isolation for development and productio... 111
5.25.2. Authentication and authorization 112
5.25.2.1. Access control 112
5.25.2.2. Permission management 112
5.25.3. Data encryption 113
5.25.4. Sensitive data protection 113
5.26. Realtime Compute 114
5.26.1. Platform security 114
5.26.1.1. Resource isolation 114
5.26.1.2. Authentication and authorization 114
5.26.1.3. Data security 114
5.26.1.4. Business process 114
5.27. Machine Learning Platform for AI 115
5.27.1. Security isolation 115
5.27.2. Authentication 116
5.27.2.1. Identity authentication 116
5.27.2.2. Access control 117
5.27.2.3. RAM and STS 117
5.27.3. Data security 119

5.27.4. Log audit 119
5.28. E-MapReduce (EMR) 119
5.28.1. Platform security 119
5.28.1.1. Access control 119
5.28.1.2. Authentication 119
5.28.1.3. Data security 122
5.28.2. Authorization control 124
5.29. DataHub 125
5.29.1. Platform security 125
5.29.1.1. Data isolation 125
5.29.1.2. Authentication and Authorization 125
5.29.1.3. Data encryption 126
5.29.1.4. Data security 126
5.29.2. Tenant security 127
5.29.2.1. Audit logging 127
5.30. Graph Analytics 127
5.30.1. Platform security 127
5.30.1.1. Security isolation 127
5.30.1.2. Authentication 127
5.30.1.3. Data security 128
5.30.1.4. Transmission encryption 128
5.30.1.5. System security 128
5.30.1.5.1. Vulnerability scanning mechanism 128
5.30.1.5.2. Update scheme for security vulnerabilities 128
5.30.1.5.3. System defense mechanism 128
5.30.1.6. Infrastructure security 129
5.30.1.7. Level-based data security 129
5.30.2. Tenant security 129

5.30.2.1. Log audit 129
5.31. Elasticsearch (on k8s) 129
5.31.1. Security isolation 129
5.31.2. Authentication and authorization 129
5.31.2.1. Identity authentication 129
5.31.2.2. Access control 129
5.31.3. Data security 129
5.31.4. Transmission encryption 130
6.Apsara Stack Security 131
6.1. Overview 131
6.2. Apsara Stack Security Standard Edition 131
6.2.1. Threat Detection Service 131
6.2.2. Traffic Security Monitoring 133
6.2.3. Server Guard 134
6.2.4. WAF 135
6.2.5. Security Operations Center (SOC) 138
6.2.6. On-premises security operations services 139
6.3. Optional security services 139
6.3.1. DDoS Traffic Scrubbing 139
6.3.2. Cloud Firewall 140
6.3.3. Sensitive Data Discovery and Protection 142

# 1.Overview

Data security and user privacy are top priorities of Alibaba Cloud Apsara Stack. Alibaba Cloud is committed to providing a public, open, and secure Apsara Stack cloud computing service platform. With technical innovation, Apsara Stack is constantly improving its computing capability and economies of scale to turn cloud computing into the infrastructure of true sense.

Apsara Stack is designed to provide users with stable, reliable, secure, and compliant cloud computing basic services and protect the availability, confidentiality, and integrity of users' systems and data.

This document introduces the Apsara Stack security system in the following parts:

- Attribution of security power and responsibilities and security capacity co-construction
- Security compliance
- Security of the Apsara Stack platform architecture
- Security features provided by Apsara Stack products
- Security services provided by Apsara Stack Security

This document also provides the best practices for secure use of Apsara Stack products and Apsara Stack Security products, which helps users make better use of the Apsara Stack platform and get an insight into the overall environment of security control.

# 2.Attribution of security power and sharing of security responsibilities

## 2.1. Attribution of security power

The products, design, model algorithm, programs, and its relevant intellectual property provided by Alibaba Cloud in various types of Apsara Stack environment all belong to Alibaba Cloud unless the contract stipulates clearly otherwise. Users have the access rights within the time period authorized by License.

The national standard *GBT 31167-2014 Information Security Technology - Cloud Computing Service Security Guide* (this document puts forward the national standard of the security control solution for the government to use cloud service. It has four deployment forms of cloud computing including Apsara Stack (private cloud). Other customers can also regard this standard as a reference when using Apsara Stack service) stipulates that "the customer owns the data, device, and other resources that the customer submits to cloud provider. The customer also owns the data and document collected, produced, and stored by customer business system on the cloud computing platform. The right of the customer to visit, use, and dominate these resources must not be limited."

In Apsara Stack environment, users are entitled to the ownership of the user data of project planning and implementation, the operation data produced during operations, and the business data that are transferred to the cloud environment. Alibaba Cloud can access the data within the scope of users' authorization, and cloud users must avoid authorizing business data to Alibaba personnel.

## 2.2. Sharing of security responsibilities

### 2.2.1. Security responsibilities of Alibaba Cloud

In Apsara Stack environment, Alibaba Cloud is responsible to provide users with cloud computing products and solutions, help users with customized deployment, or facilitates operations within the scope of the contract. Alibaba Cloud takes the following responsibilities:

- Provides users with security testimonial for compliance requirement of Alibaba Cloud Apsara Stack.
- Provides the vulnerability recognition service and technology for Apsara Stack products and helps users fix Apsara Stack on the product side.
- Protects users' Apsara Stack information system or infrastructure and provides relevant solutions and techniques, including authorization management, encryption, and auditing feature. Based on the preceding solutions and methods, Alibaba Cloud provides best practices of security management for users and puts forward the security capacity building.
- According to Alibaba Cloud security regulation and customer's requirement, the Alibaba Cloud personnel is required to sign the confidential agreement and receive proper security training and education.

### 2.2.2. Security responsibilities of users

The national standard *GBT 31167-2014 Information Security Technology - Cloud Computing Service Security Guide* stipulates that "the responsibility of information security control must not transfer to the outsourcing service partner. No matter the customer data and business are placed in the customer internal information system or on the cloud computing platform of the cloud service provider, the customer holds the responsibility for the information security." In Apsara Stack environment, the user exercises the security control based on the security solutions and technology provided by Alibaba Cloud or the third party and takes the following responsibilities for the results:

- Establishes the security control personnel, organization, security system, and operation system, which all support the Apsara Stack environment. The control object includes relevant project members of Alibaba Cloud.
- Practices admittance examination, confidentiality agreement, security training and education for relevant project members of Alibaba Cloud in the Apsara Stack environment, according to the national law, regulation, and customer requirements.
- Executes the transfer control for code and program in the Apsara Stack environment and takes responsibility for data leakage that is caused by users' fault.
- Leads the vulnerability fix process for Apsara Stack products, reviews the relevant implementation plan, and authorizes the change of plan during the upgrade process.
- Implements the account assignment, authorization, and log audit in each console in the Apsara Stack environment. Manages to achieve minimized authorization and normalized audit.
- Implements the security configuration of each console and products in the Apsara Stack environment, or authorizes the Alibaba Cloud field personnel to perform the security configuration.
- Users must perform backup and recovery drills for key data on regular basis to guarantee the business data is backed up and can be restored.

# 3.Security compliance

## 3.1. Overview

The security process of Alibaba Cloud has been recognized by authorities inside and outside China. By using years of expertise in defense against Internet security threats of Alibaba Group, Alibaba Cloud provides security protection for the Apsara Stack platform and integrates multiple compliance standards into the internal control and product design of the cloud platform. Alibaba Cloud also participates in the development of standards for various cloud platforms and contributes the best practices. Certified by more than 10 agencies inside and outside China currently, Alibaba Cloud is a cloud service provider with the most complete scope of certifications in Asia.

Certified by more than 10 agencies in and outside China, Alibaba Cloud is the cloud service provider with the most complete range of certifications in Asia. Alibaba Cloud has certifications as listed inAlibaba Cloud has the following qualifications.

### Certifications awarded to Alibaba Cloud

| Certification | Description |
| --- | --- |
| ISO 27001 | The international Information Security Management System (ISMS) Certification. It certifies Alibaba Cloud for fully performing its security duties in regard to data security, network security, communication security, and operation security. |
| CSA STAR | The International Cloud Security Management System Certification. The certification organization awarded the first cloud security gold medal to Alibaba Cloud. |
| ISO 20000 | The IT Service Management System Certification. This certifies that Alibaba Cloud has established and strictly implemented a standard service process. The standardized cloud platform services can improve IT efficiency and reduce the overall IT risk. |
| ISO 22301 | The Business Continuity Management System Certification. This certifies that Alibaba Cloud meets the requirements for business continuity planning, disaster recovery, and regular drills to enhance the stability of the cloud platform. |

| Certification | Description |
| --- | --- |
| Classified protection (level 4) | Alibaba Cloud Apsara Stack platform complies with the security and technology capabilities that are requested by Cloud computing platform classified protection 2.0 compliance specifications (level 4), which is formulated in accordance with GB/T22239 - 2019 *Information security technology - Baseline for classified protection of cybersecurity* . |
| Cloud service capability standard test by Ministry of Industry and Information Technology (MIIT) | CNAS certification for cloud products is the only product-level classified certification based on national standards. |
| Service Organization Control (SOC) audit certification | Alibaba Cloud has passed SOC3 audit, and the TYPE II of SOC1 and SOC2. |

### List of domestic Apsara stack qualifications

| Qualifications/certification | Certification authority |
| --- | --- |
| ITSS cloud computing service capability (private cloud IaaS service/level 1) | Chinese Electronics Standardization Association |
| Trusted cloud - the protection of the user data of cloud service (private cloud) | China Academy of Information and Communications Technology |
| Security level assessment report of the information system of the Ministry of Public Security(level 4, private cloud) | Information security rating center of the Ministry of Public Security |
| Security classified protection evaluation report of the Ministry of Public Security Information System Apsara stack V3.0 | Information security rating center of the Ministry of Public Security |
| Security assessment report of big data simple Apsara Stack platform of the security of the information system of Ministry of Public Security | China Academy of Information and Communications Technology |
| Cloud evaluation certificate-Cloud computing reference architecture-cloud solution | China Electronics Standardization Institute |
| Trusted cloud-open-source solutions (agile private cloud version)/virtualization and virtualization management software | China Academy of Information and Communications Technology |

# 3.2. Security compliance

**Alibaba Cloud keeps improving its management and system based on relevant standards and best practices in the industry. It is certified in a series of standard certifications, third-party audits, and self-assessment, which aims to better demonstrate its compliance practices to users.**

The overall compliance architecture of Alibaba Cloud is divided into the following parts according to compliance requirements from different perspectives, industries, and regions:

### Management system compliance

These compliance authentications demonstrate the mature management system of Alibaba Cloud and the best industry practices that Alibaba Cloud complies with:

- ISO 27001: Information Security Management Standard
- ISO 20000: IT Service Management Standard
- ISO 22301: Business Continuity Management Standard
- CSA STAR: maturity model of cloud service security
- Classified protection (level four)
- CNAS test for cloud computing standards in China

### Systematized compliance reports

These compliance authentications demonstrate the integrity and effectiveness of control in Alibaba Cloud platform, including the continuous effectiveness of system control, accuracy of separation of duties, and completeness of operations audit.

SOC 1/2 TYPE II: The Service Organization Control (SOC) reports are a series of audit reports from independent third parties to indicate the continuous effectiveness of the key compliance control and objectives of Alibaba Cloud. These reports aim to help users and their auditors learn the control measures behind operation and compliance. The SOC reports that Alibaba Cloud has are categorized into the following three types:

- SOC 1 TYPE II: internal control report over financial reporting
- SOC 2 TYPE II: reports over security, availability, and confidentiality
- SOC 3: report over security, availability, and confidentiality

## 3.3. Alibaba Apsara Stack classified protection 2.0 compliance whitepaper

According to *Cloud Computing Security Classified Protection Compliance Capability Framework* , and under the guidance of China's technology community of the cloud computing security classified protection compliance capacity specification system, Information security rating center of the Ministry of Public Security and Alibaba Cloud Computing Co., LTD. jointly compiled and issued *Apsara Stack Network Security Classified Protection 2.0 Compliance Capability Whitepaper* . The whitepaper explains in details from the technical verification architecture of classified protection capability, the compliance status of Apsara stack classified protection 2.0 to the usage recommendations for the whitepaper. With this whitepaper, customers can quickly obtain compliance protection on the Apsara stack platform side in multiple delivery scenarios. It also integrates customer-side application, security management, and protection measures such as physical environment, to jointly construct an overall security defense system of information systems to meet the needs of classified protection and customers.

# 4.Apsara Stack security architecture

## 4.1. Apsara Stack security architecture

Apsara Stack is designed with an in-depth multi-layer security defense system and provides security assurance of architecture on the cloud platform side, including infrastructure, system, distributed Apsara system security, network, application, database, cloud platform account, O&M, and operation service, and also on the cloud users (tenants) side, including network, host, application, data, and operation service.

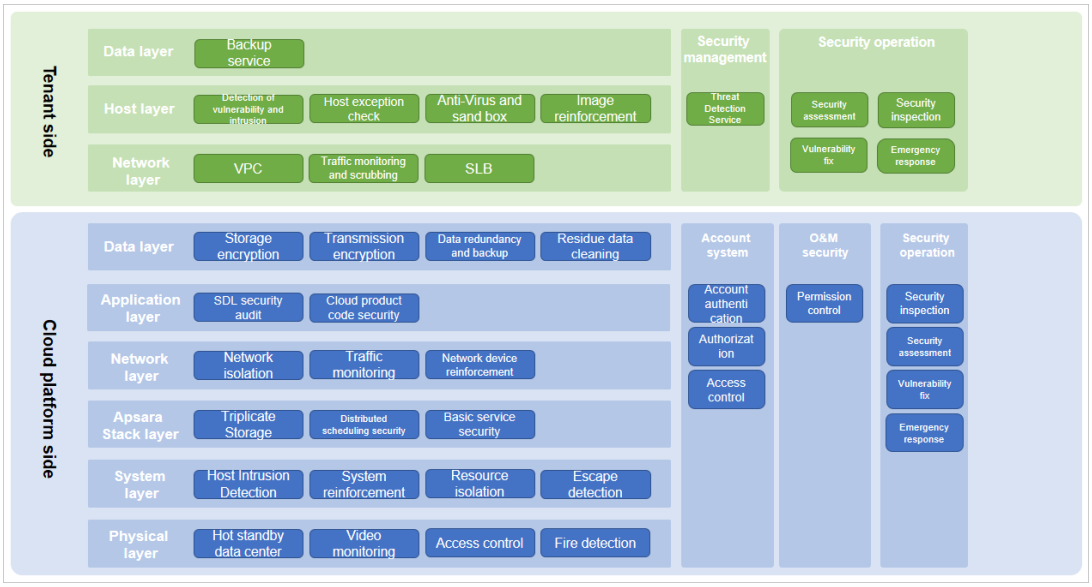


In the Apsara Stack environment, the security deployment condition of each security product is shown as follows.

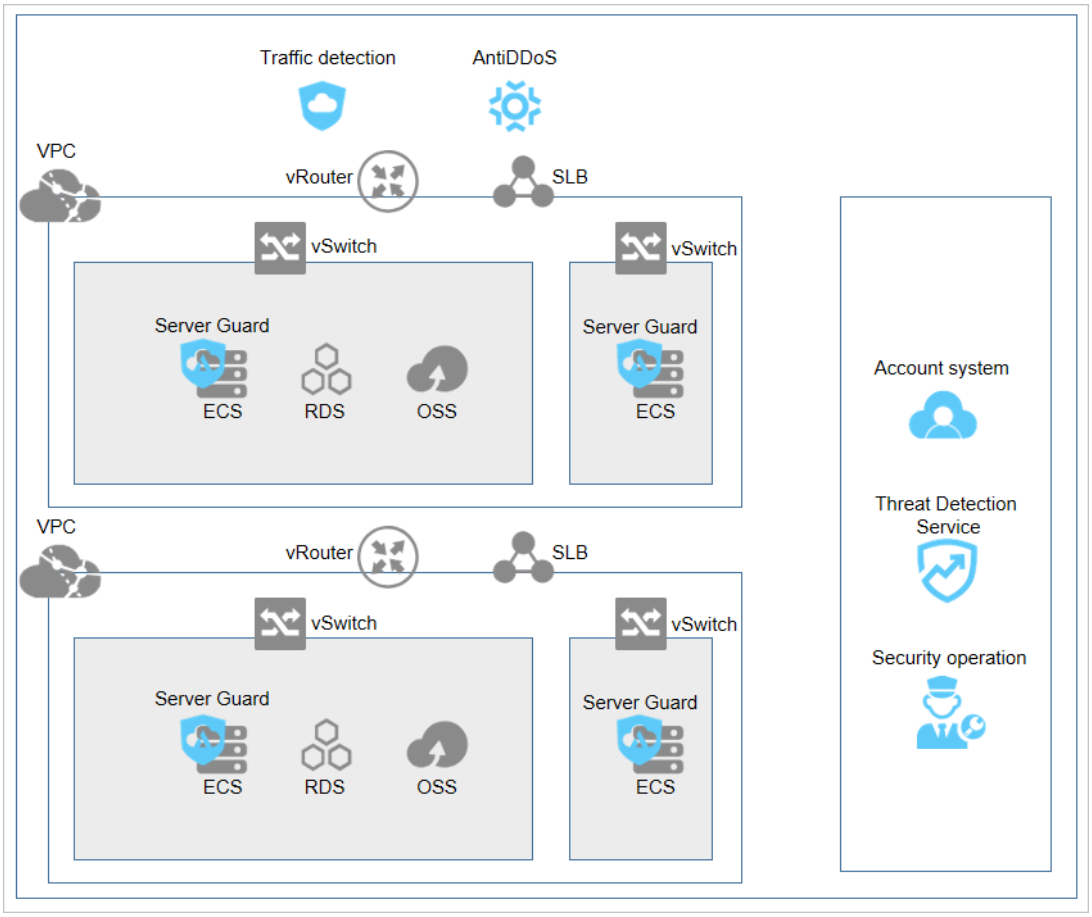


# 4.2. Cloud platform security

## 4.2.1. Apsara Stack infrastructure security

The requirements for the physical security of Apsara Stack data centers include the following security measures: dual-circuit power supply, access control, video monitoring, fire detection, and hot standby data centers.

### Dual-circuit power supply

To guarantee 24/7 non-stop services, each load in the Apsara Stack data center must be connected to two power supplies that support mutual switchover. If one power supply fails, the load is connected to the other power supply.

### Access control

Access control must be set for the Apsara Stack data center and the physical devices in the data center. For example, the access control policies must be set for the entry/exit of personnel and devices in the data center and also for the configuration, start-up, shutdown, and fault recovery of physical devices.

### Video monitoring

A video monitoring system or dedicated persons must monitor the channels or other important locations in the Apsara Stack data center around the clock. For example, the video monitoring system must monitor the entry and exit, and the alert device must collaborate with the video monitoring system or access control device to effectively monitor the monitoring sites.

### Fire detection

The Apsara Stack data center must be equipped with an automatic fire alert system, including the automatic fire detector, regional alert, and centralized alert and controller. The automatic fire alert system sends alert signals by sound, light, or point on the fire location, starts the automatic fire extinguishing device, cuts off the power, and turns off the air conditioners.

### Hot standby data centers

When a fault occurs, the faulty unit is automatically replaced by a hot standby unit based on the preset fault recovery plan to achieve automatic fault recovery.

# 4.2.2. Apsara Stack system security

## 4.2.2.1. System security of physical host

Alibaba Cloud comprehensively enhances the security of Apsara Stack physical server systems, including but not limited to account security, file permissions, system services, and host intrusion detection systems.

### Account security

Set the password length, complexity, password length, and password lifecycle for physical server accounts, delete accounts with empty passwords, and set the logon TIMEOUT parameter.

### File permission

Monitoe integrity of important directories to immediately detect intrusions when hackers tamper with or write files.

### System services

Disable unnecessary system services on the physical server to reduce attacks on the server.

## 4.2.2.2. Virtualization system security

Virtualization lays the technological foundation for the cloud computing platform and guarantees isolation between multiple tenants in a cloud computing environment by means of virtualized computing, storage, and network. Virtualization security technology of Alibaba Cloud involves tenant isolation, hotfix patches, and escape detection to guarantee the security of the virtualization layer of the Apsara Stack platform.

### Tenant isolation

The virtualization management layer plays a vital role in tenant isolation. Based on the hardware virtualization technology, virtual machine management isolates virtual machines that have multiple computing nodes at the system layer. Tenants cannot access unauthorized resources to guarantee the basic computing isolation between computing nodes. The virtualization management layer also provides storage isolation and network isolation.

- Computing isolation

  The Apsara Stack platform provides various cloud-based computing services including computing instances and services, and allows automatic scaling to meet the requirements of applications and users. These computing instances and services provide computing isolation at multiple levels to protect data and guarantee flexible configuration to meet users' needs. The computing isolation is directly provided by Hypervisor, and the key computing isolation boundaries are between the management system and users' virtual machines, and also between users' virtual machines. In the virtualized environment of Apsara Stack platform, user instances run as standalone virtual machines. The isolation is enforced with physical processor-level permissions to avoid unauthorized users' virtual machines to access physical hosts and the system resources on other users' virtual machines.

- Storage isolation

  In the basic design of cloud computing virtualization, Alibaba Cloud separates computing based on virtual machine from storage. This separation allows computing and storage to be extended independently and makes it easier to provide multi-tenant services. At the virtualization layer, Hypervisor uses the separation device driver model to implement I/O virtualization. Hypervisor intercepts and processes all I/O operations of a virtual machine to make sure that the virtual machine can only access the physical disk space allocated to it. This realizes security isolation of hard disk space between virtual machines. After the releasing of a user instance server, the original disk space is reliably cleared to guarantee the user data security.

- Network isolation

  To guarantee the network connections of virtual machine instances, Alibaba Cloud connects virtual machines to the Apsara Stack virtual network. A virtual network is a logical structure built on the physical network structure. Each logical virtual network is isolated from other virtual networks. This isolation prevents the network traffic data being accessed by other instances during deployment.

### Escape detection

A virtual machine takes two steps to perform escape attack: first it places the virtual machine controlled by the attacker on the same physical host as one of the target virtual machines. Then, it destroys the isolation boundary to steal sensitive information of the target or perform operations that compromise the functions of the target.

The virtualization management of Apsara Stack platform uses the advanced virtual machine layout algorithm to prevent virtual machines of malicious users from running on specific physical machines. At the software level of virtualization management, Alibaba Cloud also provides reinforcement, attack detection, and hotfix of virtualization management programs to prevent attacks from malicious virtual machines.

### Hotfix patches

The Apsara Stack virtualization platform supports the hotfix patch technology, which can fix system defects or vulnerabilities without restarting the system and then avoid affecting users' business.

# 4.2.3. Distributed Apsara system security

## 4.2.3.1. Security of distributed file system

The distributed file system adopts triplicate technology to store data in the system. If one of the three copies is lost, system automatically performs copy operation to maintain the three copies in the system all the time. The three copies are stored in the same physical storage medium according to security policy. They are kept separatly for operation.

All the access operation of the distributed file system must be certified by the Capability. Only the access with approved Capability is allowed to communicate with the system, which avoids unauthorized access operation.

Data stored in the distributed file system adopts binary format to avoid information leakage caused by the direct access to the plain information.

## 4.2.3.2. Security of remote process call module

The remote process call module adopts binary format for remote communication in Apsara Stack operation system. This guarantees an efficient and secure transmission and also guarantees that even if data is hijacked by any intermediary, data cannot be restored.

## 4.2.3.3. Security of job scheduling module

The job scheduling module isolates programs by using the method of sand box.

## 4.2.3.4. Security of basic service module

Basic service module deploys specific security measures for NTP and DNS servers, such as DDos attack protection, DNS zone forward, DNS amplified attack defense, and NTP amplified attack defense.

# 4.2.4. Network security

## 4.2.4.1. Basic network security

### Logical isolation

The Apsara Stack platform adopts security isolation for the management network (OPS), business network, and physical network in the Apsara Stack network environment. The OPS, business, and physical networks are logically isolated from each other by using network access control policies to prevent mutual access. Apsara Stack platform also takes network control measures to prevent unauthorized devices from connecting to the internal network of the cloud platform and prevents the physical servers of the cloud platform from connecting to external devices.

### Anti-IP/MAC/ARP spoofing

IP/MAC/ARP spoofing always challenge traditional networks. Hackers use IP/MAC/ARP spoofing to disturb the network environment and intercept network secrets. The Apsara Stack platform solves the address spoofing problem by using the underlying network technology on the physical server.

The Apsara Stack platform isolates the abnormal protocol access initiated by a server to external targets on the data link layer of the physical server, blocks the MAC/ARP spoofing of the server, and avoids IP spoofing of the server on the network layer of the host.

### Apsara Stack Security - Traffic Security Monitoring

Traffic Security Monitoring module is a millisecond(ms)-level attack monitoring product that is developed independently by Alibaba Cloud security team. With an in-depth analysis of incoming image traffic packages in the Apsara Stack environment, this module can detect various attacks and abnormal behaviors in real time.

For more information about the Traffic Security Monitoring module, see **Features > Apsara Stack Security Standard Edition > Traffic Security Monitoring** in *Apsara Stack Security Technical Whitepaper*.

## 4.2.4.2. Network device security

### Account security

Reinforces the storage encryption of the account password policies and password configuration files for network devices.

- Provides network devices with read-only accounts that can only view configurations to separate the reading configuration accounts from changing configuration accounts.
- Uses the centralized control policy to manage accounts in a unified manner.
- Uses multi-factor authentication to guarantee the account security for network devices.

### Services

Disables services on network devices to reduce attack surface of the network devices and disables features uncorrelated to the network devices.

### Log centralization

Collects and manages logs generated by network devices in a centralized manner.

# 4.2.5. Application security

## 4.2.5.1. Secure Product Lifecycle (SPLC)

Secure Product Lifecycle (SPLC) is tailored by Alibaba Cloud for cloud products, which aims to integrate security into the entire product development lifecycle. With SPLC, a complete security review module is implemented at each node from product architecture review, development, test review, to emergency response. This makes sure that the product security performance can meet the strict cloud security requirements, effectively improves security capabilities of cloud products, and reduces security risks.

The entire SPLC of a cloud product can be divided into the following six phases: product initiation, security architecture review, security development, security testing and review, application release, and emergency response.

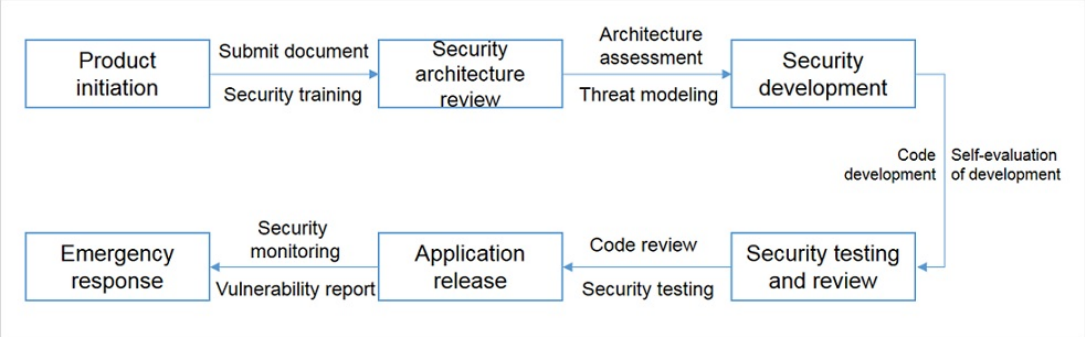


- In the product initiation phase, the security architect works together with the product team to establish a functional requirements document (FRD) and a detailed architecture diagram based on the business contents, business process, and technical frameworks. This also extracts the *Security Baseline Requirements* that is applicable to the product scope from all the security baseline requirements for the this Apsara Stack product. In this phase, specific security training courses and exams are also arranged for the product team memebers to avoid obvious security risks in subsequent product development.
- In the security architecture review phase, the security architect evaluates the security architecture of products and creates threat models of the products based on the FRD and architecture diagram established in the preceding phase. In the process of threat modeling, the security architect creates detailed models for every asset that requires protection, security requirements of assets, and scenarios where attacks may occur, and then proposes corresponding security solutions. The security architect then works with the product team to determine all the *Security Requirements* for the products, based on the preceding *Security Baseline Requirements* and the security solutions proposed during threat modeling.
- In the security development phase, the product team must abide by the secure coding standards in product development in accordance with the *Security Requirements* and achieve relevant security features and requirements of the products. To guarantee a rapid and continuous development, release, and deployment of cloud products, the product team carries out self-evaluation in this phase to confirm that the *Security Requirements* is implemented. Then, the team provides the security engineer who is responsible for testing with corresponding test information, such as the code implementation address and self-testing result report, to prepare for the security testing and review in the next phase.
- In the security testing and review phase, the security engineer implements comprehensive security reviews on the architecture design and server environment of the products according to their *Security Requirements* . The engineer also performs code review and penetration testing on the products. The product team must fix and reinforce products with security problems found in this phase.
- In the application release phase, only products that pass the security review and get the security approval can be deployed in the production environment by using a standard release system. This prevents products with security vulnerabilities from running in the production environment.
- In the emergency response phase, the security emergency team constantly monitors possible security problems in the cloud platform. They also identify security vulnerabilities by using external channels such as ASRC or internal channels, such as internal scanners and self-testing on security. If a security vulnerability is detected, the emergency team quickly rates it, determines its priority, and schedules it for fixing. The team allocates resources appropriatly to quickly fix vulnerabilities. This guarantees the security of Alibaba Cloud and its users.

#### 4.2.5.2. Web Application Firewall (on the platform side)

The Apsara Stack platform uses Web Application Firewall (WAF) to protect the security of platform applications. WAF blocks and intercepts OWASP Top 10 attacks, including Structured Query Language (SQL) injection, cross-site scripting (XSS) attack, and other attacks on Web applications to guarantee the security of platform applications.

For more information about the Web Application Firewall module, see **Features > Apsara Stack Security Standard Edition > Web Application Firewall** in *Apsara Stack Security Technical Whitepaper*.

### 4.2.6. Data security

#### 4.2.6.1. Data security system

Alibaba Cloud develops its data security system comprehensively and systematically by taking management and technical measures based on the data security lifecycle. Data security is managed and controlled during the data lifecycle, from data production, data storage, data usage, data transmission, data distribution, to data destruction.

The Apsara Stack platform has corresponding security management systems and security technologies at each stage of data security lifecycle.

#### 4.2.6.2. Data ownership

In July 2015, Alibaba Cloud initiated the first Data Protection Proposal among cloud computing service providers in China. This public proposal appeals that the ownership of data of developers, companies, governments, and social institutions on the cloud computing platforms all belongs to the users. The cloud computing platforms cannot use the data for other purposes. Platform providers have responsibility and obligation to help users protect the privacy, integrity, and availability of their data.

#### 4.2.6.3. Multi-copy redundancy storage

Apsara Stack uses the distributed storage technology to divide a file into many data fragments, stores them on different devices, and creates multiple copies for each data fragment. Distributed storage improves data reliability and security.

#### 4.2.6.4. Full-stack encryption

Apsara Stack provides full-stack encryption to guarantee the data security, namely sensitive data encryption in applications, transparent data encryption in ApsaraDB for RDS, block storage data encryption, object storage system encryption, hardware encryption modules, and network data transmission encryption. To encrypt sensitive data in applications, Apsara Stack uses encryption solutions in a hardware-trusted execution environment provided by the processor.

#### 4.2.6.5. Residual data cleanup

After memories and disks that once stored user data are released and recycled, all the residual data on them are automatically cleared.

#### 4.2.6.6. Operations data security

Without the permission of users, operations personnel cannot access unpublished data of users in any way.

Complying with the principle that production data stays within the production clusters, the Apsara Stack platform technically controls the channels where the production data flows out of the production clusters. This prevents the operations personnel from copying data from the production system.

# 4.2.7. Account system security

## 4.2.7.1. Overview

The Apsara Stack platform provides various security measures to help users protect their accounts and avoid operations of unauthorized users. These security measures include logon as a cloud account, RAM user creation, centralized management of RAM user permissions, data transmission encryption, and audit operation of RAM users. Users can use these measures to protect their cloud accounts.

## 4.2.7.2. Super administrator

The Apsara Stack platform has a default super administrator who can create system administrators and notify them of the default password by SMS or email. You must modify the password of your username as instructed when you log on to the Apsara Stack console for the first time. To improve security, the password must meet minimum complexity requirements, that is, 8 to 20 characters in length and containing at least two types of the following characters: English uppercase or lowercase letters (A to Z or a to z), numbers (0 to 9), or special characters (such as exclamation marks (!), at signs (@), number signs (#), dollar signs ($), and percent signs (%)).

## 4.2.7.3. Apsara Stack account

An Apsara Stack account is used to manage operations on the Apsara Stack platform and resources of cloud tenants.

An Apsara Stack account is the basic unit for the Apsara Stack resource ownership and the resource usage measurement. A user must register an Apsara Stack account before using Apsara Stack services. An Apsara Stack account has full permissions to all the resources it owns. By default, a resource can only be accessed by the Resource Owner. Other users must be explicitly authorized by the owner to access the resource, that is, the owner must grant the object to other users. Therefore, from the perspective of permission management, the Apsara Stack account is similar to the root or administrator account of an operating system. Sometimes the Apsara Stack account is called the root or primary account.

An authorized Apsara Stack account can have management permissions of the cloud resources or operations permissions of the cloud platform. The operations permissions of the cloud platform are managed by using OAM and the resource management permissions of the cloud tenants are managed by using RAM. RAM also supports the system of primary account and RAM users.

## 4.2.7.4. Identity credential

An identity credential is used to verify the real identity of a user. It usually refers to a user's logon password or AccessKey. Identity credentials are confidential, so users must keep their credentials secret.

- **Logon username/password**

  Users can use the logon username and password to log on to the Apsara Stack console to apply for resources and perform operations on resources.

- **AccessKey**

  Users can use the AccessKey to construct an API request (or use cloud service SDKs) to perform operations on resources.

# 4.2.7.5. RAM

Cloud tenants can use Resource Access Management (RAM) to build a system of primary account and RAM users.

RAM is an Apsara Stack service designed for user identity management and access control. You can use RAM to create and manage user accounts (such as employees, systems, and applications), and grant the accounts operation permissions to their resources. If multiple users collaboratively work with resources, RAM allows you to avoid sharing the password or AccessKey of your Apsara Stack account with other users. You can grant users the minimum permissions as required to reduce information security risks.

## RAM user identity types

RAM supports two different user identity types: RAM-User and RAM-Role.

- RAM-User

  A RAM-User is a physical identity with a fixed ID and authentication key. Generally, it corresponds to a specific person or application.

- RAM-Role

  A RAM-Role is a virtual identity with a fixed ID, but no authentication key. A RAM-Role must be associated with one or more physical identities before it becomes available. For example, it can be associated with RAM-Users under the current or another Alibaba Cloud account, with Apsara Stack services such as Elastic Compute Service (ECS), and with external physical identities such as a local enterprise account.

## Permissions

A permission is used to allow or deny a user's operation on a certain kind of resources.

Operations can be divided into two categories: resource control operations and resource use operations.

- Resource control operations are operations for lifecycle management and Operation and Maintenance (O&M) management of cloud resources, such as creating, pausing, and restarting Elastic Compute Service (ECS) instances, and creating, changing, and deleting Object Storage Service (OSS) buckets. Resource control is generally oriented to resource owners or O&M personnel in an enterprise organization.
- Resource use operations are the use of the core functions of the resources, such as user operations in an ECS instance operating system, and uploads/downloads of OSS bucket data. Resource use is oriented to applications or R&D personnel in an enterprise organization.

For elastic computing and database products, resource control operations are managed by using RAM and resource use operations are managed in each product instance, such as the permission control of ECS instance operating system or MySQL database. For storage products, such as OSS and Table Store, resource control operations and resource use operations can be both managed by using RAM.

### Authorization policies

An authorization policy is a type of simple language specification that describes a permission set.

RAM supports two types of authorization policies: system access policies managed by the Apsara Stack platform and custom access policies managed by users. For system access policies managed by the Apsara Stack platform, users can only use and cannot change the policies, and the platform updates the policy versions automatically. For custom access policies managed by users, users can create and delete policies and maintain the policy versions by themselves.

RAM allows users to create and manage multiple authorization policies under an Apsara Stack account. Each authorization policy is essentially a set of permissions. The administrator can allocate one or more authorization policies to RAM users (namely RAM-User and RAM-Role). The RAM authorization policy language can convey the authorization meaning in details, which can grant permissions to a specified API-Action and Resource-ID and can also support multiple restrictions such as the source IP address and access time.

# 4.2.8. Operations security

## 4.2.8.1. Overview

Apsara Stack provides a set of centralized operations management system, the Apsara Stack Operations system, briefly called ASO. It enables various kinds of operations roles for Apsara Stack, including field operations engineer, user operations engineer, cloud platform operations and management engineer, and operations security personnel or audit management personnel. ASO enables operations engineers to control the system operation status in time and perform corresponding operations actions.

## 4.2.8.2. OAM permission and authorization

Operation Administrator Manager (OAM) is a permission management platform for Apsara Stack Operations. OAM uses a simplified Role-Based Access Control (RBAC) model. Administrators can assign roles to operations personnel by using OAM. The operations personnel have different operation permissions to different operations systems based on their roles.

### OAM permission model

In RBAC, the administrator does not directly grant system operation permissions to specific users, but creates a role set between the sets of users and permissions. Each role corresponds to a group of permissions. After being assigned a role, a user can have all permissions of that role. Therefore, when creating a user, you are only required to assign a role to the user, without granting specific permissions to the user. The change of role permission is less frequent than that of the user permission, which simplifies permission management and reduces system overhead.

### OAM authorization system

The administrator grants permissons to operations personnel of different roles by configuring the following parameters:

- Subject: operators to the access control system. OAM subjects include users and groups.
- User: administrators and operators of the operations system.
- Group: a set of multiple users.
- Role: core of the RBAC system. Generally, a role can be considered as a set of permissions. A role can contain multiple RoleCells and/or roles.
- RoleHierarchy: In the OAM system, a role can contain other roles to form a RoleHierarchy.
- RoleCell: specific description about a permission. A RoleCell consists of resources, operation sets, and authorization options.
- Resource: description about authorization objects. For more information about resources on each operations platform, see the permission list of each operations platform.
- ActionSet: description about authorized actions. An ActionSet can contain multiple actions. For more information about actions on each operations platform, see the permission list of each operations platform.
- WithGrantOption: maximum number of authorizations in cascaded authorization, which is an integer greater than or equal to zero. If the value is not zero, the permission can be granted. If the value is zero, the permission cannot be granted.

  For example, when administrator A grants a permission to administrator B, the WithGrantOption value is 5, indicating that the permission can be granted for five times at most. When administrator B grants the permission to administrator C, the WithGrantOption value can be up to 4. If WithGrantOption is set to 0 when administrator B grants the permission to operator D, operator D can only use the permission but cannot grant the permission to others.

# 4.2.8.3. Apsara Infrastructure Management Framework permission management (data center management)

Apsara Infrastructure Management Framework is an automatic data center management system that manages the hardware lifecycles and various static resources in the Apsara Stack data center, including programs, configurations, operating system images, and data.

Apsara Infrastructure Management Framework provides a set of universal version management, deployment, and hot upgrade solutions for the Apsara system and applications and services of various Apsara Stack products. Services based on Apsara Infrastructure Management Framework can realize automatic operations in a large-scale distributed environment, which makes the operations more efficient and the system more available.

## Permission management

The permission management of Apsara Infrastructure Management Framework is based on the OAM system. The user permissions of Apsara Infrastructure Management Framework include the Admin permissions, Project permissions, and Service permissions:

- Admin permissions: Administrators can manage all pages on the Apsara Infrastructure Management Framework platform.
- Project permissions:

- The administrator must grant users the Project permissions to view the project information in **Operations > Project Operations** on the Apsara Infrastructure Management Framework platform.
- The administrator must grant users the Project permissions to view the cluster information and perform operations on the cluster in **Operations > Cluster Operations** on the Apsara Infrastructure Management Framework platform.

- Service permissions: The administrator must grant users the Service permissions to view the service information and perform operations on the service in **Operations > Service Operations** on the Apsara Infrastructure Management Framework platform.

# 4.3. Security operation service (on the platform side)

Apsara Stack Security provides multiple platform security operation services that are specified on the platform side of Apsara Stack.

### Security inspection

Apsara Stack Security investigates and sorts lists of cloud platform services, including the number of physical machines and the version of each product. Apsara Stack Security also analyzes event logs of basic security products that are provided by the cloud platform and defends against the security risks of products.

### Security evaluation and reinforcement

Apsara Stack Security evaluates the security of the cloud platform system, detects security risks of network, host, and application on the cloud platform, and then reinforces security against the detected risks.

### Vulnerability fixing

Apsara Stack Security fixes security vulnerabilities, such as password and configuration problems detected during the cloud platform running process.

### Security emergency response

If a security emergency such as an intrusion occurs, Apsara Stack Security responds to the emergency in time and analyzes the event cause.

# 4.4. Cloud user (tenant) security

## 4.4.1. Network security

## 4.4.1.1. Virtual Private Cloud

Virtual Private Cloud (VPC) helps users establish an isolated network environment based on Alibaba Cloud. It supports the customization of IP address range, Classless Inter-Domain Routing (CIDR) block, route table, and gateway. VPC can also be connected to a traditional data center by using leased lines or VPN to build a hybrid cloud.

### Security isolation

By using the tunneling technology, VPC isolates networks to the same effect as the traditional VLAN. VPC isolates broadcast domains at the NIC level, thoroughly blocks network communication by VLAN isolation, and classifies different security domains for access control.

### Access control

VPC supports flexible access control rules based on the security group firewall.

## 4.4.1.2. Distributed firewall

A security group is the distributed virtual firewall provided by Alibaba Cloud, which provides the status detection and packet filtering features.

A security group is a logical group that consists of instances in the same region with the same security requirements and mutual trust. Security groups are used to set network access control rules for one or more Elastic Compute Service (ECS) instances. As an important network security isolation measure, security groups are used to divide network security domains on the cloud.

Each instance must belong to at least one security group. Instances in the same security group can interwork by using networks. By default, instances in different security groups cannot interwork by using the intranet, but a certain source security group or source Classless Inter-Domain Routing (CIDR) block can be authorized to access a target security group to implement the interworking.

## 4.4.1.3. Server Load Balancer

Server Load Balancer (SLB) is a load balancing service that distributes traffic among multiple Elastic Compute Service (ECS) instances. It can scale up the service capability of an application by distributing traffic and make the system more available by eliminating single points of failure.

## 4.4.1.4. Apsara Stack Security - Traffic Security Monitoring

Traffic Security Monitoring module is a millisecond(ms)-level attack monitoring product that is developed independently by Alibaba Cloud security team. With an in-depth analysis of incoming image traffic packages in the Apsara Stack environment, this module can detect various attacks and abnormal behaviors in real time. Then it reports the security events to the Apsara Stack Security center and combines with other protection modules to fully protect users' system. Traffic Security Monitoring module provides robust information and basic data support.

For more information about the Traffic Security Monitoring module, see **Features > Apsara Stack Security Standard Edition > Traffic Security Monitoring** in *Apsara Stack Security Technical Whitepaper*.

## 4.4.1.5. Apsara Stack Security - DDoS Traffic Scrubbing

DDoS Traffic Scrubbing module provides automatic detecting, scheduling, and scrubbing for DDoS attacks to implement attack detection, traffic traction, and traffic scrubbing within five seconds by combining with the Traffic Security Monitoring module of Apsara Stack Security. It guarantees the business network stability of cloud tenants. For the defense trigger condition, DDoS defense system not only relies on the thresholds for traffic, but also bases on the statistical judgement of network behaviors, which guarantees accurate recognition of DDoS attack and maintains users' business availability when being attacked by DDoS.

For more information about the DDoS Traffic Scrubbing module, see **Features** > **Optional security services** > **DDoS Traffic Scrubbing** in *Apsara Stack Security Technical Whitepaper*.

#### 4.4.1.6. Cloud Firewall

Cloud Firewall is a firewall service used in cloud environments to resolve vague or undefined security boundaries caused by rapid changes in cloud businesses. Cloud Firewall sorts and isolates businesses based on visualized business data to enable access control on the east-west traffic in Apsara Stack.

For more information, see *Security Whitepaper*. Use the following path to navigate the guide: Apsara Stack Security > **Features** > **Optional security services** > **Cloud Firewall**.

### 4.4.2. Host security

#### 4.4.2.1. Operating system of ECS instances

Users have full control over the operating systems of their Elastic Compute Service (ECS) instances. Alibaba Cloud does not have any permission to access users' instances and operating systems on them. We recommend that users must access and operate ECS instances by using secure methods, such as using the SSH public/private key pair and protecting the private key well (a complex password must be used at least and it can be set upon creating instances), telnetting by using a safer SSHv2, and escalating the privilage temporarily by using sudo command.

#### 4.4.2.2. Image reinforcement

An image is an environment template for running Elastic Compute Service (ECS) instances. It generally includes an operating system and preinstalled softwares. ECS tenants can use images to create ECS instances or change the system disks of ECS instances.

Security measures for Alibaba Cloud basic images (supporting various Linux/Windows release versions) include basic security configuration, vulnerability fixing, and adding host security software by default. Basic images are configured with best security practices for the hosts by default. Alibaba Cloud host security software is added to all Alibaba Cloud basic images by default to guarantee the security of instances upon start-up.

Alibaba Cloud uses the data check algorithm and one-way hash algorithm to guarantee image integrity and prevent the images from malicious tampering. After detecting a new high-risk vulnerability, users must promptly update their basic images. Users can also upgrade the operating system or fix vulnerabilities of their ECS instances autonomously.

#### 4.4.2.3. Image snapshot

Elastic Compute Service (ECS) in the Apsara Stack platform provides snapshots and custom images. Snapshots can save the status of system data at a certain time point for data backup, which allows users to achieve disaster recovery quickly. Users can create custom images by using snapshots, which includes the whole operating system and data environment information of the snapshots in the images. Snapshots are incremental and only the changed data is copied between two snapshots.

#### 4.4.2.4. Apsara Stack Security - Server Guard

Apsara Stack Security Server Guard module provides security protection measures such as vulnerability management, baseline check, intrusion detection, and asset management for Elastic Compute Service (ECS) by means of log monitoring, file analysis, and feature scanning. The Server Guard module is divided into client side and server side. Server Guard client side works with server side to monitor attack behavior and vulnerability at the host system layer and application layer, which protects the host security in real time.

### Vulnerability management

The vulnerability management provided by Server Guard for ECS incorporates multiple scanning engines (namely network side, local side, and PoC verification) to detect all vulnerabilities in the system at a time. Features such as one-click fixing, fixing command generation, and one-click batch verification are provided to implement closed-loop vulnerability management.

### Baseline check

The baseline check provided by Server Guard can automatically detect the risk points of system, database, and account configuration for ECS and provide suggestions for fixing the risks correspondingly.

### Intrusion detection

The intrusion detection provided by Server Guard includes remote logon reminder, identification of brute force attack behaviors, Webshell detection and removal, and host exception detection.

For more information about the Server Guard module, see **Features > Apsara Stack Security Standard Edition > Server Guard** in *Apsara Stack Security Technical Whitepaper*.

# 4.4.3. Application security

## 4.4.3.1. Code security

In the Secure Product Lifecycle (SPLC) of cloud products, Alibaba Cloud security experts strictly review and evaluate the code security on each development node to guarantee the code security for Alibaba Cloud products. We recommend that enterprise users must perform black-box and white-box code security test for their online applications to prevent security vulnerabilities and improve the security robustness of their businesses.

## 4.4.3.2. Apsara Stack Security - Web Application Firewall

Web Application Firewall module is a security defense system for Web applications of cloud tenant. Based on the intelligent semantics analysis engine, WAF defends against common Web attacks such as SQL injection, XSS attack, common web server plugin vulnerabilities, trojan uploads, unauthorized access to core resources, and other common OWASP attacks. It filters out numbers of malicious access attempts to prevent the leakage of users' website assets and data, and safeguard website security and availability.

For more information about the Web Application Firewall module, see **Features > Apsara Stack Security Standard Edition > Web Application Firewall** in *Apsara Stack Security Technical Whitepaper*.

# 4.4.4. Data security

# 4.4.4.1. ApsaraDB

## Tenant layer isolation

ApsaraDB in the Apsara Stack environment isolates tenants by using the virtualization technology, which allows each tenant to have independent database permissions. Alibaba Cloud also reinforces the security of the server on which databases run. For example, users cannot access system files by reading from or writing to databases, which makes sure that users cannot access data of other users.

## Database accounts

After a user creates an ApsaraDB instance, the system does not create any initial database account for the user. The user must create a common database account in the console or by using APIs and configure database-level read/write permissions. If the user requires more fine-grained permission control, such as table/view/field-level permission control, the user can also create a super database account in the console or by using APIs, and use the database client and super database account to create a common database account. Then the user can use the super database account to configure table-level read/write permissions for the common database account.

## IP address whitelist

By default, ApsaraDB instances are set to be inaccessible from any IP addresses, that is, the IP address whitelist contains only 127.0.0.1. Users can add IP address whitelist rules by using the data security module in the console or APIs. An IP address whitelist rule can take effect without restarting ApsaraDB instances and does not affect the usage. Multiple groups can be configured in the IP address whitelist, and each group can contain up to 1,000 IP addresses or IP address segments.

## VPC isolation

Users can perform advanced network access control by using Virtual Private Cloud (VPC) in ApsaraDB. VPC is a private network environment that the user sets in the cloud platform. It strictly isolates network packets by using underlying network protocols and controls access at layer 2 of the network. Users also can connect server resources of self-built data centers to the Alibaba Cloud platform by using VPN or leased lines, and solve possible IP resource conflicts by using the IP address segments of ApsaraDB instances defined by VPC. This allows self-owned servers and ECS instances to access ApsaraDB instances simultaneously.

VPC and IP address whitelist guarantees a securer ApsaraDB instance.

## Data transfer encryption

ApsaraDB supports Secure Sockets Layer (SSL) protocol. Users can use root certificate on the server side to verify whether the target address and port database service are provided by ApsaraDB to avoid Man-in-the-Middle (MITM) attack. ApsaraDB also provides the implementation and renovation capability of SSL certificate on the server side to allow users to change SSL cerficate as required, which guarantees the security and availability of the certificate.

## Primary node and standby node

ApsaraDB adopts a high availability architecture with three nodes replica sets. Three data nodes locate in different physical servers and synchronize data automatically. Primary node and secondary node both provide service. When primary node encounters fault, the system selects new primary node automatically. When the secondary node is unavailable, the standby node takes charge.

ApsaraDB also provides automatic backup feature that supports one-click data recovery to make sure that the data is integral and reliable.

# 4.4.4.2. MaxCompute

## Authorization management

Project is the basis for implementing the multi-tenant architecture of MaxCompute in the Apsara Stack platform, and the basic unit for user data management and computing. After a project is created for a user, the user is the owner of the project. All objects such as tables, instances, resources, and UDFs in the project belong to the user. Unless authorized by the owner, no one can access objects in the project.

When the owner of a project decides to authorize another user, the owner must add the user to the project. Only users added to the project can be authorized.

A role is a collection of access permissions. A role can be used to assign the same permissions to a group of users. Role-based authorization greatly simplifies the authorization process and reduces the authorization management cost. Role-based authorization should be used preferentially when a user is authorized.

MaxCompute can assign different permissions to users or roles in a project based on four different objects, namely the project, table, function, and resource instance.

## Cross-project resource sharing

If a user is the owner or administrator of a project, other users must apply for the permission to access resources of the project. If the applicant belongs to the project team of the owner, the owner is advised to use the user and authorization management function of the project. If the applicant does not belong to the project team of the owner, the owner can use the Package-based cross-project resource sharing function.

Package is used for sharing data and resources across projects, which implements the cross-project user authorization. After Package is used, the administrator of project A can perform packaging authorization on the objects to be used in project B, namely, creating a Package, and then permits project B to install the Package. After the administrator of project B installs the Package, the administrator can determine whether to grant permissions of the Package to the users of project B as required.

## Data protection

If a project contains highly sensitive data that cannot be shared with other projects, project protection can be used by setting ProjectProtection. Project protection explicitly requires that only inbound data is allowed in the project.

# 4.4.4.3. Sensitive Data Discovery and Protection of Apsara Stack Security

The Sensitive Data Discovery and Protection (SDDP) system makes full use of big data analysis capabilities and AI-related technologies of Alibaba to achieve classification based on business needs by using intelligent sensitive data identification. The dynamic and static desensitization, global circulation monitoring, and exception detection are implemented based on precise identification to achieve precise identification, precise detection, precise analysis, and effective protection, which meets visible, controllable, and compliant security protection requirements. This product supports MaxCompute, OSS, Table Store and other Alibaba Cloud big data products.

For more information about the SDDP module, see *Apsara Stack Security Technical Whitepaper* **Features > Optional security services > Sensitive Data Discovery and Protection.**

# 4.4.5. Security management

## 4.4.5.1. Apsara Stack Security - Threat Detection Service

The Threat Detection Service (TDS) module is the big data security analysis system developed by Alibaba Cloud security team. It performs an in-depth analysis of host traffic and network traffic in the Apsara Stack environment by using machine learning and data modeling, and detect abnormal behaviors such as threat, attack, and access. From the attacker's perspective, it effectively captures vulnerability attacks and new virus attacks conducted by advanced attackers and displays ongoing security attacks to implement the visulization and awareness of business security.

Based on the Internet visualized technology, TDS also presents the result of big data threat analysis on a dashboard with a visual graph, which provides the entire security information and supports the security decision-making of Apsara Stack platform for users.

For more information about the Threat Detection Service module, see **Features > Apsara Stack Security Standard Edition > Threat Detection Service** in *Apsara Stack Security Technical Whitepaper* .

# 4.4.6. Security operation service (on the tenant side)

Alibaba Cloud provides cloud tenants with the security operation service to operate resources and management policies on the Apsara Stack platform, including configuration and hosting of security product, response of security event, accident tracking, security inspection, monitoring and scanning, and security process management. This service continuously guarantees the consecutive and secure operation of tenants' businesses.

# 4.4.7. Best security practices

We recommend that tenants must migrate previous security policies during the cloud migration and take the following best practices of security configuration offered by Alibaba Cloud as a reference to guarantee the security of their businesses:

- Cloud resource security: the security of cloud resources, which must be guaranteed by using Virtual Private Cloud (VPC).
- Apsara Stack Security: Apsara Stack Security is used to guarantee the security of tenants' businesses. The synchronization feature can synchronize the latest Apsara Stack Security rules in time. We recommend that users use Web Application Firewall (WAF) to protect Web applications.

- Security configurations of cloud products:
  - Complex passwords must be set for Elastic Compute Service (ECS) instances to prevent intrusions by brute force cracking.
  - Secure Shell (SSH) and Remote Desktop Protocl (RDP) management ports must be restricted by using security groups.
  - If high-risk ports are enabled on ECS instances, the IP address whitelist must be configured for access control.
  - Server Load Balancer (SLB) instances are prohibited to enable access of SSH, RDP, MySQL, Redis, and other high-risk port services to the Internet.
  - High-intensity passwords must be set for RDS instances, and the IP address whitelist must be configured for access control.
  - The access to Object Storage Service (OSS) instances must be restricted by using access control rules, and public read/write operations are disabled.
- Application deployment security: The compressed packages, .svn hidden directories, and .git hidden directories must be deleted before code deployment. Security of Linux, Windows, and other operating systems must be reinforced. We recommend that users use WAF to protect Web applications.

# 5.Security of Apsara Stack products

## 5.1. Elastic Compute Service (ECS)

### 5.1.1. Platform security

#### 5.1.1.1. Security isolation

Instance security isolation includes the following aspects:

##### CPU isolation

ECS supports the KVM hypervisor. Using VT-x virtualization, the hypervisor runs in vmx root mode while ECS instances run in vmx non-root mode. Hardware isolation prevents ECS instances from accessing the privileged resources of other instances.

##### Memory isolation

The hypervisor isolates memory on the virtualization layer. When ECS instances are running, extended page tables (EPT) ensure that ECS instances cannot access the memory resources of other instances.

After an ECS instance is released, all of its memory is cleared by the hypervisor to prevent other ECS instances from accessing the physical memory released by this ECS instance.

##### Storage isolation

On the virtualization layer, the hypervisor uses a device driver model to achieve I/O virtualization. ECS instances cannot directly access physical disks, and all I/O operations are intercepted and processed by the hypervisor. The hypervisor ensures that ECS instances can only access the allocated virtual disk space, thus realizing the security isolation of disk space between different ECS instances.

##### Network isolation

ECS uses Virtual Switches (VSwitches). Messages destined for an ECS instance are only sent to the VSwitch port corresponding to the virtual network interface of the ECS instance.

ECS instances, even those running in hybrid mode, are not able to receive or intercept messages intended for other ECS instances. Even if you set the network interface to the hybrid mode, the hypervisor does not transmit any traffic destined for an instance to any other instances.

Alibaba Cloud further uses Virtual Private Networks (VPCs) and security groups to isolate networks.

A security group is an additional security barrier for ECS instances provided by Alibaba Cloud. It implements a distributed virtual firewall with stateful packet inspection. A security group is independent of the operating system firewalls on the ECS instances within the group. The security group provides additional protection from outside of ECS instances. Security groups allow you to implement isolated security domains by configuring inbound or outbound policies for a single IP address or port.

A security group is a logical group that consists of a group of instances in the same region that share security requirements and have mutual access permissions. Security groups are used for network access control for one or more ECS instances and allow you to divide a cloud into separate security domains.

With the preceding isolation measures, even if two instances owned by the same user run on the same physical server, the two instances are unable to intercept each other's traffic.

In addition, we recommend that you encrypt data before saving it to ECS instance disks with either an encrypted file system or disk encryption. For more information, see ECS disk encryption.

# 5.1.1.2. Authentication

# 5.1.1.2.1. Identity authentication

Account authentication verifies the identity credentials of a user. An identity credential usually refers to a logon password or AccessKey (AK). You can create AccessKey pairs in Apsara Stack console. An AccessKey pair is composed of an AccessKey ID and an AccessKey Secret. The AccessKey ID is a public key that is used to identify a user. The AccessKey Secret is the key used to encrypt signature strings and verify those signature strings on the server. The AccessKey Secret is used to authenticate a user's identity and must be kept confidential.

ECS authenticates each request with its included signature information for both HTTP and HTTPS requests. ECS uses AccessKey pairs to implement symmetric encryption and authenticate the identity of a request sender.

AccessKey pairs are issued by Alibaba Cloud to users and can be applied for and managed through the official Alibaba Cloud website. The AccessKey ID acts as a unique identifier for a user, while the AccessKey Secret is used to encrypt and verify signature strings. AccessKey Secrets must be kept confidential.

# 5.1.1.2.2. Access control

Resource Access Management (RAM) is a centralized user management and resource access control service provided by Alibaba Cloud. Using RAM, you can create independent user accounts for your employees, systems, or applications and control their access to cloud resources. Each RAM user can log on to Apsara Stack console or call service APIs by using an independent logon password or AccessKey. By default, a newly created RAM user does not have any permissions on resources. Only an authorized RAM user can operate resources on behalf of the corresponding Apsara Stack tenant account.

With RAM, you can avoid sharing your AccessKey with other users and assign minimum permissions to different users, thus reducing security risks to the data of your enterprise. An Apsara Stack tenant account can have multiple RAM users. RAM provides features such as multi-factor authentication (MFA), strong password policies, isolation of console users from API users, custom fine-grained authorization policies, group-based authorization, temporary authorization credentials, and temporary account disabling. RAM authorization can address an API action or resource ID. You can specify limits, such as source IP address, secure access channel SSL/TLS, access time period, or MFA.

We strongly recommend that you ensure that the operating systems of ECS instances are only accessed in a secure manner. Some steps to ensure access security include keeping the private key confidential when using SSH key pairs, setting complex logon passwords when creating instances, using SSHv2 for remote logon, and using the sudo commands to temporarily escalate permissions.

ECS users can create RAM users and different groups to manage and control resource access permissions.

RAM helps you manage resource access permissions. For example, you can enhance network security by assigning an authorization policy to a group. If the original IP address of an access request is not from an intranet address specified in the policy, the access request is rejected.

You can assign different permissions to different groups to manage ECS resources:

- SysAdmins: This group needs permissions to create and manage ECS images, instances, snapshots, and security groups. You can assign to this group an authorization policy, which permits the group members to perform all ECS operations.
- Developers: This group only needs permissions to use ECS instances. You can assign to this group an authorization policy, which permits the group members to call DescribeInstances, StartInstance, StopInstance, CreateInstance, DeleteInstance, and other APIs.

If a developer becomes a system administrator and needs higher permissions, they can be easily moved from the Developers group into the SysAdmins group.

ECS can also provide the functions of a RAM user to an instance by using STS. The instance RAM role function enables ECS instances to play certain roles with certain access permissions.

The instance RAM role function allows you to associate a RAM role with an ECS instance, so that the ECS instance can access other cloud services by using the STS temporary credential from the instance. The temporary credential is updated periodically. In this way, AccessKey security is ensured and fine-grained permission control is implemented based on RAM.

# 5.1.1.3. Data security

## 5.1.1.3.1. Overview

Key Management Service (KMS) provides key management and encryption mechanisms for the storage of encrypted data necessary on the cloud platform. The encrypted data includes authorization credentials, passwords, and keys.

## 5.1.1.3.2. Triplicate storage technology

When ECS users read or write data from or to virtual disks, the operations are translated into reads or writes from or to the files stored in the Apsara Stack data system. Apsara Stack uses a flat design in which a linear address space is divided into slices, also called chunks. For each chunk, three replicas are created and stored on different nodes in the cluster to ensure the reliability of data.

Triplicate technology involves three key components: master, chunk server, and client. The write operation of an ECS user goes through several conversions, as shown in the following procedure:

1. The client determines the chunk corresponding to the write operation.
2. The client sends a request to the master to query the chunk servers on which to store the three chunk replicas.
3. The client sends a write request to the three corresponding chunk servers based on the results returned from the master.
4. If the write operation succeeds on all three chunk replicas, the client returns a success message to the user. Otherwise, the client returns a failure message.

The call distribution policy of the master takes into account the disk usage of all chunk servers in the cluster, the chunk server distribution on different switch racks, power supply, and machine load. The policy ensures that the three replicas of each chunk are distributed across different chunk servers on different switch racks. This effectively avoids data unavailability caused by the failure of chunk servers or racks where data is stored.

Triplicate backup

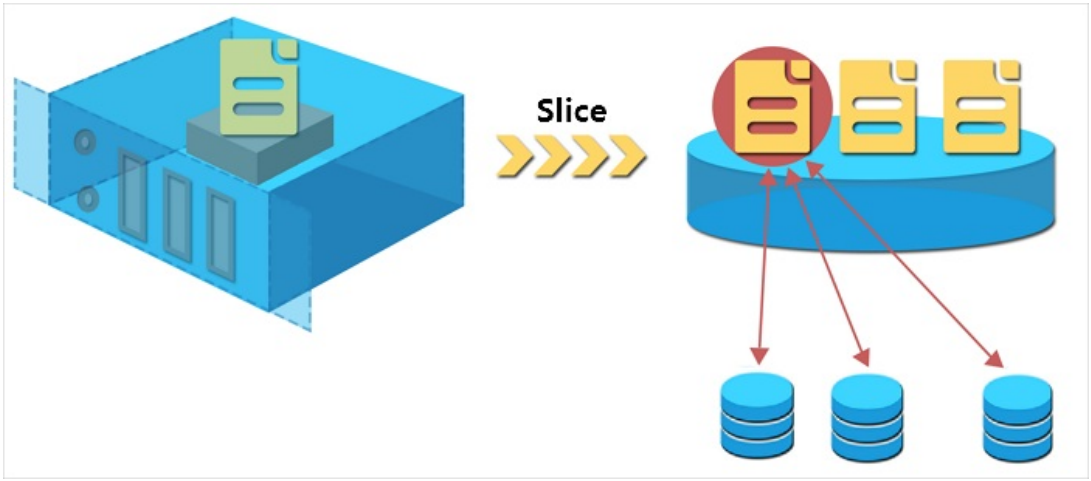


When a data node is damaged or disk faults occur on a data node, the total number of valid replicas of some chunks in a cluster becomes less than three. In such cases, the master replicates data between chunk servers to ensure that each chunk in the cluster has three valid replicas.

Automatic replication

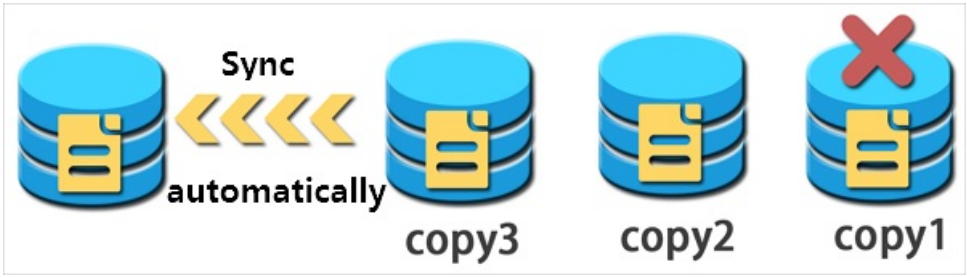


To summarize, all your (add, modify, and delete) operations on cloud disk data are synchronized to the three chunk copies at the bottom layer. This approach ensures that user data stays reliable and consistent.

When data is deleted, Apsara Distributed File System reclaims the released storage space and makes it inaccessible to any users. Data in this storage space is erased completely before it is made available for other users. This procedure ensures that user data remains secure and confidential.

## 5.1.1.3.3. ECS disk encryption

ECS disk encryption is a simple and secure encryption method that can be used to encrypt new cloud disks.

With ECS disk encryption, you do not need to create or maintain your own key management infrastructure, change existing applications and maintenance procedures, or perform additional encryption operations. Disk encryption does not have any negative impact on your business processes. The following types of data can be encrypted:

- Data on cloud disks.
- Data transmitted between cloud disks and instances. Data is not encrypted again within the instance operating system.
- All snapshots created from encrypted cloud disks. Such snapshots are encrypted snapshots.

Data transmitted from ECS instances to cloud disks is encrypted on the hosts where the ECS instances are deployed.

> Note ECS disk encryption supports Chinese cryptographic algorithms.

All available cloud disks (basic disks, ultra disks, and SSD disks) and Shared Block Storage devices (ultra Shared Block Storage devices and SSD Shared Block Storage devices) in Apsara Stack ECS can be encrypted.

#### 5.1.1.4. Encrypted transmission

Alibaba Cloud provides HTTPS encryption to ensure the security of data transmissions. If you operate data from Apsara Stack console, the console uses HTTPS for encrypted data transmission. All Alibaba Cloud services provide HTTPS-enabled APIs. Users can use AccessKey pairs to call Apsara Stack service APIs. To meet the needs of secure encrypted data transmission, Apsara Stack services use the standard SSL/TLS protocol with support for keys up to 256 bits in length.

#### 5.1.1.5. ARP spoofing prevention

ARP spoofing severely challenges traditional networks. Hackers can use ARP spoofing to imitate the routing address of another user and intercept confidential data.

To prevent ARP spoofing, Alibaba Cloud provides an ARP firewall at the Network Egress. Only MAC addresses that are allocated by the platform are authorized for communication.

### 5.1.2. Tenant security

#### 5.1.2.1. Log audit

User authentication credentials and permission control are designed to avoid security issues. Security logs can help Alibaba Cloud users better understand and diagnose a variety of security situations. Alibaba Cloud provides centralized security log management for cloud resources operations. The logon and resource access operations of each account are logged, providing information about the operator, operation time, source IP address, resource object, operation name, and operation status. With all operation records saved, users can perform security analysis, intrusion detection, resource change tracking, and compliance audit. In a compliance audit, users may need to provide detailed operation records of the Apsara Stack tenant accounts and RAM users.

#### 5.1.2.2. Secure images

Apsara Stack images integrate patches for all known high-risk vulnerabilities to prevent the host from being exposed to high risks after going online. After detecting a new high-risk vulnerability, Alibaba Cloud promptly updates images and delivers the updated images to customers. Besides, Alibaba Cloud also ensures image integrity and avoid malicious tempering by using a data verification algorithm.

Users can quickly upgrade their basic images after new high-risk vulnerabilities are detected. Moreover, users can upgrade the operating system or fix vulnerabilities of their ECS instances by themselves.

Alibaba Cloud strongly recommends that users employ Apsara Stack basic images as the first step to implement cloud migration without affecting business deployment.

#### 5.1.2.3. Block storage

Block Storage is a low-latency, persistent, and high-reliability random block-level data storage service provided by Alibaba Cloud for ECS instances. Block storage automatically replicates data within a zone to prevent unavailability caused by unexpected hardware faults and to protect your business. Just like a physical hard disk, you can format block storage attached to an ECS instance, create a file system, and persistently store data there.

Block Storage automatically encrypts block storage devices used inside of virtual machines to ensure data is secure and stored in a distributed system.

## 5.2. Container Service for Kubernetes

### 5.2.1. Platform security

#### 5.2.1.1. Security isolation

Container Service provides multiple security isolation methods to ensure cluster security.

##### Exclusive Kubernetes clusters

Kubernetes clusters you created through the Container Service console belong to you. Resources for deploying Kubernetes clusters, such as ECS and SLB instances, can be used only by the current Kubernetes clusters and are not shared with other users. Strong isolation at the physical level can avoid potential security risks arising from resource sharing.

##### ECS security group

The ECS instance used by each Kubernetes cluster belongs to the same ECS security group. Based on the least privilege principle, a security group contains only the following network access rules:

- Allow access to ECS instances over ICMP
- Allow access to ECS instances through pod CIDR blocks

##### Container network policies

In a Kubernetes cluster, pods on different nodes can communicate with each other by default. In some scenarios, the network intercommunication between different businesses is not allowed, and network policies must be introduced to reduce risks. In Kubernetes clusters, you can use the Canal network driver to implement the support for network policies.

## 5.2.1.2. Account authentication

You can use the Container Service account authentication function to secure your containerized applications.

RAM user authorization is supported for Kubernetes clusters. You can perform RAM authorization to grant permissions on Kubernetes clusters to specific RAM users. This reduces the risk of exposing Apsara Stack tenant account data.

### Role-Based Access Control (RBAC)

RBAC uses the Kubernetes built-in API group for authentication management, allowing you to manage pods corresponding to different roles and role access permissions.

## 5.2.1.3. Link security

Container Service supports TLS certificate verification for link security.

The following communication links in Container Service Kubernetes clusters are verified by TLS certificates to prevent data tampering and eavesdropping on communications:

- kubelet on worker nodes actively communicates with apiserver on master nodes.
- apiserver on master nodes actively communicates with kubelet on worker nodes.

During initialization, a master node uses SSH tunnels to connect to the SSH service of other nodes (port 22).

# 5.2.2. Tenant security

## 5.2.2.1. Application security

Container Service supports a wide range of application security policies.

### Application security policies

| Security policy | Description |
| --- | --- |
| Running containers as a non-root user | You can run applications in a container as a non-root user, so that the container cannot obtain permissions from the host by escaping a fixed state. |
| Using secure base images | You can customize base images as required and enforce the use of approved base images within your organization. You can also use secure third-party images such as Alpine-linux. Official Docker images use Alpine-linux as the base image. |

| Security policy | Description |
| --- | --- |
| Minimal image installation | Only resources necessary to run applications are installed in images. |
| Configuring TLS authentication for Docker daemons | TLS authentication is configured for Docker daemons and Docker Swarm APIs. |
| Prioritizing CPU utilization of containers | You can use the CPU sharing feature of Docker to prioritize CPU utilization of containers. The CPU sharing mechanism allows a container to take precedence over another for CPU utilization and forbids containers with a lower priority to frequently use CPU resources. This ensures high-priority containers can operate effectively and prevent CPU resource exhaustion. |
| Limiting container memory usage | A container can consume all available memory resources on the Docker host by default. You can use the memory limit mechanism to prevent a denial of service (DoS) attack arising from a single container consuming all host resources. Specifically, use -m or -memory with the docker run command to run containers. |
| Limiting container disk usage | Docker images, container rootfs, and volumes are stored in the */var/lib/docker* directory by default. They share the same file system with the host. The directory size varies depending on the content. You can mount the */var/lib/docker* directory separately to the cloud storage (such as cloud disk and OSS) without affecting the host's root file system. |
| Authentication | • Username-password pairs and certificates are used for user authentication.<br>• A logon failure processing mechanism can be used to limit the number of unauthorized logon attempts. When the number of failed logon attempts exceeds the specified threshold, the session terminates and exits.<br>• A session is locked automatically after becoming idle for a specified period of time.<br>• You must change the initial preset password when you log on to the system for the first time.<br>• The password must meet minimum complexity requirements. The new password cannot be the same as the last password. You will be prompted periodically to change your password. |

| Security policy | Description |
| --- | --- |
| Security audit | • The system provides an account security audit feature allowing you to create audit records for system account modifications. These records cannot be modified.<br>• All operations on the system platform are listed clearly, including the event date, time, initiator information, type, description, and result. Audit logs are backed up periodically and stored for at least six months. |
| Communication security and confidentiality | • Communication between system components uses TLS encryption.<br>• Confidential information in the system is saved in ciphertext and sent to specific nodes as necessary. |
| Roles and permissions | • The system has a built-in multi-tenant permission management model. The model allows you to set different access permissions based on teams and roles, and manage fine-grained permissions on clusters and applications.<br>• The system supports external permission management methods such as LDAP.<br>• An image repository allows multi-tenant, read-only permission control. |

# 5.2.2.2. Host security

## OS account requirements

An OS password must be at least 8 characters in length, and must contain at least three types of the following characters: uppercase letters, lowercase letters, numbers, and special characters. Weak passwords (such as regular or consecutive characters, employee IDs, and domain account prefixes) are not allowed. Your OS password is set to expire every 90 days by default.

## A mechanism to limit unauthorized logon attempts

A logon failure processing mechanism can be used to limit the number of unauthorized logon attempts. When the number of failed logon attempts exceeds the specified threshold, the session terminates and exits.

## Access control

Access control can be used to control user access to resources based on configured security policies.

- *passwd* file permission: 644
- *shadow* file permission: 000

- *rc3.d* file permission: 755
- *profile* file permission: 644
- *profile.d* folder permission: 755

### Disabling and deletion of default accounts

You can delete redundant and expired accounts to prevent accounts from being shared. You can disable the following default accounts: sync, shutdown, and halt.

# 5.3. Auto Scaling (ESS)

## 5.3.1. Platform security

### 5.3.1.1. Security isolation

ESS implements user account-based isolation. You can manage scaling groups, configurations, and rules in your account, such as performing create, modify, and delete operations. ESS can use ECS instance resources only in your account for automatic scaling. ESS performs symmetric encryption by using AccessKey pairs to authenticate users who manage ECS instance resources. It authenticates each access request to ensure security isolation.

### 5.3.1.2. Authentication

#### 5.3.1.2.1. Authentication

You can create an AccessKey pair in the Apsara Stack console. An AccessKey pair is composed of an AccessKey ID and an AccessKey secret. The AccessKey ID is a public key that is used to identify a user. The AccessKey secret is the key used to encrypt signature strings and verify those signature strings on the server. The AccessKey secret is used to authenticate a user's identity and must be kept confidential.

Auto Scaling authenticates each access request. Therefore, each request must contain signature information, regardless of whether it is sent through HTTP or HTTPS. Auto Scaling uses AccessKey pairs to implement symmetric encryption and authenticate the identity of a request sender.

The AccessKey ID and AccessKey secret are officially issued by Alibaba Cloud to users. You can request and manage them on the Alibaba Cloud official website. The AccessKey ID indicates the identity of a user. The AccessKey secret is the key used to encrypt the signature string and verify the signature string on the server. The AccessKey secret must be kept confidential.

#### 5.3.1.2.2. Access control

Resource Access Management (RAM) is a service that Alibaba Cloud provides to you to manage user identities and to control resource access. You can use RAM to create and manage user accounts, such as employee accounts, system accounts, and application accounts. You can also manage the operation permissions that these user accounts have on resources of your account. If multiple users in your enterprise operate resources collaboratively, RAM allows you to grant permissions to other users without sharing the AccessKey pair of your Apsara Stack tenant account with other users. Instead, you can grant users the minimum permissions necessary for them to complete their work. Thus, security risks to your enterprise information are reduced.

RAM allows you to create different roles and assign different permissions on cloud services to each role. Auto Scaling allows you to configure the RamRoleName parameter. You can configure this parameter to assign different roles to your ECS instances, allowing different instances to have permissions on different cloud services. Before configuring the RamRoleName parameter in Auto Scaling, you must ensure that the current RAM permission policy allows your ECS instance to act as the specified role. Otherwise, the scaling configuration cannot make the ECS instance available.

## 5.3.2. Tenant security

### 5.3.2.1. Log audit

ESS generates scaling activity logs that record information about each scaling activity, such as activity ID, status, status information, start time, end time, reason for activity, and details.

The states of a scaling activity include Rejected, In Progress, Successful, Warning, and Failed. Status information includes the status details. Reason for activity includes the results of scaling activities executed in a scaling group. Details include information about instances involved in a scaling activity.

# 5.4. Resource Orchestration Service (ROS)

## 5.4.1. Platform security

### 5.4.1.1. Data security

None

### 5.4.1.2. Authentication

ROS supports RAM authentication.

Resource Access Management (RAM) is a resource access control service provided by Alibaba Cloud. You can use an Apsara Stack tenant account to create RAM users and grant them permissions on resources that belong to the tenant account.

## 5.4.2. Tenant security

### 5.4.2.1. Log audit

ROS displays detailed information about historical events, including event logs for stacks. The information includes the resource name, associated resource ID, resource type, resource status, status description, and event occurrence time. Event logs provide the information about changes to stacks.

# 5.5. Object Storage Service (OSS)

## 5.5.1. Platform security

# 5.5.1.1. Security isolation

OSS slices user data and discretely stores the sliced data in a distributed file system based on specific rules. The user data and its indexes are stored separately. OSS uses symmetric AccessKey pairs to authenticate users and verifies the signature in each HTTP request sent by users. If verification is successful, OSS reassembles the distributed data. This way, OSS implements data storage isolation between multiple tenants.

# 5.5.1.2. Authentication and access control

# 5.5.1.2.1. Authentication

You can create an AccessKey pair on Apsara Stack Management Console. An AccessKey pair consists of an AccessKey ID and an AccessKey secret. The AccessKey ID is a public ID that uniquely identifies a user. The AccessKey secret is private and used to authenticate a user.

Before you send a request, you must generate a signature string for the request in the format specified by OSS. Then, you must encrypt the signature string by using your AccessKey secret to generate a verification code based on the HMAC algorithm. The verification code is timestamped to prevent replay attacks. After receiving the request, OSS finds the AccessKey secret based on your AccessKey ID, and uses the AccessKey secret to decrypt the signature string and verification code. Then, OSS calculates a verification code and compares it with the decrypted verification code. If the two verification codes are the same, OSS determines that the request is valid. Otherwise, OSS rejects the request and returns HTTP 403.

# 5.5.1.2.2. Access control

OSS supports access control list (ACL) to control access permissions. An ACL is set based on resources. You can specify ACLs for buckets or objects. You can specify an ACL for a bucket when you create the bucket or for an object when you upload the object to OSS. You can also modify the ACLs of uploaded objects and created buckets.

Access to OSS resources can be initiated by the bucket owner or third party users. An owner owns a bucket. Third-party users are other users who access resources in the bucket. Access can be either anonymous or signed. If the access is initiated with an OSS request that does not contain identification information, the access is considered to be anonymous. A signed access is a request that contains signature information in the header or a URL that contains signature information as defined in OSS API documentation.

OSS provides access control for buckets and objects.

You can configure one of the following ACLs for a bucket:

- Public read/write: All users (including anonymous users) can perform write (PutObject, GetObject, and DeleteObject) operations on objects in the bucket.
- Public read: Only the bucket owner or authorized users can perform write operations (PutObject and DeleteObject) on objects in the bucket. Other users, including anonymous users, can only perform read operations (GetObject) from the objects in the bucket.
- Only the bucket owner or authorized users can perform read and write operations (PutObject, GetObject, and DeleteObject) on objects in the bucket. Other users cannot access the objects in the bucket without authorization.

> **Note** If you do not configure the ACL of a bucket when you create the bucket, OSS sets the ACL of the bucket to private.

You can configure one of the following ACLs for an object:

- Public read/write: All users can perform read/write operations on the object.
- Public read: Only the object owner can perform read/write operations on the object. Others can perform read operations on the object.
- Private: Only the object owner can perform read/write operations on the object. Others cannot access the object.
- Default: The object inherits the ACL of the bucket.

> **Note** If you do not configure the ACL of an object when you upload the object, OSS sets the ACL of the object to default.

# 5.5.1.2.3. Support for RAM and STS

OSS supports Resource Access Management (RAM) and Security Token Service (STS) authentication.

RAM is a resource access control service provided by Apsara Stack. RAM allows you to create RAM users under an Apsara Stack tenant account. The Apsara Stack tenant account can grant access permissions on resources to RAM users.

STS is service that provides temporary access credentials. You can use STS to generate a temporary access credential for a user and specify the permission and validity period of the credential. A credential becomes invalid after it expires.

# 5.5.1.3. Data security

An error may occur when data is transferred between the client and server. OSS supports CRC and MD5 verification to secure data.

## CRC

OSS can return the CRC64 value of objects uploaded through any of the methods provided. The client can compare the CRC64 value with the locally calculated value to verify data integrity.

OSS calculates the CRC64 value for newly uploaded objects and stores the result as metadata of the object. OSS then adds the x-oss-hash-crc64ecma header to the returned response header, indicating its CRC64 value. This CRC64 value is calculated based on Standard ECMA-182.

## MD5 verification

To check whether the object uploaded to OSS is consistent with the local file, attach the Content-MD5 field value to the upload request. The OSS server verifies the MD5 value. The upload can succeed only when the MD5 value of the object received by the OSS server is the same as the Content-MD5 field value. This method can ensure the consistency between objects.

# 5.5.1.4. Data encryption

#### 5.5.1.4.1. Server-side encryption

OSS supports server-side encryption for uploaded data. When you upload data, OSS encrypts the data by using AES256 and permanently stores the encrypted data. When you download the data, OSS automatically decrypts the data, returns the original data, and declares in the header of the returned HTTP request that the data had been encrypted on the server.

To encrypt an object on the OSS server when you upload the object, you only need to add the x-oss-server-side-encryption header in the PutObject request and set its value to AES256.

#### 5.5.1.4.2. Client-side encryption

OSS allows you to use client-side encryption to encrypt data before the data is sent to the server while the data encryption key (DEK) used is kept only on the local client. Other users cannot obtain the raw data without the DEK and enveloped data key (EDK), even if the data is leaked. OSS uses functions provided by SDKs to encrypt the data on local clients before the data is uploaded to the OSS bucket.

## 5.5.2. Tenant security

### 5.5.2.1. Key management

Apsara Stack Key Management Service (KMS) is a secure and highly available service that integrates hardware and software, and provides a key management system that can be extended to the cloud. KMS uses customer master keys (CMKs) to encrypt OSS objects and uses KMS API operations to generate data encryption keys (DEKs) in a centralized manner. You can define policies in KMS to control and monitor key usage. You can use these keys to protect data in OSS buckets.

### 5.5.2.2. Log audit

OSS automatically saves access logs. After access logging is enabled for a source bucket, OSS generates an object that contains access logs for that bucket (by hour), names the object based on predefined naming rules, and writes the object into the bucket specified by the user. These logs are used for later auditing and behavior analysis. Request logs contain information such as the request time, source IP address, request object, return code, and processing duration.

### 5.5.2.3. Configure hotlink protection

To prevent additional fees caused by unauthorized access to the resources in your bucket, you can configure hotlink protection for your buckets on the Apsara Stack console or by using API operations.

You can set the following parameters to configure hotlink protection:

- Referer Whitelist: Only specified domain names are allowed to access OSS resources.
- Allow Empty Referer: If this parameter is disabled, a request is allowed to access OSS resources only if the request includes the Referer field configured in the HTTP or HTTPS header.

For example, for a bucket named oss-example, you can add http://www.aliyun.com/ to the Referer whitelist. Requests in which the Referer field is http://www.aliyun.com/ can access the objects in this bucket.

# 5.6. Apsara File Storage NAS

## 5.6.1. Platform security

### 5.6.1.1. Security isolation

#### Network isolation

NAS provides the permission group mechanism to control the networks over which NAS instances can be accessed. You can add rules to a permission group of a NAS instance to allow users from specified IP addresses or address segments to access the NAS instance with different permissions. In this way, networks are isolated from each other.

#### Storage isolation

In NAS, each mount point instance of a file system is mapped to a storage unit in the server storage pool. The storage units corresponding to different mount point instances are isolated from each other.

The access control module on the NAS server verifies the I/O requests of users based on the mapping between VPCs and NAS mount point instances. The module checks the storage unit information carried by each request against the storage unit information on the server for consistency. This ensures storage isolation on the server.

### 5.6.1.2. Authentication

#### Permission control

NAS allows you to perform standard directory or file permission operations on a NAS instance. You can also configure read, write, or execution permissions for a user or user group. NAS supports two types of mount points: VPC and classic network. If you configure a VPC mount point for a NAS instance, only the ECS instances within the same VPC as the mount point can access the NAS instance. If you configure a classic network mount point, only the ECS instances under the same account as the mount point can access the NAS instance.

In NAS, the permission group acts as an IP address whitelist. You can add rules to the permission group of a NAS instance to allow users from specified IP addresses or address segments to access the NAS instance with different permissions.

VPC Default Permission Group is automatically generated for each account by default. This permission group allows all IP addresses in the VPC to access the mount point with full permissions. Full permissions include read and write permissions with no restrictions on the root user.

> **Note**
>
> - Mount points in a classic network do not have a default permission group.
> - When you add a permission group rule for a mount point in a classic network, you can only set the authorized IP address to a single IP address. You cannot set the authorized IP address to an IP address segment.

A permission group rule has four attributes, as listed in the following table.

## Permission group rule attributes

| Attribute | Value | Definition |
| --- | --- | --- |
| Authorized IP Address | An IP address or IP address segment (You must specify an IP address for a permission group rule of a mount point in a classic network.) | The IP address or IP address segment of the authorized object. |
| Read and Write Permission | • Read Only<br>• Read/Write | The operation permissions of the authorized object on the NAS instance. |
| User Permission | • Do Not Limit root User<br>• Limit root User<br>• Limit All Users | Whether to restrict the permissions of the authorized object's Linux system users on the NAS instance.<br>Description:<br>• Do Not Limit root User allows the root user to access the NAS instance.<br>• Limit root User considers the root user as nobody.<br>• Limit All Users considers all users including root as nobody. |
| Priority | Valid values: 1 to 100. 1 indicates the highest priority. | When an authorized object matches multiple rules, the rule with the highest priority takes effect. |

## Access control

NAS can work with RAM. You can make RAM settings in the NAS console to complete RAM authorization.

RAM allows you to grant the permissions on NAS instances to RAM users.

## NAS operation permissions that can be granted to RAM users

| Action | Description |
| --- | --- |
| DescriptFileSystems | Lists NAS instances. |

| Action | Description |
| --- | --- |
| DescriptMountTargets | Lists the mount points of a NAS instance. |
| DescriptAccessGroup | Lists the permission groups of a NAS instance. |
| DescriptAccessRule | Lists permission group rules. |
| CreateFileSystem | Creates a NAS instance. |
| CreateMountTarget | Adds a mount point to a NAS instance. |
| CreateAccessGroup | Creates a permission group. |
| CreateAccessRule | Adds a permission group rule. |
| DeleteFileSystem | Deletes a NAS instance. |
| DeleteMountTarget | Deletes a mount point. |
| DeleteAccessGroup | Deletes a permission group. |
| DeleteAccessRule | Deletes a permission group rule. |
| ModifyMountTargetStatus | Disables or enables a mount point. |
| ModifyMountTargetAccessGroup | Modifies the permission group of a mount point. |
| ModifyAccessGroup | Modifies a permission group. |
| ModifyAccessRule | Modifies a permission group rule. |

# 5.6.1.3. Data security

## Multi-copy data storage

NAS maintains multiple data copies to ensure data security.

User data: The NAS server stores three copies of user data. Services can continue running even when two copies are lost. The server monitors the number of copies in real time. When a data node is corrupted or a hard drive on a data node fails, the number of valid copies of some data in the cluster becomes less than 3. In this case, the server activates the replication mechanism to replicate data. This mechanism ensures that there are always three valid copies of each piece of data in the cluster.

The server prevents accidental silent errors by verifying that the stored data matches the check data. When the server detects a silent error, it replicates healthy copies to ensure the availability of three valid copies. This improves data reliability.

## Data reclamation

The server reclaims the storage space that is released by the Delete operation. This reclaimed storage space is inaccessible to all users. The data stored in the storage space is erased before the space is reused. This fully guarantees data security.

# 5.6.2. Tenant security

## 5.6.2.1. Log audit

The NAS management system logs NAS instance operations, including creating and deleting NAS instances.

NAS logs are generated in real time and automatically stored on the server. A log contains detailed information about an operation, such as the executor and execution time. This information can be used for failure investigation and analysis.

## 5.6.2.2. Directory-level ACLs

This topic describes how to configure directory-level access control lists (ACLs). Only directory-level ACLs are available for Apsara File Storage NAS file systems.

### Prerequisites

- You must mount NFSv4 file systems on all clients.
- You must use the alinas-acl tool to configure ACLs. We recommend that you do not change the file mode creation mask or use commands such as `chmod` to change file permissions. Otherwise, you may not obtain the expected results.

### Procedure

1. The syntax of the mount command is **sudo mount -t nfs -o vers=4.0 <the domain name of a mount target>:<the directory of an NAS file system> <a local directory>**. For example, you can use **mount -t nfs -o vers=4.0 014544bbf6-wdt41.cn-hangzhou.nas.aliyuncs.com:/ /mnt** to mount an NFSv4 file system.

   > **Note**
   > - The value of the vers parameter changes based on the client version. If an error occurs when you set vers to 4.0, set vers to 4 instead.
   > - In some cases, ACLs are not enabled for a file system by default. To ensure ACLs are available for the file system, you can mount the file system again to enable ACLs.

2. Install the `nfs4-acl-tools` tool in CentOS.

   ```
   sudo yum -y install nfs4-acl-tools
   ```

3. Make sure that Python 2.7 is installed.

   ```
   python --version Python 2.7.5
   ```

4. Use the alinas-acl tool to configure an ACL.

```
./alinas_acl set ./foo --add --user Alice --rule r #Grant the Alice user the read-only access to the f
oo file.
./alinas_acl set ./foo -a -u Alice -r r #You can use the command to perform the same operation as
the preceding command.
./alinas_acl set ./dir --add --group Staff --rule rwx #Grant the Staff group the read, write, and exe
cute access to the dir directory.
./alinas_acl set ./foo --add --user EVERYONE@ --rule none #Grant the EVERYONE@ principal no acc
ess to the foo file.
./alinas_acl set ./foo --add --user 1001 --rule none #Grant the 1001 user principal no access to the
foo file.
./alinas_acl set ./dir -d -u Bob #Revoke the Bob user access to the dir directory.
```

> **Note**
>
> - To avoid a decrease in performance, we recommend that you configure an ACL for a directory rather than each file in the directory.
> - We recommend that you add a maximum of 10 access control entries (ACEs) to an ACL.

5. View the ACL.

```
./alinas_acl get ./foo #View the permissions on the foo file. # file: foo/ # owner:
        root # group: root OWNER@::rw- GROUP@::r-- EVERYONE@::--- Alice::r-- Staff:g:rwx
        1001::---
```

> **Note** When you configure an ACL, three special principals named OWNER@, GROUP@, and EVERYONE@ are automatically generated. These principals correspond to the user, group, and others classes of a file mode creation mask, respectively. The permissions that you specify for the file mode creation mask can be different from the permissions that you specify for an ACL. The actual permissions change based on the client version.

# 5.7. Tablestore

## 5.7.1. Platform security

### 5.7.1.1. Security isolation

This topic describes the security isolation methods of Tablestore, including network and storage isolation.

#### Network isolation

Tablestore supports instance-level VPC access control. The following types of VPC access settings are supported:

- Allows all network access: Access from the Internet and VPCs bound to the instance is allowed.
- Allows access from specific VPCs: Only access from VPCs bound to the instance is allowed.
- Allows access from the console and specific VPCs: Only access from VPCs bound to the instance and the Tablestore console is allowed. Access from other sources is denied.

### Storage isolation

Tablestore uses a shared storage mechanism. This mechanism allows the instances of different users to share the same cluster resource. Tablestore uses data partitions as the smallest unit and supports the load balancing mechanism at the data partition level to isolate the impact between different instances.

## 5.7.1.2. Authentication

This topic describes the authentication methods of Tablestore, including authentication and access control.

### Authentication

Tablestore authenticates requests based on AccessKey pairs. Each valid Tablestore request must contain the correct AccessKey pair information. Tablestore authenticates each request from applications to prevent unauthorized data access and ensure data security.

### Access control

Tablestore supports Security Token Service (STS), which allows you to control access of RAM users. STS is a temporary access credential service provided by Apsara Stack. It provides temporary access control. You can use STS to generate a temporary access credential. You can specify the permissions and validity period of the credential. The credential becomes invalid when it expires.

Tablestore supports authorization based on tables and API operations.

## 5.7.1.3. Data security

This topic describes the data security policies of Tablestore.

Tablestore is built on the Apsara Distributed File System and provides linear storage space. Linear addresses are sliced into chunks. For each chunk, three replicas are created and stored in different nodes in the cluster to ensure data reliability.

In Tablestore, data is serialized before it is written to the disks. Each data block is written to one or more chunks.

Apsara Distributed File System evaluates the disk usage of all nodes, the distribution of these nodes on different racks, the power supply, and the host loads to ensure that the chunk replicas are distributed to different hosts across different racks. This prevents host or rack faults from affecting service availability.

When a data node is damaged or a disk fault occurs on a data node, the chunk replica number becomes smaller than three. The Apsara Distributed File System starts the automatic replication process to replicate data among different service nodes when the replica number is smaller than three. This ensures that each chunk in the cluster has three valid replicas.

Write operations in Tablestore can be returned only after all three replicas are written to the disks. This ensures strong data consistency.

# 5.7.2. Tenant security

## 5.7.2.1. Key management

This topic describes how to manage keys. You can use keys to protect data stored in Tablestore.

Apsara Stack Key Management Service (KMS) is a secure and high-availability service that integrates hardware and software and provides a key management system that can be extended to the cloud.

KMS uses customer master keys (CMKs) to encrypt Tablestore tables and uses the KMS API to generate data encryption keys (DEKs) in a centralized manner. You can define policies in KMS to control and monitor key usage. You can use CMKs to protect data stored in Tablestore.

# 5.8. ApsaraDB for RDS

# 5.8.1. Platform security

## 5.8.1.1. Secure isolation

### Tenant isolation

ApsaraDB for RDS uses virtualization technology to isolate tenants. Each tenant can maintain their own database permissions independently. Alibaba Cloud also implements increased security for servers that run databases to prevent other users from accessing your data. For example, databases cannot read or write system files.

## 5.8.1.2. Authentication

ApsaraDB for RDS secures data through authentication.

### Identity authentication

Account authentication uses your logon password or AccessKey pair to verify your identity. You can create an AccessKey pair from Apsara Stack Management Console. An AccessKey pair consists of AccessKey ID and AccessKey Secret. AccessKey ID is a public key used for identification. AccessKey Secret is used to encrypt signature strings sent from the client and verify signature strings sent by the server. You must keep your AccessKey Secret confidential.

The ApsaraDB for RDS server authenticates the sender identity of each access request. Because of this, each request must contain signature information, regardless of whether it is sent using HTTP or HTTPS. ApsaraDB for RDS uses AccessKey ID and AccessKey Secret to implement symmetric-key encryption and authenticate the identity of a request sender. AccessKey pairs can be applied for and managed from the Apsara Stack. The AccessKey Secret will only be known to you, so it is necessary to take precautions to keep it confidential.

### Permission control

ApsaraDB for RDS does not automatically create initial database accounts for a newly created instance. You can use the console or API to create a standard database account and configure database-level read and write permissions. To implement fine-grained permission control, such as table-level, view-level, or field-level permissions, you can use the console or API to create a master database account. You can then use the database client and master database account to create standard database accounts. A master database account can configure table-level read/write permissions for standard database accounts.

### Access control

All ApsaraDB for RDS instances that are created by an Apsara Stack tenant account are managed as resources by that account. By default, an Apsara Stack tenant account is granted full operation permissions on all resources belonging to the account.

ApsaraDB for RDS supports Resource Access Management (RAM). You can use RAM to allow RAM users to access and manage RDS resources under your account. ApsaraDB for RDS can also provide short-term access permissions with temporary credentials provided through STS.

## 5.8.1.3. Data security

ApsaraDB for RDS secures data through hot standby, data backups, and log backups.

High-availability ApsaraDB for RDS instances implement two database nodes for hot standby. When the primary node fails, the secondary node immediately takes over services. Database backups can be initiated anytime. To improve data traceability, ApsaraDB for RDS can restore data to any previous point in time based on the backup policy.

Automatic backup at regular intervals is required to guarantee the integrity, reliability, and restorability of databases. ApsaraDB for RDS provides two backup functions: data backup and log backup.

## 5.8.1.4. Data encryption

### SSL

ApsaraDB for RDS provides Secure Sockets Layer (SSL) for MySQL, SQL Server, PolarDB, and PostgreSQL. You can prevent man-in-the-middle attacks by using the server root certificate to verify whether the destination database is an ApsaraDB for RDS instance. ApsaraDB for RDS also allows you to enable and update SSL certificates for servers to ensure security and validity.

Although ApsaraDB for RDS can encrypt the connection between an application and a database, SSL cannot run properly until the application authenticates the server. SSL consumes extra CPU resources, which affects the throughput and response time of instances. The severity of the impact depends on the number of user connections and the frequency of data transfers.

## 5.8.1.5. DDoS attack prevention

ApsaraDB for RDS prevents DDoS attacks by using the traffic scrubbing and black hole filtering features.

When you access an ApsaraDB for RDS instance from the Internet, the instance is vulnerable to DDoS attacks. When a DDoS attack is detected, the RDS security system first scrubs inbound traffic. If traffic scrubbing is insufficient or if the black hole threshold is reached, black hole filtering is triggered.

Triggering conditions for traffic scrubbing and black hole filtering are listed as follows:

- **Traffic scrubbing**

  Traffic scrubbing only targets traffic from the Internet. Traffic is redirected from an IP address to the scrubbing device, which then checks whether the traffic is normal. Abnormal traffic is discarded and traffic to the server is limited by the scrubbing device to mitigate damage on the server. These operations may have an impact on normal traffic.

  ApsaraDB for RDS triggers and stops traffic scrubbing automatically. Traffic scrubbing is triggered for a single ApsaraDB for RDS instance if any of the following conditions are met:

  - Packets per second (PPS) reaches 30,000.
  - Bits per second (BPS) reaches 180 Mbit/s.
  - The number of new concurrent connections per second reaches 10,000.
  - The number of active concurrent connections reaches 10,000.
  - The number of inactive concurrent connections reaches 10,000.

- **Black hole filtering**

  Black hole filtering only targets traffic from the Internet. If an RDS instance is undergoing black hole filtering, the instance cannot be accessed from the Internet and connected applications will not be available. Black hole filtering is triggered for a single ApsaraDB for RDS instance if any of the following conditions are met:

  - BPS reaches 2 Gbit/s.
  - Traffic scrubbing is ineffective.

  Black hole filtering is automatically stopped 2.5 hours after being triggered. Then, the instance will undergo traffic scrubbing. If the DDoS attack is still occurring, black hole filtering is triggered again. Otherwise, the system restores the normal state.

# 5.8.2. Tenant security

## 5.8.2.1. Log audit

ApsaraDB for RDS can audit logs to identify security issues.

ApsaraDB for RDS allows you to view SQL transactions and periodically audit the SQL server to identify and resolve issues. RDS Proxy records all SQL statements sent to ApsaraDB for RDS, including the IP address, database name, user account used for execution, execution period, number of returned records, and execution time of each statement.

## 5.8.2.2. IP address whitelist

ApsaraDB for RDS uses the IP address whitelist to prevent access from invalid IP addresses.

ApsaraDB for RDS instances can be accessed from any IP address by default. Because of this, the IP address whitelist contains only the entry 0.0.0.0/0. You can add IP address whitelist rules through the data security module in the console or by calling an API. The IP address can be updated without restarting the ApsaraDB for RDS instance. Whitelist updates will not affect the normal operation of the instance. Multiple groups can be configured in the IP address whitelist. Each group can contain up to 1,000 IP addresses or IP address segments.

#### 5.8.2.3. Software update

ApsaraDB for RDS supports post-restart update and mandatory update for software.

ApsaraDB for RDS automatically provides you with new versions of installed database software. In most cases, it is not required to update software immediately. Only when you manually restart an ApsaraDB for RDS instance does the system update the database software to the latest compatible version.

In rare cases such as critical bugs and security vulnerabilities, ApsaraDB for RDS will force the database to update during the maintenance period of the instance. Such mandatory updates only result in temporary database disconnections, and will not have any adverse impact on the application if the database connection pool is configured properly.

You can use the console or API to change the maintenance schedule to prevent a mandatory update from occurring during peak hours.

# 5.9. Cloud Native Distributed Database PolarDB-X

## 5.9.1. Platform security

### 5.9.1.1. Security isolation

#### Network isolation

PolarDB-X supports advanced control of network access by using a Virtual Private Cloud (VPC).

A VPC is a private network environment that you set. It strictly isolates network packets through underlying network protocols, and it controls access at the network layer. The VPC and IP address whitelist together greatly improve the security of PolarDB-X instances.

### 5.9.1.2. Authentication

PolarDB-X provides a system to manage accounts and permissions, similar to that of MySQL. This system supports commands and functions such as GRANT, REVOKE, SHOW GRANTS, CREATE USER, DROP USER, and SET PASSWORD.

When you create a PolarDB-X database, by default, you can specify an account with all permissions. You can use this account to create one or more new accounts.

- You can grant permissions at the database and table levels. Currently, global permissions and column-level permissions are not supported.
- These eight statements of associated basic permissions are supported: CREATE, DROP, ALTER, INDEX, INSERT, DELETE, UPDATE, and SELECT.
- You can use `user@'host'` to match and verify access to a host.

> **Note** However, if the business host is in a Virtual Private Cloud (VPC) network, the IP address cannot be obtained due to technical restrictions. In this case, we recommend that you change the format to `user@'%'`.

# 5.9.2. Tenant security

## 5.9.2.1. IP address whitelist

PolarDB-X provides IP address whitelists to ensure secure access. You can configure an IP address whitelist for each PolarDB-X database.

The default setting of PolarDB-X instances allows access from any IP address. You can add IP addresses to the whitelist on the **Whitelist Settings** page in the console. You are required to restart PolarDB-X instance after you update the IP address whitelist, and your operations on the instance are not affected. You can set IP addresses or CIDR blocks in the IP address whitelist.

> **Note** If the business host is in a Virtual Private Cloud (VPC) network, the IP address cannot be obtained due to technical restrictions. We recommend that you remove the IP address whitelist.

## 5.9.2.2. Protection against high-risk SQL operations

PolarDB-X prohibits high-risk operations such as full table deletion and full table update by default. You can temporarily skip this restriction by adding a hint. The following statements are prohibited by default:

- DELETE statements that do not contain the WHERE or LIMIT conditions.
- UPDATE statements that do not contain the WHERE or LIMIT conditions.

For example, the following statement is prohibited:

```
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]
```

After a hint is added, the statement is successfully executed.

```
mysql> /*TDDL:FORBID_EXECUTE_DML_ALL=false*/delete from tt;
Query OK, 10 row affected (0.21 sec)
```

## 5.9.2.3. Slow SQL audit

In the PolarDB-X console, you can query the slow SQL statements sent by an application to PolarDB-X. Slow SQL statements increase the response time (RT) of the entire link and reduce the throughput of PolarDB-X.

Contents of a slow SQL statement include the execution start time, database name, SQL statement, client IP address, and execution time. You can query details of slow SQL statements in the PolarDB-X console for optimization and adjustment.

#### 5.9.2.4. Performance monitoring

The PolarDB-X console provides monitoring metrics in different dimensions. You can perform related operations based on the monitoring information.

There are two types of PolarDB-X monitoring information:

- Monitoring information about resources, including the CPU, memory, and network.
- Monitoring information about engines, including the logical queries per second (QPS), physical QPS, logical response time (RT) in milliseconds, physical RT in milliseconds, number of connections, and number of active threads.

The QPS and CPU performance of a PolarDB-X instance are in positive correlation. When PolarDB-X encounters a performance bottleneck, the CPU utilization of the PolarDB-X instance remains high. If the CPU utilization exceeds 90% or remains above 80%, the PolarDB-X instance faces a performance bottleneck. If there is no bottleneck in the PolarDB-X instance, the curent type of the PolarDB-Xinstance cannot meet the QPS performance requirements of the business. In this case, upgrade the instance.

# 5.10. AnalyticDB for MySQL

## 5.10.1. Platform security

### 5.10.1.1. Security isolation

AnalyticDB for MySQL isolates tenants by database. The Apsara Stack tenant account that is used to create a database is the owner of the database. To allow other users to access the database, the database owner must first grant access permissions to the other users. The database that is created by each Apsara Stack tenant account runs on a dedicated process, which isolates databases at the process level.

AnalyticDB for MySQL uses a multi-tenancy architecture to provide the database of each Apsara Stack tenant account with a dedicated process. In the multi-tenancy architecture, physical resources, such as CPU, memory, and storage space, are isolated between databases.

AnalyticDB for MySQL allows you to manage the version of each database, scale database resources, and start and stop database services.

### 5.10.1.2. Authentication

The following figure shows how AnalyticDB for MySQL implements identity authentication and access control.

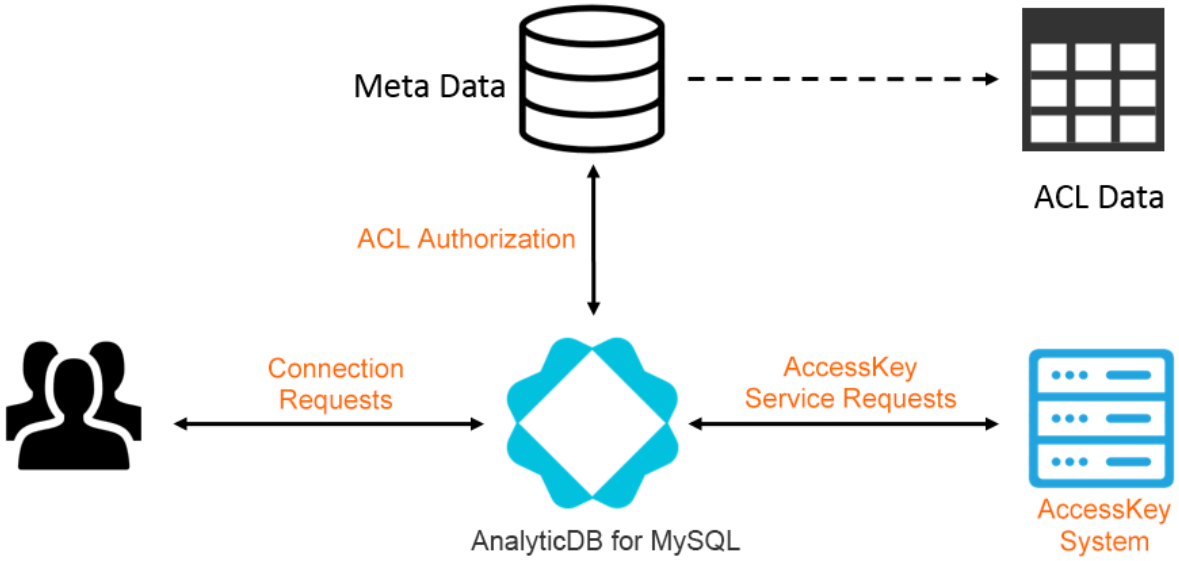


## Identity authentication

AnalyticDB for MySQL provides a MySQL protocol-based identity authentication system with username and password authentication.

AnalyticDB for MySQL uses an AccessKey pair to implement identity authentication. You can connect to AnalyticDB for MySQL by using an AccessKey pair. You can also use the AccessKey pair to establish a Java Database Connectivity (JDBC) connection or an Open Database Connectivity (ODBC) connection to AnalyticDB for MySQL.

You can create an AccessKey pair in the Apsara Stack Cloud Management (ASCM) console. Each AccessKey pair consists of an AccessKey ID and an AccessKey secret. These credentials are similar to a username and password. The AccessKey ID can be publicly shared and used to identify a user. The AccessKey secret is used to authenticate the identity of the user and must be kept confidential.

You can connect to AnalyticDB for MySQL by using the AccessKey ID and AccessKey secret of your Apsara Stack tenant account or a Resource Access Management (RAM) user.

## Permission control

AnalyticDB for MySQL uses an access control list (ACL) to manage table-level permissions. The ACL rules are similar to those of MySQL. However, AnalyticDB for MySQL does not allow you to use an ACL to grant permissions to specific hosts, which is different from MySQL.

An ACL lists authorized users as well as their authorization objects and operation permissions. ACL data is stored in the AnalyticDB for MySQL metadata system and uses ApsaraDB for RDS to ensure data persistence. AnalyticDB for MySQL caches metadata to accelerate authorization for Data Manipulation Language (DML) and Data Definition Language (DDL) operations.

After you connect to AnalyticDB for MySQL, AnalyticDB for MySQL uses the ACL metadata to control your operation permissions on database objects. For example, AnalyticDB for MySQL determines whether you can perform SELECT, INSERT, DELETE, CREATE, SHOW, DROP, ALTER, DESCRIBE, LOAD DATA, or DUMP DATA operations on a specific table or column.

AnalyticDB for MySQL provides the following authorization objects:

- Database: specifies a database or all tables in a database, such as `db_name.*` or `*` (default

database).

- Table: specifies a table, such as `db_name.table_name` or `table_name`.
- Column: specifies a column in a specified table. It consists of `column_list` and `Table`.

### Access control

AnalyticDB for MySQL implements access control by using RAM. Access control based on Security Token Service (STS) is not supported.

RAM allows you to create RAM users by using an Apsara Stack tenant account. You can then grant resource access permissions to RAM users. RAM users are affiliated with their parent Apsara Stack tenant account. All resources that are created by RAM users belong to their parent Apsara Stack tenant account.

## 5.10.1.3. Data security

### Multi-tenancy

AnalyticDB for MySQL uses a multi-tenancy architecture to isolate resources such as CPU, memory, disk space, and network bandwidth between databases.

### Data reliability

All AnalyticDB for MySQL data is stored in Apsara Distributed File System. Apsara Distributed File System ensures high reliability and data persistence by using three-replica redundancy or erasure code (EC). After Data Manipulation Language (DML) operations such as INSERT and DELETE are performed on a real-time table, updates are synchronized to Apsara Distributed File System. Full data is written to Apsara Distributed File System during batch loading.

### Data consistency

AnalyticDB for MySQL uses the Multi-Version Concurrency Control (MVCC) method to store changes to real-time tables when you perform INSERT and DELETE operations. If the table you query has concurrent data updates, MVCC ensures that data is queried from the snapshot of data that was taken when the query was initiated.

> **Note** You can clear outdated data versions at regular intervals.

# 5.10.2. Tenant security

## 5.10.2.1. Log audit

You can enable log audit to record information about all SQL operations in AnalyticDB for MySQL. The information includes the following items:

- Query time
- IP address of the client
- Executed SQL statements

You can use SQL statements to query historical data.

Sample audit log:

```
[2017-10-10 13:37:57,351] INFO [pool-31-thread-22] c.a.c.a.f.l.AccessLog.info - Client=127.0.0.1 Total_time=1044 Exec_time=1043 Queue_time=1 - [2017-10-10 13:37:56 308] 1 SQL Statement \;process=20171010133756010003163108099998380 42\;CLUSTER=ayads-bjyz
```

# 5.11. AnalyticDB for PostgreSQL

## 5.11.1. Platform security

### 5.11.1.1. Security isolation

#### Network isolation

In Apsara Stack, you can use IP address whitelists to control access. You can also use a VPC to control network access.

A VPC is a private network environment that you can set in Apsara Stack. It strictly isolates network packets by using underlying network protocols and controls access at the network layer.

By default, AnalyticDB for PostgreSQL instances in a VPC are only accessible from the ECS instances within the same VPC. You can also apply for a public IP address to receive access requests from the Internet. This method is not recommended. The requests include but are not limited to:

- Access requests from ECS elastic IP addresses (EIPs).
- Access requests from the Internet egress of your data center.

> **Note** IP address whitelists apply to all connections to AnalyticDB for PostgreSQL instances. We recommend that you configure whitelists before you apply for a public IP address.

#### Tenant isolation

AnalyticDB for PostgreSQL uses virtualization technology to isolate tenants. Each tenant can maintain their own database permissions independently. Apsara Stack also enhances security for servers that run databases. For example, users cannot use the databases to read or write operating system files. This prevents other users from accessing your data.

### 5.11.1.2. Authentication

The AnalyticDB for PostgreSQL instances that you create by using your Apsara Stack tenant account are owned by the account. By default, Apsara Stack tenant accounts have full access permissions on their resources.

AnalyticDB for PostgreSQL supports Resource Access Management (RAM) and Security Token Service (STS). You can use RAM to grant access and management permissions on the AnalyticDB for PostgreSQL resources of your account to other RAM users. You can use STS to issue temporary access credentials to RAM users for short-term access to resources.

#### 5.11.1.3. Primary and secondary nodes

Each AnalyticDB for PostgreSQL instance consists of a coordinator node and multiple compute nodes. Each node uses a primary/secondary architecture. If the primary node fails, the service is quickly switched to the secondary node. You can back up databases at any time. AnalyticDB for PostgreSQL can restore data from backup sets based on backup policies to improve data traceability.

### 5.11.2. Tenant security

#### 5.11.2.1. Database account

After you create an instance, you can create a superuser account in the console or by using an API operation. You can execute the `GRANT` statement to authorize other database accounts.

#### 5.11.2.2. IP address whitelists

By default, AnalyticDB for PostgreSQL instances block access from all IP addresses. The default IP address whitelist contains only `127.0.0.1`. You can add IP addresses to a whitelist on the Security Controls page of the console or by using an API operation. The IP address whitelist can be updated without restarting the AnalyticDB for PostgreSQL instance. Whitelist updates do not affect the normal operation of the instance. You can configure multiple IP address whitelists. Each whitelist can contain up to 1,000 IP addresses or CIDR blocks.

#### 5.11.2.3. SQL audit

AnalyticDB for PostgreSQL allows you to view SQL details. You can audit SQL operations on a regular basis to identify problems in a timely manner. The Proxy module records the information of all SQL statements that are sent to AnalyticDB for PostgreSQL, including the connected IP addresses, names of the accessed databases, accounts used for statement execution, SQL statements, execution duration, number of returned records, and execution time points.

#### 5.11.2.4. Backup and restoration

For data integrity and reliability, a database must automatically back up data on a regular basis to ensure that data can be restored. AnalyticDB for PostgreSQL allows you to restore instances from backup sets.

#### 5.11.2.5. Software upgrade

- AnalyticDB for PostgreSQL provides new versions of database software on a regular basis.
- Software upgrade is optional. It is carried out only upon your request.
- If the current database version that you are using has critical security risks, the AnalyticDB for PostgreSQL team will notify you and recommend that you schedule the upgrade. The AnalyticDB for PostgreSQL team provides full support throughout the upgrade.
- The update of AnalyticDB for PostgreSQL is typically completed within five minutes. During the upgrade, network interruptions may occur and the database may become read-only for about one minute. If the database reconnection settings or connection pools are properly

configured, the upgrade has minimal impact on applications.

# 5.12. KVStore for Redis

## 5.12.1. Platform security protections

### 5.12.1.1. Security isolation

#### Tenant isolation

KVStore for Redis uses the virtualization technology to isolate tenants. Each tenant can maintain independent database permissions. Alibaba Cloud also increases security protections for the servers that run databases. For example, you cannot read from or write to system files by using the databases, so you cannot access other users' data.

#### Network isolation

In Apsara Stack, in addition to the whitelist, you can use Virtual Private Cloud (VPC) to restrict connections.

A VPC is a private network that you specify in Apsara Stack. The VPC strictly isolates your network packets based on network protocols and restricts connections at the network layer. You can use a virtual private network (VPN) or a leased line to connect server resources in your IDC to Alibaba Cloud, and use CIDR blocks in a VPC to prevent IP conflicts. In this way, your own servers and ECS instances can connect to KVStore for Redis instances at the same time. Protections based on the VPC and IP address whitelist improve the instance security.

By default, ECS instances in a VPC can only connect to KVStore for Redis instances in the same VPC. You can also request a public IP address to accept connections over a public network. We recommend that you do not use this connection method. The connection requests include but are not limited to:

- Those from ECS Elastic IP addresses (EIPs).
- Those from the public IP addresses in your own IDC.

> **Notice** The IP whitelist is applicable to all types of connections to KVStore for Redis instances. We recommend that you set the whitelist before requesting the public IP address.

### 5.12.1.2. Authentication

The instances that you create by using your Alibaba Cloud account are the resources under this account. By default, the Alibaba Cloud account is granted full operation permissions on all the resources under the account.

KVStore for Redis supports Resource Access Management (RAM) and Security Token Service (STS) services. By using RAM, you can create and manage RAM users. You can grant access and management permissions on KVStore for Redis resources under your Alibaba Cloud account to the RAM users. By using STS, you can manage short-term permissions granted to RAM users. You can use STS to grant permissions to temporary users.

#### 5.12.1.3. Transmission encryption

KVStore for Redis provides secure encryption based on the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols. You can use the server root certificate from KVStore for Redis to verify that KVStore for Redis provides database services based on the target IP address and port. This can effectively prevent man-in-the-middle attacks (MITM). Also, KVStore for Redis allows you to enable and update SSL and TLS certificates for servers. Therefore, you can replace the SSL or TLS certificate to ensure security and validity.

> Note
>
> - To use the transmission encryption feature, you must enable server verification in your application.
> - Transmission encryption consumes extra CPU resources and affects the throughput and response time of KVStore for Redis instances. The performance depends on the number of connections and the data transfer frequency.

### 5.12.2. Tenant security protections

#### 5.12.2.1. Database account

To connect to KVStore for Redis, you must pass password authentication. The password is the access credential. KVStore for Redis optimizes the performance of transient connections. Therefore, when you enable password authentication, the performance of KVStore for Redis instances is not affected.

#### 5.12.2.2. IP address whitelist

KVStore for Redis allows you to use an IP address whitelist to restrict connections and secure your data. You can set an IP address whitelist for each KVStore for Redis instance.

By default, KVStore for Redis instances block connections from all IP addresses or CIDR blocks. In this case, the IP address whitelist is set to `127.0.0.1`. To add an IP address or CIDR block to the whitelist, in the KVStore for Redis console, choose **Instance Information** > **Change Whitelist**. After you modify the IP address whitelist, you do not need to restart the instance, so you can still run the instance normally. You can specify multiple IP address groups for a whitelist. Each group contains a maximum of 1,000 IP addresses or CIDR blocks.

#### 5.12.2.3. Backup and recovery

Databases require regular automatic backups to guarantee data integrity, reliability, and restorability. KVStore for Redis supports instance recovery based on backup sets.

#### 5.12.2.4. Software upgrade

- KVStore for Redis regularly provides database upgrades.
- The upgrades are not mandatory. Databases upgrade to the specified version only when you request.

- When the KVStore for Redis team determines that your version has major security risks, KVStore for Redis notifies you to enable the upgrade. The KVStore for Redis team supports the whole upgrade process.
- KVStore for Redis completes the upgrade within five minutes. During the upgrade, temporary disconnections may occur, and the instance may stay in read-only status for one minute. If you have correctly configured the database reconnection or connection pool for your application, the upgrade does not affect your application.

# 5.13. ApsaraDB for MongoDB

## 5.13.1. Platform security

### 5.13.1.1. Isolation

#### Network isolation

ApsaraDB for MongoDB allows you to use a VPC for higher-level network isolation.

A VPC is a private network that you configure on Apsara Stack. It strictly isolates your network packets by using underlying network protocols to implement access control at the network layer.

#### Tenant isolation

ApsaraDB for MongoDB uses virtualization technologies to isolate tenants. Each tenant has separate database permissions. Apsara Stack also enhances the security of database servers. For example, Apsara Stack prohibits read and write operations on system files by using a database. This ensures that you cannot access the data of other users.

### 5.13.1.2. Authentication

#### Identity authentication

Account authentication uses an identity credential to verify the real identity of a user. An identity credential usually refers to a logon password or AccessKey. You can create AccessKey pairs in the ApsaraDB for MongoDB console. An AccessKey contains an AccessKey ID and an AccessKey Secret. The AccessKey ID is public and represents the user identity. The AssessKey Secret is used to encrypt the signature string. The server also uses the AccessKey Secret to verify the signature string. Make sure that you keep the AccessKey Secret confidential.

ApsaraDB for MongoDB performs identity authentication on each access request. Therefore, all HTTP and HTTPS requests must contain the signature information. ApsaraDB for MongoDB implements symmetric encryption through the Access Key ID and Access Key Secret to authenticate the request sender. You can apply for an AccessKey ID and an AccessKey Secret in the Apsara Stack console. The AccessKey ID represents the user identity. The AssessKey Secret is used to encrypt the signature string. The server also uses the AccessKey secret to verify the signature string. Make sure that you keep the AccessKey secret confidential.

#### Permission control

To log on to an ApsaraDB for MongoDB instance, you must pass the username and password authentication. After an ApsaraDB for MongoDB instance is created, a root account is created by default. You can either specify the password for the root account when you create an instance or reset the password after you create an instance.

The root account has all management permissions on an ApsaraDB for MongoDB instance. You can log on to the database as the root user to add or delete accounts, and grant permissions to other accounts.

### Access control

ApsaraDB for MongoDB instances that you create by using your Alibaba Cloud account are considered resources under this account. An Alibaba Cloud account has full operation permissions on its resources by default.

ApsaraDB for MongoDB supports Resource Access Management (RAM). RAM allows you to grant access and management permissions on the ApsaraDB for MongoDB resources under your account to RAM users.

## 5.13.1.3. Data security

ApsaraDB for MongoDB uses a high-availability architecture that features a three-node replica set. The three data nodes are located on different physical servers. The secondary and standby nodes automatically synchronize data from the primary node. Services are provided by the primary and secondary nodes. When the primary node fails, the system automatically selects a new primary node. When the secondary node fails, the standby node takes over the services.

ApsaraDB for MongoDB can automatically create backups. You can quickly restore data to ensure data integrity and reliability.

You can set the frequency to create full physical backups during one week (at least twice per week) and the start time and end time of the backups. In addition, you can perform full physical backups in the ApsaraDB for MongoDB console, or through APIs at any time based on your O&M requirements.

The system automatically backs up incremental logs generated by an ApsaraDB for MongoDB instance. The combination of full backups and incremental logs enables you to restore data to a specific second in time within the backup retention period.

## 5.13.1.4. Data encryption

### SSL

ApsaraDB for MongoDB provides Secure Sockets Layer (SSL). You can prevent man-in-the-middle attacks by using the server root certificate to verify whether the destination database is an ApsaraDB for MongoDB instance. ApsaraDB for MongoDB also allows you to enable and update SSL certificates for servers to ensure data security and validity.

## 5.13.1.5. Anti-DDoS

This feature monitors inbound network traffic in real time. When large volumes of malicious traffic is identified, it scrubs traffic through IP filtering. If traffic scrubbing fails, it triggers the black hole threshold.

# 5.13.2. Tenant security

## 5.13.2.1. Log audit

This feature records all operations that a client performs on a connected database. It provides references for fault analysis, behavior analysis, and security audit. This feature helps you obtain data execution information for analysis. Audit logs will gradually become an essential regulatory requirement of Finance Cloud and other core businesses.

## 5.13.2.2. IP address whitelists

ApsaraDB for MongoDB allows you to configure IP address whitelists for each ApsaraDB for MongoDB instance to implement network access control.

The default whitelist of an ApsaraDB for MongoDB instance contains 0.0.0.0/0, which indicates that the instance is accessible from all IP addresses. You can use OpenAPI Explorer or the ApsaraDB for MongoDB console to configure an IP address whitelist. Updating an IP address whitelist does not restart the instance, so your business is not affected.

# 5.14. ApsaraDB for OceanBase

# 5.14.1. Platform security

## 5.14.1.1. Security isolation

### Multi-tenant data isolation

ApsaraDB for OceanBase implements multi-tenant data isolation in databases. This allows one ApsaraDB for OceanBase cluster to serve multiple tenants. Cross-tenant data access is not allowed. This eliminates the risks of data leakage and data breach for each tenant. Each tenant has exclusive use of the resources that are allocated to the tenant. The frontend applications of each tenant offer stable performance. The performance is measured by the following performance metrics: response time, transactions per second (TPS), and queries per second (QPS). These performance metrics for each tenant are not affected by the service loads of the other tenants.

## 5.14.1.2. Authentication

### Strong authentication support

When you create a tenant, the system automatically creates a super account that has all the permissions. To log on to your database, use the "<Account>@<Tenant>#<Cluster name>" format. You can use the account that has all the permissions to create accounts and grant the created accounts different permissions.

- You can grant permissions based on the following levels:
  - Tenant level: The permissions apply to all the databases of the tenant.
  - Database level: The permissions apply to all the objects in the specified database.

  - Table level: The permissions apply to all the columns in the specified table.

- You can grant the created accounts the following basic permissions: CREATE, DROP, ALTER, INDEX, INSERT, DELETE, UPDATE, and SELECT.

## 5.14.1.3. High-availability architecture

In ApsaraDB for OceanBase, each data record is stored in more than 50% of at least three servers. For example, if three servers are used, each data record must be stored in two of the three servers. Each write transaction is valid only if the transaction is stored in more than 50% of all the servers. Therefore, no data loss occurs if only a minority of all the servers fail. This ensures that a recovery point objective (RPO) of zero can be achieved. In addition, ApsaraDB for OceanBase uses the Paxos protocol at the underlying layer to ensure high availability. If the primary server fails, a new primary server is automatically elected by the remaining servers based on the Paxos protocol. This ensures automatic switchovers and service continuity. In production environments, the recovery time objective (RTO) can be less than 30 seconds.

ApsaraDB for OceanBase retains multiple replicas and uses the Paxos protocol. This allows you to deploy ApsaraDB for OceanBase across data centers in different regions and implement high-availability features such as active geo-redundancy. ApsaraDB for OceanBase supports the following typical deployment solutions: Three Data Centers in One Region, Three Data Centers Across Two Regions, and Five Data Centers Across Three Regions. This allows ApsaraDB for OceanBase to meet the various business requirements for disaster recovery across data centers and zones.

## 5.14.1.4. Compatibility

ApsaraDB for OceanBase is compatible with most of the MySQL 5.6 features. This allows you to migrate MySQL-based services to ApsaraDB for OceanBase based on minimal or zero code modifications. In ApsaraDB for OceanBase, you can create partitioned tables and use subpartitions. This serves as an alternative to the sharding solutions of traditional databases. These features improve the efficiency of application development and data migration.

ApsaraDB for OceanBase is compatible with MySQL in the following aspects:

- APIs: Java Database Connectivity (JDBC) and Open Database Connectivity (ODBC) APIs are used in ApsaraDB for OceanBase. This allows you to access ApsaraDB for OceanBase by using MySQL drivers.
- Data objects: ApsaraDB for OceanBase supports standard SQL objects and objects that are specific to MySQL. The standard SQL objects include databases, tables, views, and auto-increment columns. ApsaraDB for OceanBase implements multitenancy in database systems.
- SQL statements:
  - ApsaraDB for OceanBase supports the standard SQL statements that you can execute to add, delete, modify, and query data.
  - ApsaraDB for OceanBase supports MySQL-specific statements that are frequently used in applications, such as REPLACE and INSERT ON DUPLICATE KEY UPDATE.
  - ApsaraDB for OceanBase supports MySQL-specific options, such as the IGNORE option in DML statements and the HINT option that specifies the indexes to be used in SELECT statements.
- System objects: System objects include system views, variables, and functions.
- Transactions: ApsaraDB for OceanBase uses the multi-version concurrency control (MVCC) protocol. This allows you to perform data reads or writes in parallel. ApsaraDB for OceanBase

supports the read committed isolation level.

# 5.14.2. Tenant security

## 5.14.2.1. Database accounts

When you log on to the ApsaraDB for OceanBase console, username and password authentication is required. After you log on to the console, you can create tenants, manage clusters, and perform operations and maintenance (O&M) tasks on clusters.

## 5.14.2.2. IP address whitelists

ApsaraDB for OceanBase allows you to use whitelists to implement security access control. You can configure a whitelist for each ApsaraDB for OceanBase tenant.

By default, a newly created ApsaraDB for OceanBase instance is accessible from all IP addresses. You can execute the following statement to modify the default access setting: ALTER TENANT tenantname SET VARIABLES ob_tcp_invited_nodes = '192.168.0.0/16,10.125.227.255/255.255.252.0'.

## 5.14.2.3. Log audit

ApsaraDB for OceanBase allows you to audit SQL logs. This provides an efficient method for you to identify issues. The gv$sql_audit file records the details of each SQL statement that is routed to ApsaraDB for OceanBase. The details include the corresponding server IP address, database name, user account, SQL statement, start time of execution, execution duration, and queuing time.

## 5.14.2.4. Software upgrades

- ApsaraDB for OceanBase releases software versions on a regular basis.
- Software upgrades are optional. Software upgrades are implemented only after you send requests to upgrade software to the specified versions.
- If the assessment results of the ApsaraDB for OceanBase team show that your version is experiencing major security risks, the team notifies you of the risks. The team also recommends that you schedule and perform software upgrades.
- The required time for upgrades depends on cluster sizes. In most cases, each upgrade requires dozens of minutes to an hour. Your services are not affected during the upgrade processes.

# 5.14.3. Notes on MySQL-related vulnerabilities in ApsaraDB for OceanBase

**Notes on MySQL-related vulnerabilities in ApsaraDB for OceanBase**

To Alibaba Cloud Apsara Stack users:

A recent security audit report revealed that ApsaraDB for OceanBase used MySQL 5.6.25. This MySQL version is at the risk of being attacked.

The version number 5.6.25 is used as a logical version number to implement the compatibility between ApsaraDB for OceanBase and MySQL. ApsaraDB for OceanBase is compatible with MySQL 5.6.25 and does not use MySQL 5.6.25. ApsaraDB for OceanBase uses a proprietary kernel and is not exposed to the security vulnerabilities that are caused by MySQL 5.6.25.

The explanation applies if MySQL server-related vulnerabilities are detected in ApsaraDB for OceanBase.

Beijing Ant Yun Financial Services Co., Ltd.
May 16, 2019

# 5.15. Data Transmission Service (DTS)

## 5.15.1. Platform security

#### 5.15.1.1. Security isolation

DTS uses independent processes and files to isolate instances and data between tenants. For example, users are not allowed to read/write OS files of instances so that users cannot access data of other users.

#### 5.15.1.2. Authentication

You can use your Alibaba Cloud account to create a DTS instance. The resources of the DTS instance are owned by the Alibaba Cloud account. The account has full access permissions on its DTS resources by default.

DTS supports RAM for Alibaba Cloud. You can assign permissions to access and manage DTS resources to RAM users. RAM enables you to assign permissions as needed and helps enterprises minimize information security risks.

#### 5.15.1.3. Transmission security

To enhance data transmission security, DTS-defined log formats are used.

In DTS, data is encrypted for secure transmission. For example, data is encrypted during incremental data synchronization between the data reading module and the data synchronization module.

DTS also supports HTTPS to effectively improve access security.

#### 5.15.1.4. Data security

When you use DTS to synchronize or subscribe to incremental data, many pieces of incremental data are stored on the DTS server. The incremental data is serialized and stored based on the storage format defined in DTS. The DTS-defined storage format helps to enhance data security.

> **Note** After data is written to the DTS server, it is stored for seven days and automatically deleted after this period of time expires.

## 5.16. Data Management (DMS)

### 5.16.1. Platform security

#### 5.16.1.1. Security isolation

VPCs are isolated from each other, and are suitable for scenarios which require high security. Therefore, VPCs are used as the infrastructure for database instances.

DMS supports access to instances in VPCs. This ensures network security and data operation convenience.

#### 5.16.1.2. Authentication

##### Authentication of Apsara Stack tenant accounts and database accounts

Before using DMS, you must log on to the Apsara Stack console or other Apsara Stack consoles with your Apsara Stack tenant account and password. If your logon session expires, or if you switch to another account, you can no longer access DMS. You must use your Apsara Stack tenant account to log on to DMS again. You can use DMS only after you are logged on with your Apsara Stack tenant account.

### Database account permission control

When you are logged on with your Apsara Stack tenant account and attempt to connect to a database through DMS, DMS will check the permissions of your account. To access resources within the database, you must be the owner of the database resources or be authorized to access the database.

## 5.16.1.3. Transmission security

**HTTPS and SSL support**

DMS supports HTTPS and SSL connections between user-side browsers and DMS servers, which prevents interception and data thefts during data transmission.

# 5.16.2. Tenant security

## 5.16.2.1. ActionTrail

### Operation audit

DMS provides an audit function, which records operations such as user logon, logoff, SQL operations, table structure changes, table data changes, import, export, and operation results. Additionally, you can use the log audit function (Monitoring System) of Apsara Infrastructure Management Framework to query user operation logs. The logs record details such the access, operations, and SQL statements of a user in a specific instance.

# 5.17. Server Load Balancer (SLB)

# 5.17.1. Platform security

## 5.17.1.1. Authentication

SLB instances created using an Apsara Stack tenant account are owned by the account. By default, the tenant account has full permissions on these resources.

You can manage SLB instances by using RAM. You can authorize a RAM user to access and manage SLB resources owned by the tenant account. SLB also supports STS, which enables temporary access to SLB resources.

# 5.17.2. Tenant security

## 5.17.2.1. HTTPS

SLB supports HTTPS load balancing to forward HTTPS requests.

SLB supports HTTPS, SSL, and TLS load balancing:

- SLB provides centralized certificate and key management for services that require certificate authentication. This eliminates the need to deploy certificate management systems on individual ECS instances.
- All decryption operations are performed on SLB, reducing the CPU usage of backend ECS instances.

SLB provides centralized certificate management that allows you to store certificates and keys. All private keys uploaded to the certificate management system are encrypted.

### 5.17.2.2. IP address whitelists

SLB masks the IP addresses of backend servers and provides only the IP address of the SLB instance for external use.

SLB also provides the whitelist function. By adding a whitelist, you can control which IP addresses can access the SLB service.

### 5.17.2.3. Log management

Server Load Balancer (SLB) provides the log management feature that allows you to view the operation and health check logs of an SLB instance.

# 5.18. Virtual Private Cloud (VPC)

## 5.18.1. Platform security

### 5.18.1.1. Security isolation

VPCs are isolated through tunneling technology. The isolation effect between VPCs is the same as that of the traditional VLANs. Broadcast domain isolation can be achieved on ECS instances and NICs. VPCs, like VLANs, are isolated at the network layer. Meanwhile, VPCs divide different security domains for access control.

Each VPC is identified by a unique tunnel ID.

A unique tunnel ID is generated when tunnel encapsulation is performed on each data packet transmitted between the ECS instances within a VPC. Then, the data packet is transmitted over the physical network.

ECS instances in different VPCs cannot communicate with each other. They have different tunnel IDs and therefore are on different routing planes.

### 5.18.1.2. Access control

VPCs support RAM. RAM allows you to grant access and management permissions on your VPC resources to RAM users.

VPCs also support STS, which issues temporary credentials to grant temporary access permissions.

# 5.18.2. Tenant security

## 5.18.2.1. Security groups

VPCs use security groups to divide network security domains and implement Layer 3 access control. The security groups act as virtual firewalls and provide stateful inspection and packet filtering features.

VPCs are isolated by default and can be connected to each other through peering connections.

# 5.19. Log Service

# 5.19.1. Platform security

## 5.19.1.1. Security isolation

Logtail supports multi-tenant isolation. Compared with mainstream open-source collection agents, Logtail has a more refined architecture. A fixed number of threads are used by Logtail to discover events, read data, parse data, and send data. The number of threads does not increase as the number of configurations increases. All configurations operate in the same execution environment. However, Log Service uses multiple technical methods to guarantee processing isolation of configurations, fair scheduling of configurations, reliability and controllability of data collection, and high cost performance of resources.

The characteristics of Logtail multi-tenant isolation are as follows:

- Logtail schedules data collection based on time slices to guarantee isolation and fairness of configuration data endpoints.
- Logtail supports multi-level feedback queues for resource usage to guarantee the isolation and fairness of processing flows and configurations with extremely low resource consumption.
- Logtail supports non-blocking mode of event processing to guarantee high reliability even if log files are rotated when configuration is blocked or data collection is stopped.
- Logtail supports different throttling mechanisms for various configurations, policies of stopping collection, dynamic configuration updates to guarantee a high level of control for data collection.

## 5.19.1.2. Authentication

As with the “authentication” content, support for ram can be removed later, using a separate topic to ensure the security of log data, all HTTP requests of the Log Service API must undergo security authentication. The security authentication is based on the Alibaba Cloud AccessKey and is done by using the symmetric encryption algorithm.

Its process is as follows:

1. The requester generates a signature string based on the API request content (including HTTP header and body).
2. The requester uses Alibaba Cloud AccessKey pair (AccessKey ID and AccessKey Secret) to sign the signature string generated in the first step. A digital signature is generated for this API request.
3. The requester sends both the API request content and digital signature to the server.

4. After receiving the request, the server repeats step 1 and step 2 to compute the expected digital signature for this request.

   > **Note** The server retrieves the AccessKey pair used by this request from the background.

5. The server compares the expected digital signature with the digital signature sent from the requester. If they are the same, the request passes security authentication. Otherwise, the request is immediately rejected.

## 5.19.1.3. Data security

The operation of collecting Log Service user data is mapped to reading and writing operations of files which are stored in the Apsara Stack data system.

Apsara Stack uses a flat network in which a linear address space is divided into slices, which are also called chunks. Each chunk is duplicated into three copies. Each copy is stored on different nodes in the cluster to ensure the reliability of data.

The triplicate technology used for the Apsara Stack data system involves three key components: master, chunk server, and client. Each write operation performed by an ECS user undergoes several processes before the client executes the operation. The following section describes how the client executes this operation:

1. The client determines the location of the chunk corresponding to the write operation.
2. The client sends a request to the master to query the storage locations (chunk servers) of all three chunk replicas.
3. The client sends write requests to the chunk servers that are returned from the master.
4. If the write operation succeeds on all three chunk replicas, the client returns a success message. Otherwise, the client returns a failure message.

The distribution strategy of the master takes all factors of the entire system into account, such as chunk server disk usage, chunk server distribution across different racks, power distribution conditions, and machine workloads. This strategy ensures that each chunk replica is distributed on different chunk servers across different racks, which effectively prevent data unavailability if a rack or chunk server fails.

In the cases of data node damage or hard disk failures on data nodes, the total number of valid replicas of some chunks become less than three. In these cases, the master replicates data between chunk servers to maintain three valid replicas of all chunks in the cluster.

All user operations including addition, modification, and deletion of data are synchronized to the three copies. This mechanism ensures the reliability and consistency of user data.

Furthermore, when you delete data, the released storage space is reclaimed by Apsara Distributed File System and is not accessible to users. Data erasure is performed before the storage space is reused. This mechanism provides the highest level of data security.

## 5.19.1.4. Encrypted data transmission

Log Service ensures your data security in transmission by using the following methods:

- You use your Alibaba Cloud AccessKey to perform Log Service authentication. To prevent data tampering during data transmission, the Logtail client obtains your Alibaba Cloud AccessKey to sign all log data packets before they are sent. The Logtail client also uses the HMAC SHA1

signature algorithm for authentication.

> **Note** The Logtail client obtains your Alibaba Cloud AccessKey over HTTPS to ensure the security of your AccessKey.

- The API layer uses the signature authentication mechanism to control access permissions and ensure the security of your data.
- Log Service applies HTTPS and SSL to the network connection between a client and a server to ensure that data is not monitored or stolen during transmission. Log Service communicates with the client over HTTPS to ensure data security.

# 5.19.2. Tenant security

## 5.19.2.1. Service monitoring

Log Service monitors machine group status and log collection status of Logtail in real time.

- Machine group status

  Log Service monitors the heartbeat status of all the servers in your machine group in real time. The server status can be **OK** or **Fail**. **Fail** indicates that the machine group is in an abnormal state and cannot collect logs.

- Log collection status

  When you use Logtail to collect logs, Log Service sends alerts through **Collection Error Diagnosis** if errors such as failed regular expression parsing, invalid file path, and insufficient shard capacity occur. Alerts contain information about the errors, such as the time of occurrence, the IP address of the server, the number of errors, and the types of errors.

# 5.20. Key Management Service (KMS)

# 5.20.1. Platform security

## 5.20.1.1. Security isolation

This topic describes the security isolation of KMS.

> **Note** No instances are deployed in KMS. Therefore, resource isolation caused by instance virtualization do not occur.

CMKs are the only resources in KMS. You can use CMKs by calling API operations but cannot directly access the CMK data. Security isolation is implemented at the network layer of the API.

## 5.20.1.2. Authentication

## 5.20.1.2.1. Identity authentication

You can create an AccessKey pair in the Apsara Stack Cloud Management (ASCM) console. An AccessKey pair consists of an AccessKey ID and an AccessKey secret. The AccessKey ID is public and identifies a user, whereas the AccessKey secret is private and is used to authenticate a user.

Before you send a request to KMS, you must generate a signature string for the request in the format specified by KMS. Then, you must encrypt the signature string by using your AccessKey secret to generate a verification code based on the HMAC algorithm. The verification code is timestamped to prevent replay attacks. After KMS receives the request, KMS finds the AccessKey secret based on your AccessKey ID and uses the AccessKey secret to decrypt the signature string and verification code. Then, KMS calculates a verification code and compares it with the decrypted verification code. If the two verification codes are the same, KMS determines that the request is valid. Otherwise, KMS rejects the request and returns HTTP status code 403.

## 5.20.1.2.2. Access control

Resource Access Management (RAM) is used for access control in KMS. You can use RAM policies to define different identity types and grant RAM users permissions on KMS.

Permissions on KMS have the following basic elements:

- Action: the Action parameter in KMS API requests. For example, you can specify this parameter to create, delete, modify, or query keys, as well as to encrypt or decrypt data by using keys. Each API operation corresponds to an action. You can grant permissions on each action to an identity.
- Resource: Keys are the only resources of KMS. Key IDs are used to identify resources.

## 5.20.1.2.3. RAM and STS support

KMS supports RAM and Security Token Service (STS) authentication.

RAM is a resource access control service provided by Alibaba Cloud. It allows you to create RAM users under your Apsara Stack tenant account and grant them permissions to access resources.

STS is an Alibaba Cloud service that provides temporary access credentials. It is used for short-term access control. You can use STS to generate a temporary access credential. You can specify the permissions and validity period of the credential. The credential expires automatically upon its expiration date.

## 5.20.1.3. Data security

KMS uses multiple mechanisms to ensure the security of your data.

Data in KMS refers to the CMKs that are created and managed in KMS. CMKs are stored in ApsaraDB for RDS servers in primary-secondary mode. Each primary or secondary server has its own redundancy and backup mechanism. In this way, ApsaraDB for RDS implements hierarchical redundancy for CMK data.

The key material of CMKs is encrypted by KMS before it is stored on disks. KMS provides a hierarchical key architecture and automatically rotates upper-layer keys. KMS allows you to use trusted platform modules (TPMs) to protect the root KMS key and ensure the privacy of your data.

## 5.20.1.4. Transmission encryption

KMS implements end-to-end encryption for data transmission.

When you send a request to KMS, you must use HTTPS to ensure the privacy and integrity of the exchanged information.

# 5.21. Apsara Stack DNS

## 5.21.1. Tenant security

### 5.21.1.1. Tenant isolation

Apsara Stack DNS isolates data by tenant. After receiving a DNS request, Apsara Stack DNS determines whether the request is authorized based on the AccessKey pair contained in the request. If it is, Apsara Stack DNS processes the data. Tunnel IDs are used to associate VPCs with zones. The backend DNS server responds to the users' DNS requests based on tunnel IDs. This helps isolate data among tenants.

### 5.21.1.2. Network security hardening

Apsara Stack DNS provides recursive resolution. To protect against potential security risks from the Internet, Apsara Stack DNS reinforces the security of domain name resolution. Specifically, only outbound traffic is allowed, and all inbound traffic from the Internet is discarded.

### 5.21.1.3. Log audit

Log audit is an essential step to ensure security. Apsara Stack DNS provides you with detailed log information. All operations are logged in real time. If a security breach occurs, you can query log entries to trace the activities of the attacker and analyze the security breach.

# 5.22. API Gateway

## 5.22.1. Platform security

### 5.22.1.1. Security isolation

API Gateway isolates resources by tenant. Resources belong only to the tenant account they are created in. Resources are isolated between different tenants.

### 5.22.1.2. Authentication

#### 5.22.1.2.1. Authentication

You can generate an AccessKey pair in the Apsara Stack console. An AccessKey pair contains an AccessKey ID and an AccessKey Secret. The AccessKey ID is public and uniquely identifies a user. The AccessKey Secret is private and is used to authenticate the user identity.

Before you send a request to API Gateway, you must generate a signature string for the request in the format specified by API Gateway. Then, you need to encrypt the signature string by using the AccessKey Secret (based on the HMAC algorithm) to generate a verification code. The verification code is timestamped to prevent replay attacks. After receiving the request, API Gateway finds the AccessKey Secret corresponding to the AccessKey ID, and extracts the signature string and verification code in the same way. If the calculated verification code is the same as the one received, the request is valid. Otherwise, API Gateway rejects the request and returns an HTTP 403 error.

## 5.22.1.2.2. API access control

API Gateway users consist of API providers and third-party users (API callers). The provider of an API can request a third-party user to provide an AppId. After the API provider authorizes the application specified by AppId to call the API, the third-party user can immediately initiate an access request to the API by using the AppKey and AppSecret of the application.

An access request must carry the signature of the user who accesses the API. A signature-based access request is a request that contains signature information in the header as stipulated in the API Gateway documentation.

## 5.22.1.2.3. RAM and STS support

API Gateway allows you to manage APIs through RAM and STS.

Resource Access Management (RAM) is a resource access control service provided by Alibaba Cloud. You can use an Apsara Stack tenant account to create RAM user accounts and grant them permissions to access the resources that belong to the tenant account.

Security Token Service (STS) is a temporary access credential service provided by Alibaba Cloud. It provides temporary access control. You can use STS to generate a temporary access credential. You can determine the permissions and validity period of the credential. The access credential expires automatically upon its expiration date.

## 5.22.1.3. Data security

API Gateway uses signature authentication to ensure the consistency and integrity of user data when an API is called. In addition, API Gateway provides the data cleansing function. During data cleansing, invalid parameters are cleaned to ensure the security of requests.

API Gateway requires the caller to add user identity information to the API request, and add an encrypted signature to the data for transmission. After receiving an API request, API Gateway verifies the user identity and checks data for integrity and consistency.

In addition, API Gateway can clean invalid parameters that are not preset by users. This ensures that only valid and secure API requests are approved.

## 5.22.1.4. Transmission encryption

API Gateway supports HTTPS to ensure the security of data during transmission.

# 5.22.2. Tenant security

#### 5.22.2.1. Log audit

Log Service is a log management service provided by Alibaba Cloud. API Gateway can be used together with Log Service to provide you with access, monitoring, and audit information query and display functions. API Gateway records API requests in real time and synchronizes the logs to Log Service on a regular basis. Request logs contain information such as the request time, source IP address, requested object, returned code, and processing duration.

#### 5.22.2.2. IP address-based access control

API Gateway allows you to set a blacklist and a whitelist based on the ClientIP of a caller.

An IP address-based access control policy takes effect immediately after you bind it to an API. This can prevent API requests from unauthorized IP addresses.

- Blacklist: prohibits API requests from the specified IP addresses.
- Whitelist: permits the API requests from only the specified IP addresses.

## 5.23. Enterprise Distributed Application Service (EDAS)

### 5.23.1. Platform security

The platform-side security design mainly includes authentication and transmission encryption.

#### 5.23.1.1. Authentication

Authentication Enterprise Distributed Application Service (EDAS) mainly includes permission control, access control, and API authentication.

##### ACL

- EDAS Agent

  EDAS Agent implements authentication and authorization at two layers: it restricts authentication from the consumer to provider and implements the principle of least privilege by authenticating the AccessKey ID and AccessKey secret of delivered commands.

- DAuth
  - Security credentials: EDAS DAuth can generate the App_KEY and App_SECRET for EDAS access requests to implement authentication control.
  - Authentication settings: Granular policy settings for EDAS authorization are supported, including authorization enabling, signature verification, log switch, custom logs, log cache, and log detection.

- Diamond Server

  Diamond Server uses the App_KEY and App_SECRET generated by DAuth to authenticate access requests to HTTP and HTTPS APIs.

- Config Server

- Throttling: HTTP requests sent to Config Server can be throttled to prevent Config Server from being unstable under excessive workloads.
- Authentication: Config Server uses the App_KEY and App_SECRET generated by DAuth to authenticate access requests to HTTP and HTTPS APIs.

### Access control

EDAS RAM authorization can retrieve your on-demand AccessKey ID and AccessKey secret through STS. You can create ECS instances by using the on-demand AccessKey ID and AccessKey secret to access your ECS API. On-demand AccessKey ID and AccessKey secret have a validity period and must be regenerated after expiration. The ECS API permissions are restricted. You can call API operations only with the specified API operation permissions. The restrictions avoid excessive permissions caused by the use of ECS RAM FullAccess.

### API authentication

API access permissions are authenticated using the AccessKey ID and AccessKey secret of the RAM user. The EDAS API uses the custom key-based Hash Message Authentication Code (HMAC) HTTPS solution for authentication. To authenticate a request, you must combine certain elements in the request to form a string. Then, use the EDAS key to calculate the HMAC of the string. Generally, this process is called request signing. The output HMAC algorithm is called signature because the algorithm simulates the security attributes of real signatures. Finally, you can use the EDAS API syntax to add the signature as a request parameter.

After the system receives the authenticated request, it extracts the EDAS key and uses the key to calculate the signature of the received message by using the same method. Then, the system compares the calculated signature with the signature of the requester. If both signatures match, the system considers that the requester has the access permission on the EDAS key and issues the key to the requester as the issuance authority of the delegator. If the signatures do not match, the request is discarded, and the system returns an error message.

## 5.23.1.2. Transmission encryption

Enterprise Distributed Application Service (EDAS) uses the TLS certificate deployment solution to realize full link encryption. Considering the validity period of the certificate, EDAS supports certificate update.

### EDAS management security

EDAS management involves the following roles:

- EDAS Console: allows you to create clusters and deploy packages in the EDAS console.
- EDAS Agent: is an EDAS client deployed on your Elastic Compute Service (ECS) instance and supports various EDAS operations. EDAS Agent consists of two components: StarAgent and EDAS Agent.
- EDAS Server: receives commands from EDAS Console and delivers them to EDAS Agent.

The security process of EDAS management is as follows:

1. When EDAS Console receives your deployment command (for example, deploying a WAR package), EDAS Console performs API authentication and uses the AccessKey ID and AccessKey secret of a Resource Access Management (RAM) user to connect to the EDAS Server API.
2. EDAS Server sends the encrypted commands to EDAS Agent through the encrypted channel.

### EDAS RPC security

The EDAS remote procedure call (RPC) process involves the following roles:

- EDAS DAuth: generates the AccessKey ID and AccessKey secret of an EDAS user and performs actions such as authentication.
- EDAS Config Server: provides information including service registration, IP addresses, and called APIs.
- EDAS Dubbo and High-speed Service Framework (HSF): are distributed RPC products.
- EDAS Pandora: is a lightweight container isolation service for class isolation and loading.

The EDAS RPC process is as follows:

1. Both the provider and consumer start the Pandora process. Pandora provides call services for HSF and Dubbo.
2. On the ECS instance to which it belongs, the provider registers a service with EDAS Config Server. During service registration with EDAS Config Server, TLS is used to protect link security, and the EDAS AccessKey ID and AccessKey secret are used for authentication. Only specified users can publish called APIs to register with EDAS Config Server.
3. The consumer pulls related information from EDAS Config Server, including the service name and service IP address provided by the provider.
4. The consumer directly calls the Pandora process of the provider to access the service, uses the EDAS AccessKey ID and AccessKey secret for authentication to call the service.

### EDAS Docker security

The EDAS Docker call process is as follows:

1. Create a Docker cluster in EDAS by calling the Docker API. EDAS uses the RAM Security Token Service (STS) on-demand token for authentication.
2. Create cluster certificates. One certificate is generated for each cluster.
3. Create resources, including ECS and Server Load Balancer (SLB) instances, and use the RAM STS on-demand token to call the APIs for ECS or SLB instances. Authorize the required APIs.
4. Configure nodes.
    i. Generate node certificates. The system generates a root certificate for a cluster and uses cloud-init to directly generate node certificates based on the root certificate.
    ii. Use cloud-init to install Docker.
    iii. Install system services such as EDAS Agent by using the EDAS installation scripts. Then you can apply the EDAS management process to manage these services.

## 5.23.2. Tenant security

API auditing is the main security function on the tenant side.

ActionTrail records the operations performed on resources under users' cloud accounts, supports operation record query, and stores record files in user-specified OSS buckets. The operation records stored by ActionTrail can be used for security analysis, resource change tracking, and compliance audit.

ActionTrail collects API call records of cloud services, including the records of API calls initiated in the console). It standardizes the operation records and saves them to user-specified OSS buckets as files. You can manage the record files by using all the management functions provided by OSS, such as authorization, lifecycle management, and archiving management.

# 5.24. MaxCompute

## 5.24.1. Platform security

### 5.24.1.1. Security isolation

MaxCompute is designed to handle security issues in multi-tenant scenarios. It uses the Alibaba Cloud account authentication system, which authenticates users based on symmetric AccessKey pairs. It also verifies the signature information in each HTTP request. MaxCompute stores user data separately in Apsara Distributed File System to achieve data isolation between users. This allows MaxCompute to meet the requirements for multi-user collaboration, data sharing, data confidentiality, and data security, implementing true resource isolation between tenants.

MaxCompute executes all computing tasks in individual sandboxes. The sandboxes are multi-layered, starting at the kernel level. System sandboxes are combined with an authentication management mechanism to ensure data security and prevent server failures caused by human error or malicious operations.

Sandbox protection

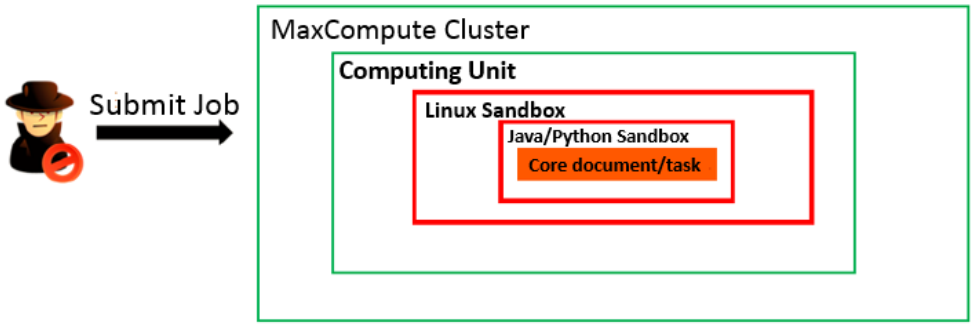


### Network isolation

MaxCompute is a big data platform provided by Alibaba Cloud to process massive amounts of data. Because MaxCompute processes and stores large amounts of user data, it must comply with security isolation standards to ensure the security of user data. MaxCompute is designed with the support for Virtual Private Cloud (VPC) and configured with limits for MaxCompute, that is, restrictions on VPC-capable MaxCompute.

At present, MaxCompute supports VPCs in the following ways:

- Classic networks, VPC networks, and the Internet are isolated from each other. Each network can only access endpoints and virtual IP addresses (VIPs) within themselves.
- Projects that do not have VPC IDs or IP address whitelist configured do not have any access restrictions. These projects can be accessed by their domain names over all three networks.
- Projects that have VPC IDs configured are only accessible from the VPC with the same ID.

- Projects that have IP address whitelists configured are accessible from the machines whose IP addresses are in the whitelist.
- If a proxy is used to send an access request, the request will be granted or denied based on the IP address or VPC ID of the last-hop proxy server.

**Elasticsearch on MaxCompute** is an enterprise-level massive data retrieval system developed by Alibaba Cloud. It also needs to meet security isolation requirements. Therefore, the MaxCompute team added VPC-capable **Elasticsearch on MaxCompute** based on the original VPC-capable MaxCompute, and configured limits for **Elasticsearch on MaxCompute**, that is, restrictions on VPC-capable Elasticsearch.

At present, **Elasticsearch on MaxCompute** supports VPCs in the following ways:

- Classic networks, VPC, and the Internet are isolated from each other. Users can access only the endpoints and VIPs on their networks.
- Projects without a VPC ID whitelist or IP address whitelist are accessible to users from valid domains over the three types of networks.
- To start an Elasticsearch service instance, MaxCompute must use the same VPC list as the instance. They share a VPC whitelist and have the same VPC restrictions.
- Starting an Elasticsearch instance occupies all resources by default. You must scale up the MaxCompute instance or scale down the Elasticsearch instance if you start more Elasticsearch instances.

Specific usage scenario: An Elasticsearch instance is started for each project by default when a MaxCompute is deployed in the Apsara Stack environment. You can start your Elasticsearch instance in your project, apply for a domain name and VIP after the instance is started, and verify VPC settings in the Elasticsearch frontend.

# 5.24.1.2. Authentication and authorization

## Identity authentication

You can create an AccessKey pair in the Apsara Stack Cloud Management (ASCM) console. An AccessKey pair consists of an AccessKey ID and an AccessKey secret. The AccessKey ID identifies a user, while the AccessKey secret is used to authenticate the user identity. The AccessKey secret must be kept strictly confidential.

Before you send a request to MaxCompute, you must generate a signature for the request. To do this, you must generate a signature string in the format specified by MaxCompute and encrypt the signature string by using your AccessKey secret. After MaxCompute receives the request, it identifies the AccessKey secret that corresponds to the AccessKey ID. Then, it extracts the signature string and verification code in the same way. If the extracted verification code is the same as the one received, the request is valid. Otherwise, MaxCompute rejects the request and returns an HTTP 403 error.

## Access control

You can use an Apsara Stack tenant account or a RAM user to access MaxCompute resources. You can use an Apsara Stack tenant account to create different RAM users for use in different scenarios. MaxCompute applies access control policies when you access resources by using an Apsara Stack tenant account or a RAM user.

- If you access a resource by using an Apsara Stack tenant account, MaxCompute verifies whether the account is the resource owner. The resource owner is the only account that can

access the resource.

- If you access a resource by using a RAM user, MaxCompute checks whether the Apsara Stack tenant account of the RAM user is the resource owner and whether the RAM user is granted permissions on that resource.

> **Note** The preceding descriptions apply only to an unauthorized Apsara Stack tenant account and RAM users. If the Apsara Stack tenant account and RAM users are granted permissions on a resource, they can access the resource that otherwise can only be accessed by the resource owner.

MaxCompute supports the following two authorization mechanisms for RAM users:

- **ACL-based authorization**: an object-based authorization mechanism. An access control list (ACL) contains the permission data of an object. It is a resource of the object. ACL-based authorization can be performed only on objects that exist. If an object is deleted, the ACL of the object is also deleted. ACL-based authorization is similar to the authorization mechanism that is implemented by using the GRANT and REVOKE statements defined in SQL-92. You must execute statements to grant or revoke permissions on an object.

  > **Note** In the current version, ACL-based authorization supports independent management of field-level permissions. ACL fields and tables are independent authorization objects and contain complete authorization information, including permissions, expiration time, and conditions. You can authorize separate permissions on ACL fields and tables and specify expiration time and conditions. You can view authorization information and independently revoke permissions on them. ACL field-level permissions are in the OR relationship. These permissions are independent of table-level permissions in the existing authentication process. If table-level permissions do not exist, the system checks whether field-level permissions are granted.

- **Policy-based authorization**: a policy-based authorization mechanism. An access control policy contains the permission data of both an object and a subject. It is a resource of the subject. Policy-based authorization can be performed on objects and subjects that do not exist or are uncertain. The system does not automatically change or delete policies related to an object when the object is deleted. MaxCompute uses a custom language to specify access control policies for objects.

MaxCompute also supports various other access control mechanisms.

## Hierarchical authorization of management permissions

MaxCompute supports hierarchical authorization of management permissions. Permission objects and operations are used to define permissions on management operations. You can use existing access control mechanisms, such as policy-based or ACL-based authorization, to control permissions on management operations.

Permission objects include policies, ACLs, project configurations, and projects. Permission operations are the actions that are performed on permission objects, such as the Put action in Put Policy and the Create action on projects.

Fine-grained management permissions support more complex access control scenarios, such as allowing only specific clients to perform management operations.

## Column-level access control

Label-based security (LabelSecurity) is a mandatory access control (MAC) policy at the project level. It allows project administrators to control user access to sensitive data at the column level.

LabelSecurity classifies both data and users who want to access the data into different levels. The data is classified into the following levels based on its sensitivity:

- Level 0: unclassified data
- Level 1: confidential data
- Level 2: sensitive data
- Level 3: highly sensitive data

Project owners must define their own standards to determine which level their data is classified as and what level of access each level is permitted. The default sensitivity level of data is 0. The default access level of all users is also 0.

LabelSecurity provides data sensitivity classification at the column level. Administrators can label columns in a table with sensitivity levels. A table can have columns with individual sensitivity levels. Administrators can also label views with sensitivity levels. By default, the sensitivity level of a new view is 0. The sensitivity levels of a view and its base table are independent of each other.

LabelSecurity applies the following default security policies based on the levels of data and users:

- No-ReadUp: Users are not allowed to read data that has a higher sensitivity level than their own, unless they are explicitly authorized.
- Trusted-User: Users are allowed to write data into columns regardless of data sensitivity levels. The default sensitivity level of a new column is 0.

> **Note**
>
> - In some traditional MAC systems, other complex security policies are implemented to prohibit unauthorized data operations in a project. For example, the No-WriteDown policy prohibits users from writing data with a sensitivity level that is not higher than the user level. No-WriteDown is not included as a default security policy in MaxCompute to reduce management costs for data sensitivity levels. Project administrators can set the SetObjectCreatorHasGrantPermission parameter to false to implement a policy similar to No-WriteDown.
> - If you want to prevent data flows between projects, you can enable ProjectProtection. After the settings take effect, users are only able to access data within their own projects and data cannot flow to other projects.

LabelSecurity is disabled by default. Project owners can enable it as needed. After LabelSecurity is enabled, the preceding default security policies take effect. After these policies take effect, users must have the Select permission and required access level to read sensitive data in tables.

Run the following command to enable or disable LabelSecurity:

```
Set LabelSecurity=true|false;
  -- This command is used to enable or disable the LabelSecurity mechanism. The default value is false
.
  --Only project owners can run this command. Other operations can be performed by project administr
ators.
```

## Fine-grained column-level authorization and authentication based on permission model 2.0

To achieve fine-grained permission management and enhance data security, MaxCompute provides column-level authorization and authentication based on permission model 2.0.

MaxCompute provides centralized management and query interfaces for permissions based on permission model 2.0 to implement column-level authorization and authentication.

- MaxCompute and DataWorks allow for fine-grained tenant authorization and access control of shared package resources across tenants.
- MaxCompute authorizes users to access tables and track the queries on those tables.

The following section describes how to use column-based table operations to configure fine-grained column-level authorization based on permission model 2.0.

> **Note** Fine-grained authorization cannot be implemented by using access control policies because they cannot ensure security. The following examples implement fine-grained authorization by using ACLs.

**Syntax for ACL-based fine-grained authorization**

- Grant or revoke permissions on columns in a table within a project

  ```
  grant/revoke <privileges> on table <name>(<column_list>) to|from USER/ROLE <user/role name>;
  ```

  Description:

  i. If you grant permissions on a table, permissions are granted on all columns in the table, including:
     a. Added columns
     b. Renamed columns
  ii. If you grant or revoke permissions on columns, permissions on the columns can be differentiated from other permissions or merged with similar permissions.
     a. If you grant the Select permission on col1 and col2 and grant the Describe permission on col2 and col3, both authorization statements are valid.
     b. If you grant the Select permission on col1 and col2 and then grant the Select permission on col3 and col4, the Select permission takes effect on col1, col2, col3, and col4.
  iii. Only a project owner and other authorized users can grant permissions.

- Add columns to a package across projects

```
add table <name>(<column list>) to package pkgdel1 with privileges <privilege list>;
```

> **Note** If you add columns to a package multiple times, permissions on the columns can be differentiated from other permissions or merged with similar permissions.

Description:

i. The syntax does not change when you install the package.

ii. The syntax does not change when you install or uninstall the package. However, permissions on columns added to the package must be taken into account.

iii. Only a project owner and other authorized users can grant permissions.

- Grant or revoke permissions on columns across projects

```
grant/revoke <privileges> on table <name>(<column_list>) to|from USER|ROLE <user/role name>
PRIVILEGEPROPERTIES("refobject"="true", "refproject"="<project_name>", "package"="<package nam
e>");
```

Description:

i. The rules used to grant or revoke permissions on a table within a single project take effect for cross-project authorization only if a table is added to more than one package.

ii. You can grant or revoke permissions only on columns that are added to the packages. Grant appropriate permissions on columns in the tables that are added to more than one package.

iii. Only a project owner and other authorized users can grant permissions.

**Authentication policy:** Policy-based authentication uses the same logic as ACL-based authentication. Make sure that you consider how these policies interact with column-level permissions.

**Permission query**

- Query permissions within a project

```
show grants for <user|role name>; #The syntax does not change. However, the result is displayed a
t the column level.
show grants for table <name>(columns);
show grants on table <name>(columns) for user|role <name>;
```

> **Note** If a column is specified, only permissions on that column are displayed.

- Query package permissions

```
describe package <pkg name>;
describe package <pkg name> PRIVILEGEPROPERTIES ("allowedonly"="true");
describe package <pkg name> PRIVILEGEPROPERTIES ("contentonly"="true");
```

> **Note** The preceding describe package commands return results at the column level.

- Query permissions across projects

```
show grants for <user|role name> PRIVILEGEPROPERTIES ("refobject"="true", "refproject"="<project>
"); #The syntax does not change. However, the result is displayed at the column level.
show grants for table <name>(columns) PRIVILEGEPROPERTIES("refobject"="true", "refproject"="<pr
oject>");
show grants on table <name>(columns) for user|role <name> PRIVILEGEPROPERTIES ("refobject"="tr
ue", "refproject"="<project>");
```

**Audit policy**: Relevant information is included in audit logs.

### Access control

MaxCompute supports RAM authorization.

**Resource Access Management (RAM)** is a resource access control service provided by Alibaba Cloud. You can use your Apsara Stack tenant account to create RAM users and grant them permissions to access specific resources owned by the account.

## 5.24.1.3. Data security

Apsara Stack uses a flat design in which a linear address space is divided into slices, also called chunks. Each chunk is replicated into three copies. Each copy is stored on a different node in the cluster, which ensures data reliability.

Triplicate technology used for Apsara Stack data storage involves three key components: master, chunk server, and client. Write operations on MaxCompute undergo several processes before the client executes the operation. The processes are described as follows:

1. The client determines the location of a chunk corresponding to the write operation.
2. The client sends a request to the master to query the storage locations (chunk servers) of the three chunk replicas.
3. The client sends write requests to the chunk servers that are returned from the master.
4. If the write operation succeeds on all three chunk replicas, the client returns a success message. Otherwise, the client returns a failure message.

The distribution strategy of the master takes factors of the entire system into account, such as chunk server disk usage, chunk server distribution across racks, power distribution conditions, and machine workloads. This strategy ensures that each chunk replica is distributed to different chunk servers on different racks. This effectively prevents data unavailability caused by the failure of a chunk server or rack.

When a data node is damaged or disk faults occur on a data node, the total number of valid replicas of some chunks in a cluster becomes less than three. In these cases, the master replicates data between chunk servers to ensure that the number of valid replicas of all chunks in the cluster is three.

In short, all data operations in MaxCompute (add, modify, and delete operations) are synchronized to three replicas. This approach ensures user data reliability and consistency.

Furthermore, when data is deleted, the released storage space is reclaimed by Apsara Distributed File System. During this period, the space is not accessible from any users. Data erasure is performed on the space before it is released for further usage. This mechanism provides a high level of protection for user data.

# 5.24.1.4. KMS-based storage encryption

## Background information

As MaxCompute is deployed more and more around the world, its security requirements increase and the need to protect sensitive user data such as financial information becomes more and more crucial. MaxCompute must also implement data-at-rest encryption to comply with security regulations.

Therefore, MaxCompute provides storage encryption. It uses Apsara Distributed File System to encrypt, store, and decrypt user data, and uses Key Management Service (KMS) to ensure the security of user data and keys.

The MaxCompute storage encryption feature provides an added layer of security and minimizes damage caused by data loss. Even if encrypted data is lost or stolen, its content cannot be extracted.

> Note
>
> - The MaxCompute storage encryption feature enables transparent encryption and decryption of data by using customer master keys (CMKs) to simplify user operations.
> - MaxCompute allows you to configure storage encryption settings by project. After encryption settings are configured for a project, all subsequent data writes to the project are encrypted to reduce security risks.
> - The MaxCompute storage encryption feature is backward compatible. It allows unencrypted projects to be encrypted and allows both encrypted and unencrypted data to coexist in a project.

## Description

This section describes related features.

- MaxCompute uses projects as its basic operational unit and stores the table data of projects in encrypted form. Only full table encryption is supported. Resources and volumes cannot be encrypted.
- The types of tasks that support storage encryption are SQL Task 2.0, including the service mode, MergeTask, and Tunnel. After the storage encryption feature is enabled for a project, table data written by using these types of tasks is stored in encrypted form.
- AES-CTR, AES-256, and RC4 encryption algorithms are supported.
- MaxCompute is connected to KMS to ensure the security of keys. You must activate KMS to generate and manage keys for encryption and decryption. After you submit a request to enable storage encryption, MaxCompute automatically connects to KMS to generate the keys required for encryption.

> Note Projects that belong to the same project owner share the same key.

- You can read encrypted and unencrypted data without the need to change task types. Encrypted and unencrypted data can coexist in a project.

## Procedure to encrypt data

1. Send a request to activate KMS.
2. Send a request to enable the storage encryption feature of MaxCompute.

   > **Note** When you enable the storage encryption feature, the system asks KMS to create a CMK. The CMK is used to protect the data key used for encryption.

3. After KMS is activated and storage encryption is enabled, submit jobs in MaxCompute for processing. After the jobs are processed, MaxCompute uses Apsara Distributed File System to encrypt the stored data.
4. Apsara Distributed File System provides KMS with the created CMK to obtain the data key used for encryption.
5. The data key obtained from KMS consists of a Data Key (DK) and an Enveloped Data Key (EDK). DK is the plaintext key used to encrypt data, and EDK is the ciphertext key generated by using envelope encryption on the DK. After Apsara Distributed File System encrypts data by using the DK, it stores the encrypted data and EDK to complete the data encryption process.

The following figure shows how to encrypt data.

Procedure to encrypt data

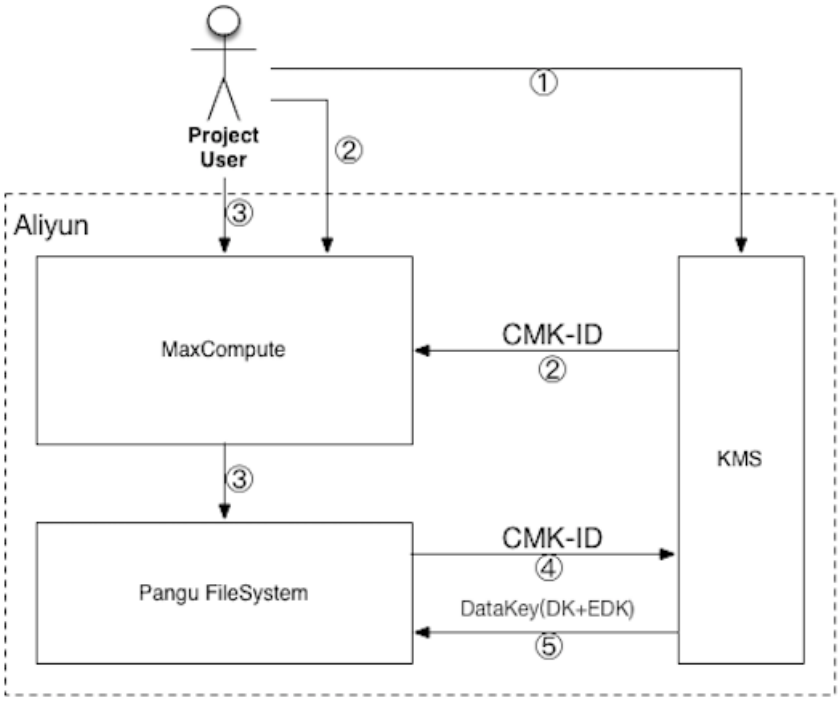


## Procedure to process encrypted data

If you use MaxCompute to process encrypted data, the system automatically decrypts the data. You do not need to perform any other operations to decrypt data. Procedure:

1. Submit a MaxCompute task to process encrypted data.
2. Apsara Distributed File System reads the EDK of the encrypted data and sends it to KMS to obtain the DK.

   > Note To ensure security, you cannot directly use the EDK to decrypt data.

3. Apsara Distributed File System decrypts data based on the received DK.
4. MaxCompute processes the decrypted data and returns the result.

The following figure shows how to process encrypted data.

Procedure to process encrypted data

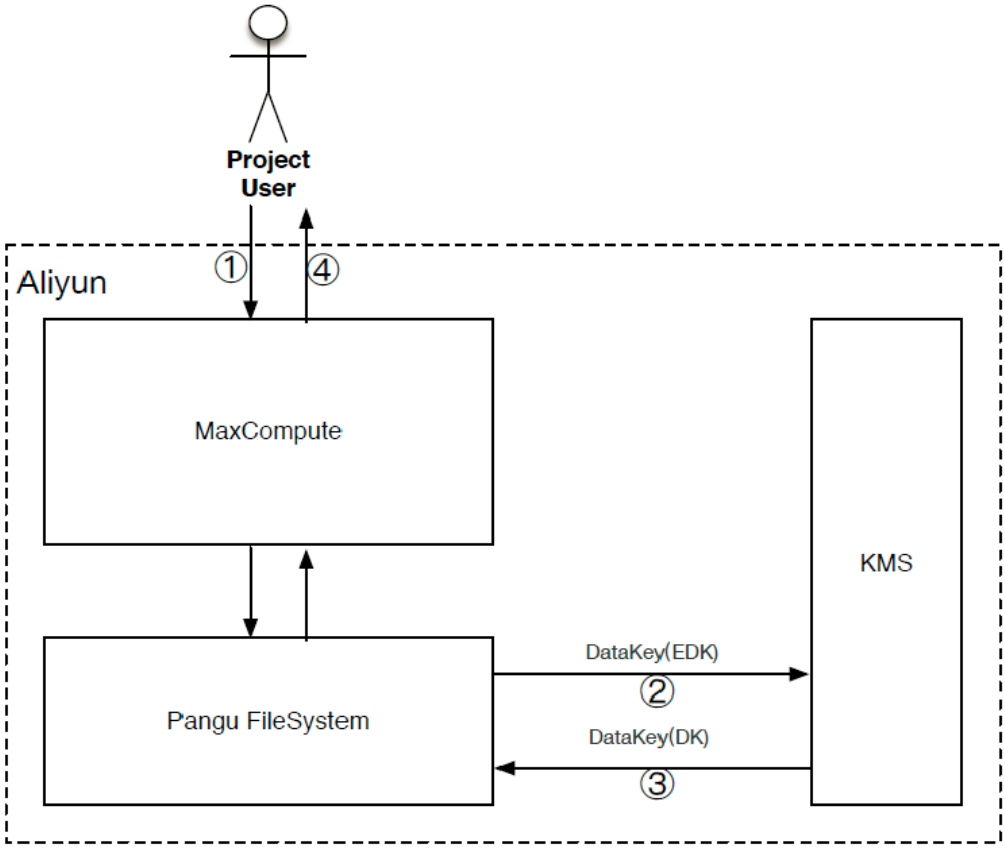


## Use of CMKs for storage encryption

MaxCompute supports CMKs to suit your business and security requirements in different scenarios. When you create a project, you can specify a specific CMK to encrypt data.

Procedure to use a CMK to encrypt and decrypt data:

1. When you create a project, send a request to enable the storage encryption feature.
2. Specify the CMK ID used by the project.

   You can use MaxCompute to create a CMK or use the CMK that you created or uploaded in KMS.

   If you use a CMK that you created or uploaded in KMS, you must authorize MaxCompute to use the CMK and create and use the related DK to encrypt and decrypt data.

3. Select an encryption algorithm.
4. Configure the other settings required to create the project.
5. After the project is created, the storage encryption feature takes effect. Data written to MaxCompute by using SQL and Tunnel tasks is stored in encrypted form.

### Precautions

When you enable the storage encryption feature for a project, take note of the following rules:

- KMS must be activated by using your Apsara Stack tenant account. If KMS is not activated, a message appears, indicating that the storage encryption feature cannot be enabled.
- When you submit a request to enable the storage encryption feature for a project, you must specify an encryption algorithm. The default encryption algorithm is AES-CTR.
- A MaxCompute production engineer must enable the storage encryption feature for a project and specify the encryption algorithm by using AdminConsole. Then, AdminConsole accesses KMS to generate the CMK required for encryption.
- After the project configuration takes effect, table data that was generated or imported by using tasks for which the storage encryption feature is enabled is stored in encrypted form.
- User data is not automatically decrypted unless the storage encryption feature is enabled. If existing data must be decrypted, it must be overwritten in encrypted form.

### Special notes

- After the storage encryption feature takes effect, data encryption and decryption operations are automatically performed. No additional operations are required to use the data.
- In addition to the supported task types mentioned in the "Description" section, other task types such as OpenMR can be used but are unavailable for data encryption.
- After the storage encryption feature takes effect for a project, you can query your keys in KMS. However, you cannot change the keys or encryption algorithm.
- The storage encryption feature can be disabled. After the feature is disabled, new data is no longer stored in encrypted form. Encrypted data remains encrypted until it is overwritten. The existing encrypted data can still be read by using KMS.
- MaxCompute must use a temporary STS token of a RAM role to access your keys. Therefore, you must evaluate the load on RAM.

### Remarks

You can configure storage encryption settings in the Apsara Stack Cloud Management (ASCM) console and Apsara Stack Operations (ASO) console. You can enable the storage encryption feature during project creation in the ASCM console and encrypt data during project management in the ASO console.

## 5.24.1.5. Transmission encryption

MaxCompute provides RESTful APIs for transmission and uses HTTPS to ensure security.

# 5.24.2. Tenant security

## 5.24.2.1. Cross-project resource sharing

Assume that you are the owner or administrator (admin role) of a project, and you receive a resource access application from a user to use resources in your project. If the applicant belongs to your project, we recommend that you grant permissions by using the user authorization management function. If the applicant does not belong to your project, you can use the package-based resource sharing function to grant cross-project access to the applicant.

A package is a mechanism to share data and resources across projects. It is used to solve cross-project user authorization problems.

The administrator of project A can create a package that includes all objects required by project B. Then, the administrator of project A can grant project B the permissions to install the package. After the package is installed in project B, the administrator of project B can determine whether to grant the permissions on the package to the users of project B.

The following section describes how to use packages:

- A package creator can run the following commands to perform package-related operations:

```
create package <pkgname>
-- Create a package
```

> **Notice**
>
> - Only project owners have the permissions to perform this operation.
> - A package name cannot exceed 128 characters.

```
add project_object to package package_name [with privileges privileges]
remove project_object from package package_name
project_object ::= table table_name |
instance inst_name |
function func_name |
resource res_name
privileges ::= action_item1, action_item2, ...
-- Add resources to a package
```

> **Note**
>
> - A project is not a valid object. A project cannot be added to a package.
> - When an object is added to a package, the permissions on the object are also added to the package. If you do not include [with privileges privileges] to specify permissions when you add an object to the package, the read-only permissions are granted on the object by default. Read-only permissions allow you to perform read, describe, and select operations on resources. Objects and their permissions are processed as a whole and cannot be changed after being added to a package. If you want to change the permissions on an object, you must delete the object from the package and add it again.

```
allow project <prjname> to install package <pkgname> [using label <number>]
-- Grant access permissions to other projects
```

```
disallow project <prjname> to install package <pkgname>
-- Revoke access permissions for other projects
```

```
delete package <pkgname>
-- Delete a package
```

```
show packages
-- View the list of packages
```

```
describe package <pkgname>
-- View details of a package
```

- A package user can run the following commands to perform package-related operations:

```
install package <pkgname>
-- Install a package
```

> **Note**
> - Only project owners have the permissions to perform this operation.
> - To install a package, you need to enter pkgName in the following format: *<projectName>.<packageName>*.

```
uninstall package <pkgname>
-- Uninstall a package
```

> **Note** To uninstall a package, you need to enter pkgName in the following format: *<projectName>.<packageName>*.

```
show packages
-- View the list of packages that have been created and installed
```

```
describe package <pkgname>
-- View details of a package
```

An installed package is an independent object in MaxCompute. To access resources in a package (resources shared by other projects), you must have the read permissions on the package. If you do not have the read permissions on a package, you can submit an application for the permissions to the project owner or administrator. The project owner or administrator can grant permissions to you by using the ACL or policy authorization method.

**Example:** You can use ACL authorization to allow a user (odps_test@aliyun.com) to access resources in a package.

```
use prj2;
install package prj1.testpkg;
grant read on package prj1.testpackage to user
aliyun$odps_test@aliyun.com;
```

You can use policy authorization to allow any user in project prj2 to access resources in the package.

```
use prj2;
install package prj1.testpkg;
put policy /tmp/policy.txt;
```

**Note** /tmp/policy.txt content is as follows:

```
{
"Version": "1",
"Statement":
[{
"Effect":"Allow",
"Principal":"*",
"Action":"odps:Read",
"Resource":"acs:odps:*:projects/prj2/packages/prj1.testpkg"
}]
}
```

# 5.24.2.2. Data protection mechanism (Project Protection)

Users authorized to access multiple projects can perform cross-project data access operations to transfer project data. If a project contains highly sensitive data that cannot be shared with other projects, the administrator can enable Project Protection to only allow inbound data.

You can enable Project Protection as follows:

```
set projectProtection=true
-- Set the Project Protection rule to allow only inbound data flow.
```

**Note** The Project Protection parameter is set to false by default. You need to manually enable Project Protection by setting this parameter to true.

## Outbound data flow after Project Protection is enabled

After you enable Project Protection for a project, you may encounter this situation: A user requests to export data from a table in the project. The table is evaluated, and verified that it does not contain any sensitive data. MaxCompute provides two methods to export data after Project Protection is enabled.

**Configure an exception policy**

A project owner can configure an exception policy when enabling Project Protection. The command is as follows:

```
SET ProjectProtection=true WITH EXCEPTION <policyFile>
```

> **Note** An exception policy is different from authorization despite the fact that both operations have the same syntax. This exception policy implements an exception in the Project Protection mechanism. Any access requests that meet the description of the exception policy can ignore Project Protection rules.

Exception policy example:

```
{
   "Version": "1",
   "Statement":
   [{
      "Effect":"Allow",
      "Principal":"ALIYUN$Alice@aliyun.com",
      "Action":["odps:Select"],
      "Resource":"acs:odps:*:projects/alipay/tables/table_test",
      "Condition":{
         "StringEquals": {
            "odps:TaskType":["DT", "SQL"]
         }
      }
   }]
   }
-- Allow the user Alice@aliyun.com to perform a select operation on the alipay.table_test table throug
h an SQL task and transfer data from the alipay project.
```

Note

- An exception policy is different from typical authorization. Assume that the user Alice in the preceding example does not have the permissions to perform select operations on the alipay.table_test table. Although the exception policy is configured, Alice cannot export data from the table.
- Project Protection is a method to control the flow of data, not a method to control the access to data. Data flow control is effective only when users can access the data in question.

**Set Trusted Projects**

Assume that there is a protected source project and a destination project. The destination project is set as the Trusted Project of the source project. Data flow from the source project to the destination project does not break the Project Protection rules. If any two projects are mutually set as Trusted Projects, these projects then form a Trusted Project Group. Data flow within this group is allowed, but data flow out of the group is prohibited.

You can run the following commands to manage Trusted Projects:

```
list trustedprojects;
-- View all Trusted Projects in the current project.
add trustedproject <projectname>;
-- Add a Trusted Project to the current project.
remove trustedproject <projectname>;
-- Remove a Trusted Project from the current project.
```

## Resource sharing and data protection

In MaxCompute, package-based resource sharing and project protection are mutually independent mechanisms that take effect at the same time, but their functions are mutually restrictive.

MaxCompute gives priority to resource sharing over data protection. If a data object is made accessible from users from other projects through resource sharing, the object is not subject to Project Protection rules.

## Prevent data flow out of projects

To prevent outbound data flow from projects, you must enable Project Protection and verify the following settings:

- Make sure that no Trusted Projects are added. If any Trusted Projects are added, you must assess the potential risks.
- Make sure that no exception policies are configured. If any exception policies are configured, you must assess the potential risks, especially data leakage due to TOC2TOU.
- Make sure that no data sharing packages are used. If any data sharing packages are used, you must ensure that no sensitive data exists in the package.

# 5.24.2.3. Log audit

MaxCompute allows you to audit various kinds of log data generated for different users. Log data (such as static data, running records, and security information) is stored in a metadata warehouse in MaxCompute.

- Static data: data that is not automatically deleted after it is generated.
- Running record: records of the process of a running task. The records exist in only one partition.
- Security information: data that originates from Table Store, such as whitelists and ACL tables.

Metadata warehouse: When you use MaxCompute to analyze its running status, the metadata stored in MaxCompute is aggregated into this table for convenient analysis.

## 5.24.2.4. IP address whitelists

MaxCompute security provides multiple levels of access control for projects, including multi-tenancy and security authentication. Only users with AccessKey pairs (AccessKey ID and AccessKey Secret) that have been authorized can access and perform operations on the data. This topic describes how to configure access control policies by using an IP address whitelist based on the preceding access authentication.

Note

- You can obtain the IP address as follows:
  i. If you use MaxCompute AdminConsole (the odpscmd tool) in a cluster such as AG, you can directly obtain the IP address of a server.
  ii. If you use an application system such as DataWorks or DataX to access project data, you must obtain the IP address of the server where DataWorks or DataX is deployed.
  iii. If you use a proxy server or the multi-hop proxy server to access the MaxCompute service, you must obtain the IP address of the last-hop proxy server.
  iv. If you access the MaxCompute service from an ECS instance, you need to obtain the NAT IP address.
- The IP address formats are as follows:

  Three formats of IP address ranges are supported. If you want to enter multiple IP addresses, separate them with commas (,). 1. Separate IP addresses 2. IP address range, starting and ending IP addresses separated by a hyphen (-) 3. IP address range specified by a subnet mask

  Examples:

```
10.32.180.8,10.32.180.9,10.32.180.10
-- Separate IP addresses
10.32.180.8-10.32.180.12
-- IP address range
10.32.180.0/23
-- IP address range specified by a subnet mask
```

The procedures to configure an IP address whitelist at the group, project, and system levels are as follows:

## Configure a project group-level IP address whitelist

After you configure an IP address whitelist for a project group, projects in the group are subject to this configuration restriction.

The configuration method is as follows:

1. In the AdminConsole, choose **MaxCompute Configuration > Group Management**. Select the group to be configured and double-click the group name.
2. In the Configuration dialog box that appears, configure the whitelist, and then click **Complete Modification**.
3. You can view the configuration result on the Project Group Attributes tab.

## Configure a project-level IP address whitelist

If a project does not belong to a project group, you can configure the whitelist at the project level.

The configuration method is as follows:

1. In the AdminConsole, choose **MaxCompute Configuration > Project Management**. Select the project to be configured and click the whitelist settings icon on the right side.


2. In the Configuration dialog box that appears, configure the whitelist, and then click **Save**.
3. You can view the configuration result on the Project Attributes tab.

> **Note** A project owner can also run the setproject command to set project attributes. For example, setproject odps.security.ip.whitelist="IP list separated by commas".

## Configure a system-level IP address whitelist

Some upper-level business systems such as DataWorks need to access all projects in MaxCompute instances. If there is no system-level whitelist configured, changes to the IP addresses will have to be configured manually for every project and project group whitelist one by one. This process is very prone to making configuration errors. MaxCompute supports system-level IP address whitelists to globally configure a whitelist for all MaxCompute instances. System-level whitelists are configured at the application level.

The configuration method is as follows:

1. In the AdminConsole, choose **MaxCompute Configuration > System-Level Whitelist Management**. The Configuration dialog box appears.
2. In the Configuration dialog box, configure the whitelist, and then click **Save**.
3. After completing the configuration, you can click **Add** to configure an IP address whitelist for another application.

## Precautions

1. The URL of Apsara Stack MaxCompute AdminConsole is *http://{odps_ag}:9090*, which corresponds to port 9090 of MaxCompute AG.

2. When configuring a whitelist for the first time, you must ensure that the IP address of the local server is included in the whitelist. Otherwise, the local server is not accessible. If the whitelist is configured incorrectly, the system administrator must modify the whitelist configuration from a management system like AdminConsole.
3. After a whitelist is configured for a project or project group, the corresponding project can only be accessed by IP addresses that are included in the whitelist. If some basic systems such as DataWorks need to access the project, ensure that the IP address of the server where DataWorks resides is included in the whitelist.
4. For added information security, you can also restrict access to MaxCompute from a service through policies even if the IP address whitelist allows the service to access MaxCompute. This is another level of fine-grained access control.
5. To access the MaxCompute service through a proxy server, add the IP address of the last-hop proxy server to the whitelist.

### Impact and effect

1. Before the configuration, there are no restrictions on MaxCompute service IP addresses accessing a project.
2. After the configuration, only the IP addresses and IP address ranges that comply with the configured rules are able to access the project. Based on the original AccessKey ID and AccessKey Secret authentication, the IP address rule check is superimposed.
3. To allow basic systems such as DataWorks, DataX, and DPC to access MaxCompute projects, you must also find the IP addresses of the servers where the systems reside and add them to the IP address whitelist.

# 5.25. DataWorks

## 5.25.1. Permission isolation for development and production environments

DataWorks manages code and configurations by workspace. Two workspace modes are available, which are Basic Mode and Standard Mode.

A workspace created in Standard Mode can isolate the development and production environments. Take the MaxCompute engine as an example. A workspace created in Standard Mode requires two MaxCompute projects: one for the development environment and the other for the production environment. Data in the two environments is completely isolated from each other.

Developers can only operate development environment data on the DataStudio page of DataWorks. Changes to production environment data take effect only after an administration expert publishes the changes. A workspace created in Standard Mode allows you to strictly control table permissions. Developers are prohibited from arbitrarily operating tables in the production environment to guarantee data security.

The development and production environments are integrated for a workspace created in Basic Mode. This type of workspace features fast iteration. Code takes effect immediately after being submitted, without requiring you to publish it. However, permissions of the development and production environments are not isolated.

# 5.25.2. Authentication and authorization

## 5.25.2.1. Access control

### Logon control

You can use your Alibaba Cloud account to log on to the Resource Access Management (RAM) console and create multiple RAM users. By creating policies and attaching them to RAM users, you can enable RAM users who meet specified conditions to access DataWorks. For example, you can specify that only RAM users who use the specified IP address or Classless Inter-Domain Routing (CIDR) block, enable multi-factor authentication (MFA), and use the HTTPS access protocol can access DataWorks.

By specifying the IP addresses or CIDR blocks that have access to DataWorks, you can further prevent unauthorized access and ensure data and business security. For example, when your AccessKey is inadvertently lost or stolen, DataWorks can prevent access from unauthorized IP addresses (such as IP addresses that are not in your internal network) before you create a new AccessKey.

### Sandbox isolation

A workspace is a basic unit for isolating user data in DataWorks. All nodes in the workspace run in the sandbox to prevent data leakage. Sandbox isolation also prevents developers from jeopardizing third-party data stores on the public network due to unauthorized use of external resources. By default, DataWorks only allows the following access scenarios:

- In DataStudio, developers can only access a specified compute engine.
- In Data Integration, developers can only access data stores that have been registered.

If developers need to access external resources outside the workspace in addition to the above two scenarios, the workspace administrator must add the resources to the sandbox whitelist in advance.

## 5.25.2.2. Permission management

### Role management

DataWorks defines seven roles in permission management, including owner, administrator, developer, O&M engineer, deployment engineer, security administrator, and guest.

| Role | Permission description |
| --- | --- |
| Owner | Indicates the user who owns a workspace. The owner has all permissions of a workspace. |
| Administrator | Indicates an administrative user entrusted by the owner. An administrator has all permissions of a workspace except for deleting the workspace. |

| Role | Permission description |
| --- | --- |
| Developer | Indicates a user who operates the development environment. A developer has the permissions to develop nodes and workflows and operate data in the development environment. |
| O&M engineer | Indicates a user who operates the production environment. An O&M engineer has the permissions to terminate, rerun, and deploy nodes in the production environment. |
| Deployment engineer | Indicates a user who connects the development and production environments. A deployment engineer has the permissions to publish code from the development environment to the production environment. |
| Security administrator | Indicates a data security manager. A security administrator has the permissions to manage the configuration in the Data Protection module. |
| Guest | Indicates a user with the minimum permission. A guest can only view code instead of performing any other operations. |

### Permission management

DataWorks allows you to manage data permissions on a workspace. You can authorize permissions by table or field and view and audit permissions.

### Data download control

DataWorks gives you full control over configuration data download to reduce the risk of data leakage and ensure data security.

## 5.25.3. Data encryption

All underlying data of DataWorks is encrypted for storage and transmission, including user code, workflow configuration, and data store connection information. Only authorized users can view, use, and modify the data.

## 5.25.4. Sensitive data protection

DataWorks supports data identification, sensitive data discovery, data classification and grading, desensitization, access monitoring, risk discovery and alerting, and audit.

- Data identification: automatically identifies sensitive data in a workspace based on preset rules.
- Data classification and grading: allows you to define different levels of data confidentiality and provide separate access control permissions for each level.
- Desensitization: desensitizes sensitive data by masking, aliases, and hashing.
- Access monitoring: monitors the access to and export of sensitive data.

- Risk discovery: monitors sensitive data access behavior in specific scenarios.

# 5.26. Realtime Compute

## 5.26.1. Platform security

### 5.26.1.1. Resource isolation

Realtime Compute projects are isolated based on permissions. Only authorized users can access or perform operations on a project and its sub-products.

Realtime Compute allows you to isolate resources at the project level. For example, if a sharp increase in streaming data inputs results in a higher CPU usage for the job of a user's project, the CPU usage of your job is not impacted. This is enabled by the application of virtualization technologies at the underlying layer of Realtime Compute.

### 5.26.1.2. Authentication and authorization

#### Realtime Compute accounts

You can only log on to the Realtime Compute console using Alibaba Cloud accounts, which are managed based on the username, password, and signature key. The accounts comply with the existing security system of Alibaba Cloud. To ensure the security of accounts, the HTTPS protocol is used for transmission.

#### Data store accounts

In Realtime Compute, the accounts of data stores are required to create source and result tables. We provide Resource Access Management (RAM) and Security Token Service (STS) to prevent your business data from leaking due to the loss of account information.

### 5.26.1.3. Data security

Realtime Compute ensures the security of Realtime Compute system data and business data.

#### System data security

The security of system data is ensured by Realtime Compute. The following measures have been taken to ensure the security of system data:

- The HTTPS protocol is used for transmission.
- The Advanced Encryption Standard (AES) is used to encrypt the information about the connection with data stores. This helps to prevent sensitive information from leaking.
- Comprehensive attack tests have been performed to ensure high-level security.

#### Business data security

Realtime Compute does not store business data of users. The security of business data is ensured by external Alibaba Cloud data stores. For more information, see security models and best security practices of Alibaba Cloud data stores.

### 5.26.1.4. Business process

### Business process security

Realtime Compute defines a strict development process for streaming data analysis, and provides separate pages for data development and administration in its console. This guarantees a complete and secure business process while minimizing adverse effects on user experience.

- Code versions

  Realtime Compute allows you to compare code versions and roll back to an earlier version. This helps you trace the code and troubleshoot faults.

- Standalone debugging tool in the IDE

  Realtime Compute offers a debugging tool in the integrated development environment (IDE), which allows you to debug the code without affecting online data. With this tool, you can specify data for source tables, dimension tables, and result tables to create a job, and then debug the data offline. This ensures that running jobs are not affected.

- Job publishing process

  Realtime Compute offers a job publishing process that prevents running jobs from being affected by the changes of offline code. After you debug the new code, you can publish the job and view it on the Administration page of the Realtime Compute platform. The running job does not automatically use the new code. Instead, you must confirm the changes, terminate the running job, and start the job again using the new code. In this way, you can exercise a complete control over the code that is used for the job to be published.

# 5.27. Machine Learning Platform for AI

## 5.27.1. Security isolation

### Tenant isolation

When users use Apsara Stack Machine Learning Platform for AI, the computing and storage resources assigned to each user are isolated and secured with the multi-tenancy technology.

- Network virtualization plays a crucial role in tenant isolation. Based on the virtual network established at the upper layer of the physical network, Machine Learning Platform for AI isolates the algorithm jobs of different users. This allows each job to manage its own computing and storage resources in the distributed cluster. When running a distributed deep learning job, the system uses the PCIe technology to connect a GPU instance to a Docker instance. This guarantees GPU computing performance while isolating resources for security.

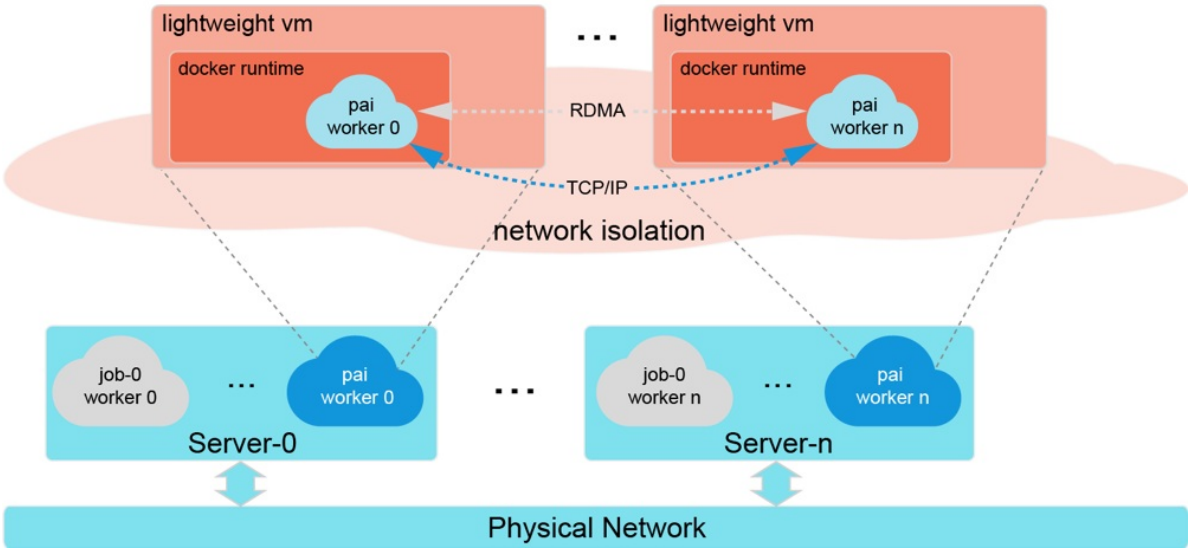


- When you activate Machine Learning Platform for AI, the DataWorks service is also activated. Based on centralized resource management, DataWorks creates a dedicated gateway resource group for each user. The gateway resource group allows the user to run SQL, MapReduce, and other types of tasks securely. Tasks performed by the user do not affect other users.

### Secure isolation for online model service

The online model service stores the AccessKey pair of the user who deploys the service in ApsaraDB for RDS. It verifies each request and rejects if it fails the authentication. The online model service uses the container technology to isolate tenant resources, including memory and CPU resources. This way, other services deployed on the same server are not affected.

# 5.27.2. Authentication

## 5.27.2.1. Identity authentication

### Machine Learning Platform for AI authentication

When a user accesses a machine learning application from a browser, the following identity authentication procedure is performed:

1. All requests from the browser are transmitted to LoginFilter. LoginFilter calls the API of the DataWorks SSO module to perform identity authentication. If the user has not logged on to the system, the user is redirected to the Apsara Stack logon page. Related cookies are written to the system after logon.
2. All requests contain a ctoken to prevent Cross-Site Request Forgery (CSRF) attacks.
3. All requests to add, delete, modify, or query resources are verified to ensure that the user has permissions to access the resources. This prevents unauthorized actions.
4. When the DataWorks scheduling module or other third-party modules call the API of Machine Learning Platform for AI, the token center of DataWorks is used to verify permissions in all requests.

### Identity authentication for online model service

A user can obtain an AccessKey pair in the Apsara Stack console. An AccessKey pair consists of an AccessKey ID and an AccessKey Secret. The AccessKey ID is public and is used to identify a user. The AccessKeySecret is private and is used to authenticate a user.

When a user sends a request to the online model service, the online model service authenticates the user as follows:

1. Convert the request into a signature string in the format specified by the online model service.
2. Use the AccessKey Secret and the HMAC algorithm to encrypt the signature string and generates a verification code. The verification code is time stamped to prevent replay attacks.
3. After the online model service receives the request, it locates the AccessKey Secret based on the AccessKey ID, uses the AccessKey Secret to decrypt the signature string, and generates a verification code.
   - If the generated verification code is the same as the one in the request, the online model service considers the request to be valid.
   - If the generated verification code is different from the one in the request, the online model service rejects the request and returns an HTTP 403 error message.

## 5.27.2.2. Access control

After your machine learning application has been deployed using the online model service, the system generates a token for accessing the prediction API. This token is required when you access the deployed application. This token is private. You must ensure that it is kept by the application user.

## 5.27.2.3. RAM and STS

Resource Access Management (RAM) is a resource access control service provided by Apsara Stack. With RAM, you can avoid sharing the AccessKey pair of your Apsara Stack account with other users, and assign minimum operation permissions to different RAM users based on the least privilege principle. RAM allows you to use the primary account to create RAM user accounts. The primary account can assign resource access permissions to RAM users.

Security Token Service (STS) is a temporary access token service provided by Apsara Stack to manage temporary access permissions. STS can generate a temporary access token for the user. The access permission and expiration date of the token are defined by the user. The access token expires automatically upon the expiration date.

When using the deep learning module, you can go to the OSS quick authorization page in the Apsara Stack console to perform quick authorization:

1. OSS quick authorization grants an OSS bucket read and write permissions to MaxCompute service account odps.aliyuncs.com.
2. After the authorization, role **AliyunODPSPAIDefaultRole** is created in RAM. Each role has a globally unique resource descriptor called **RoleArn**. The format is `acs:ram::$accountID:role/$roleName`.
3. After the role is created, Machine Learning Platform for AI uses the AccessKey pair of account odps.aliyuncs.com to call the AssumeRole operation of STS.
4. If the call is successful, you will obtain a temporary AccessKey pair (AccessKeyId and

AccessKeySecret) and a temporary STS token (SecurityToken). The default validity period is 3,600 seconds. With the temporary AccessKey pair and STS token, you can read and write the OSS bucket in the corresponding project.

The role information contained in **AliyunODPSPAIDefaultRole** is as follows:

```
{
 "Statement": [
  {
   "Action": "sts:AssumeRole",
   "Effect": "Allow",
   "Principal": {
    "Service": [
     "odps.aliyuncs.com"
    ]
   }
  }
 ],
 "Version": "1"
}
```

The role authorization policy is as follows:

```
{
 "Version": "1",
 "Statement": [
  {
   "Action": [
    "oss:GetObject",
    "oss:ListObjects",
    "oss:DeleteObject",
    "oss:ListParts",
    "oss:PutObject",
    "oss:AbortMultipartUpload"
   ],
   "Resource": "*",
   "Effect": "Allow"
  }
 ]
}
```

## 5.27.3. Data security

### Data security of Machine Learning Platform for AI applications

Machine Learning Platform for AI uses MaxCompute for big data computing and storage. MaxCompute projects are used to store data of different users to ensure that the users cannot access one another's data.

For financial institutes and other scenarios that require higher data security, MaxCompute allows you to enable or disable protection for data stored in projects. You can use this feature to allow or prohibit the export of data. For more information, see Project data protection.

### Data security of online model service

The online model service uses three types of data:

- Online model service metadata: The metadata is stored in and secured by ApsaraDB for RDS.
- Kubernetes (K8s) metadata: The metadata is stored in etcd, which provides services and data backup with three replicas. To access etcd, you need certificates kept by the K8s cluster administrator.
- Monitoring data: The monitoring data is stored in Alibaba Cloud disks, the type of which decides the type of data backup.

## 5.27.4. Log audit

Machine Learning Platform for AI provides a request log. User access records are automatically written into a designated file in the specified format based on the frequency of user access. The user access log is used for auditing or action analysis. The request log contains information such as the request time, source IP address, request method, request URL, request user ID, processing duration, and error code.

# 5.28. E-MapReduce (EMR)

## 5.28.1. Platform security

### 5.28.1.1. Access control

E-MapReduce (EMR) uses Resource Access Management (RAM) to grant multiple levels of permissions to different RAM users and control access to resources. You can create authorization policies and attach them to RAM users to control their data access. Users must use the credentials of their RAM users to log on to the EMR console.

For data stored in Object Storage Service (OSS), the following limits apply:

- All buckets in OSS are listed, but you can only access authorized buckets.
- You can only view the data of authorized buckets.
- A job can only read and write data from and to an authorized bucket.

### 5.28.1.2. Authentication

EMR supports Kerberos authentication. If you run open source components on a cluster in the safe mode of Kerberos, only authenticated clients can access services on the cluster, such as HDFS.

Kerberos is a network authentication protocol. EMR uses Hadoop Authentication Service (HAS). Open source big data ecosystems (Hadoop/Spark) only support built-in Kerberos for security authentication. HAS provides a new authentication method, Kerberos-based token authentication. This authentication method enables Hadoop or Spark to support authentication methods other than Kerberos by connecting to existing authentication and authorization systems. It also hides the complexity of Kerberos from end users. Kerberos-based token authentication supports most components in big data ecosystems and requires little or no changes to components.

The original Kerberos authentication method provided by HAS is still available for all components.

## Kerberos identity authentication

Kerberos is an identity authentication protocol based on symmetric-key cryptography. As an independent third-party identity authentication service, Kerberos provides identity authentication and supports single sign-on (SSO). After a principal is authenticated, the principal is allowed to access multiple application services, such as HBase and HDFS.

Kerberos authentication includes two steps: 1. Key Distribution Center (KDC) authenticates a client. 2. A service authenticates the client.

Kerberos authentication

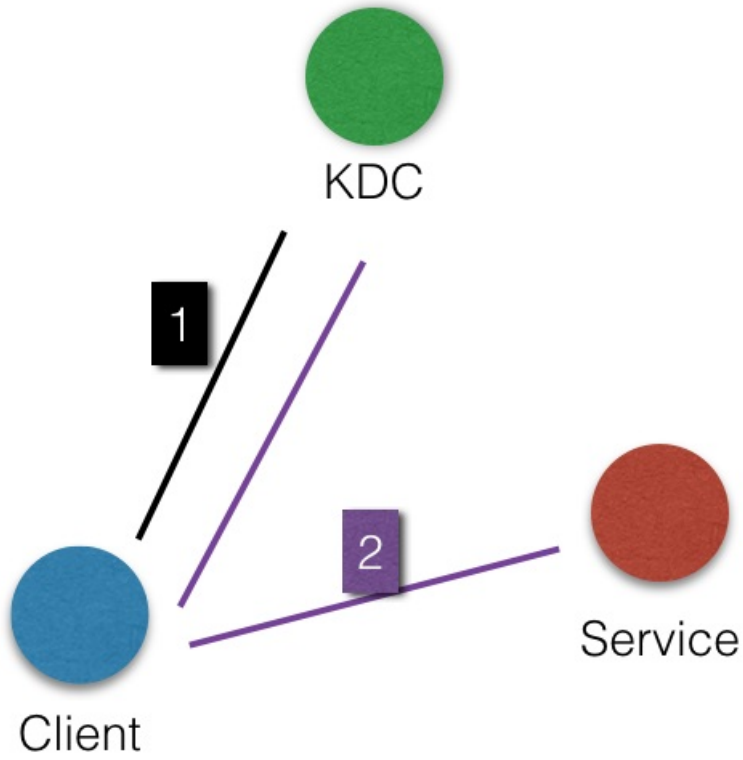


- KDC: indicates a Kerberos server.
- Client: indicates a client (principal) that needs to access an application service on a cluster.

KDC and the service authenticate the identity of the principal.

- Service: indicates a service with Kerberos authentication enabled, such as HDFS, YARN, or HBase.

**Authenticate the identity of a client by KDC**

A client can access a service with Kerberos authentication enabled only after it is authenticated by KDC. If the identity authentication succeeds, the client gets a Ticket Granting Ticket (TGT), which can be used to access a service integrated with Kerberos.

**Authenticate the identity of a client by a service**

After a client obtains a TGT, they can use the TGT to access application services. The client uses the TGT and the name of the service to be accessed such as HDFS to obtain the Service Granting Ticket (SGT) from KDC, and then uses the SGT to access the service. The service uses relevant information to authenticate the client. After the authentication succeeds, the client can access the service.

## Authorization

After you create users, you need to grant permissions on components to those users. Authorization control is available regardless of whether authentication is enabled.

- HDFS authorization

  When HDFS authorization is enabled, only users granted the relevant permissions can access HDFS and perform operations such as reading data or creating folders.

- YARN authorization

  YARN supports service-level and queue-level authorization.

  Service-level authorization

  - Service-level authorization controls the access of specified users to cluster services, such as job submissions.
  - Authorization policies are configured in the hadoop-policy.xml file.
  - Service-level authorization must be verified before all other authorizations, such as the authorization of the permissions to perform actions on HDFS or submit a YARN job for queuing.

  Queue-level authorization

  - YARN uses queues to control the actions that can be performed on resources. The following types of queue schedulers are available: capacity scheduler and fair scheduler.

- Hive authorization

  Hive has the following two built-in authorization schemes. You can configure both of them because they are compatible with each other.

  - Storage-based authorization
  - SQL standards-based authorization

- HBase authorization

If authorization is disabled, all accounts can access HBase clusters to perform any action, such as disabling tables, dropping tables, and performing major compaction. For clusters without Kerberos authentication, users can use a forged identity to access cluster services. This is the case even if HBase authorization is enabled. We recommend that you create clusters with Kerberos authentication enabled to ensure high security.

- Kafka authorization

  If Kafka authorization is enabled but Kafka authentication (or simple username-password cryptography) is disabled, users can use a forged identity to access cluster services. We recommend that you create Kafka clusters with Kerberos authentication enabled to ensure high security.

- Ranger authorization

  Apache Ranger is a centralized security framework that manages fine-grained access control across Hadoop components, such as HDFS, Hive, YARN, Kafka, Storm, and Solr. You can manage access policies in the Apache Ranger console.

## 5.28.1.3. Data security

This topic describes the data security of EMR.

Hadoop Distributed File System (HDFS) stores files as data blocks. Each data block has multiple replicas (three by default). We recommend that you store replicas on multiple racks. For example, if you set the replication factor to 3, you can store two replicas on different nodes in the local rack and store another replica on a node in a remote rack.

HDFS scans all replicas periodically. If a replica is found missing, HDFS quickly generates a new replica. If a node fails, HDFS quickly recovers all data on that node. If you use disks on Alibaba Cloud, three replicas are created for each disk. If an error occurs on one of the replicas, HDFS copies data from another replica to the failed replica to ensure data reliability.

HDFS is proven to be a reliable data storage system. You can also integrate HDFS with Object Storage Service (OSS) based on cloud features to back up data. This ensures higher data reliability.

Data synchronization

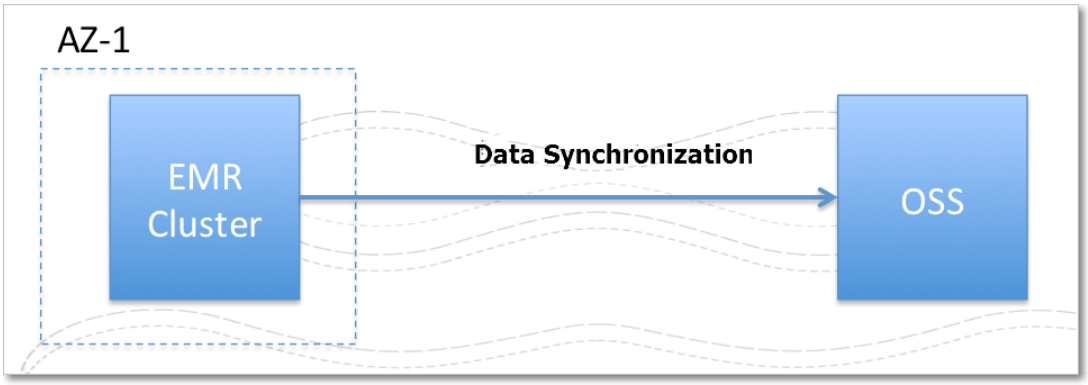


You can use EMR schedulers or third-party schedulers to synchronize data to OSS on a schedule.

The synchronization interval is also the maximum time range allowed for lost data. For example, if data is synchronized hourly, the maximum time range allowed for lost data is also one hour. If data is synchronized at 30-minute intervals, the maximum time range allowed for lost data is also 30 minutes.

Data backup

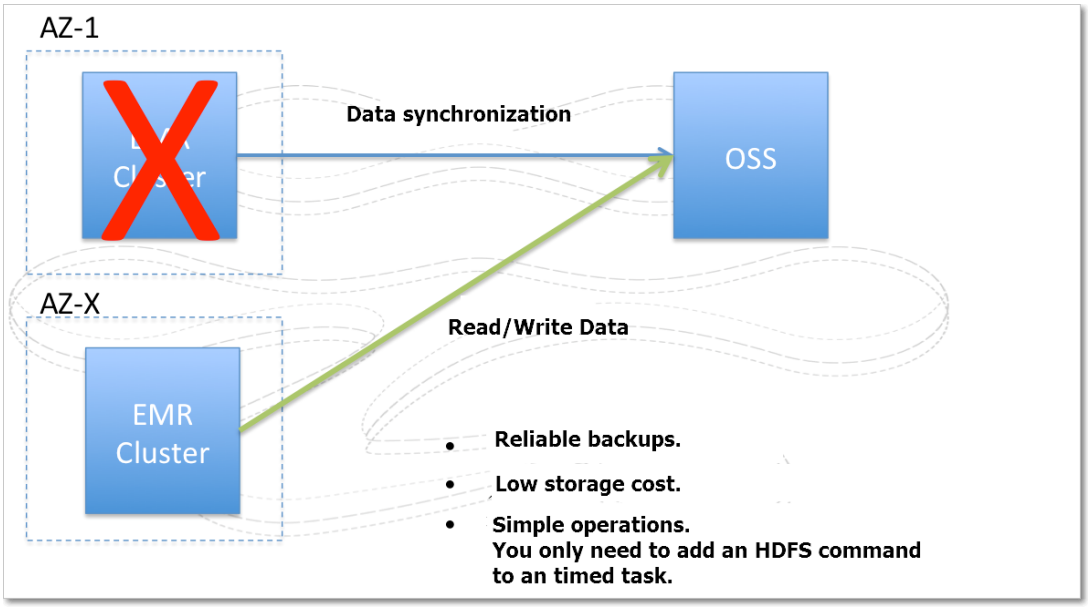


If a cluster fails, EMR creates a new cluster and reads the data of the original cluster from OSS. If the original cluster contains metadata, EMR also creates metadata in the new cluster to ensure service continuity. After the original cluster recovers, EMR synchronizes the updated data from OSS to the cluster.

Disaster recovery for services

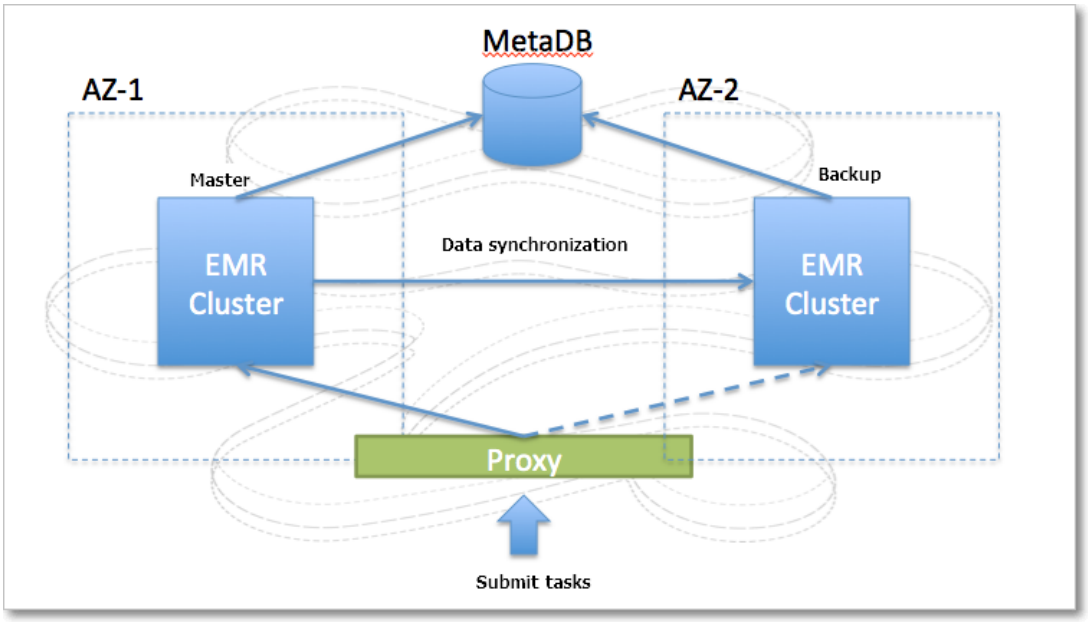


You can create two clusters with the same compute capability and storage capacity in different zones in primary/secondary mode. This way, the clusters share the same metadatabase. The metadatabase uses a geo-redundancy architecture, which is similar to the three-node architecture of ApsaraDB for RDS. EMR synchronizes data on the primary cluster to the secondary cluster in near real time to ensure data consistency. You can use the DistCp tool to synchronize data on a schedule. However, this method has low time-efficiency in terms of disaster recovery. Another method is to monitor data updates on the primary cluster. When a data update occurs, EMR synchronizes incremental data to the secondary cluster within seconds.

A proxy is used to submit jobs. When you submit a job, you do not need to specify or know the destination cluster. The job is submitted to the primary cluster. However, if the primary cluster is inaccessible, the job is submitted to the secondary cluster.

In cases where external data sources are used, EMR must move all external data to the secondary cluster when a primary/secondary cluster switchover occurs. Therefore, you must make sure that the external databases also support disaster recovery.

# 5.28.2. Authorization control

This topic describes the authorization control of EMR.

After you create users, you need to grant permissions on components to users. Authorization control is available regardless of whether authentication is enabled.

- **HDFS authorization**

  When HDFS authorization is enabled, only users granted the relevant permissions can access HDFS and perform operations such as reading data or creating folders.

- **YARN authorization**

  YARN supports service-level and queue-level authorization.

  - **Service-level authorization**
    - Service-level authorization controls the access of specified users to cluster services, such as job submissions.
    - Authorization policies are configured in the hadoop-policy.xml file.
    - Service-level authorization must be verified before all other authorizations, such as the authorization of the permissions to perform actions on HDFS or submit a YARN job for queuing.
  - **Queue-level authorization**

    YARN uses queues to control the actions that can be performed on resources. The following types of queue schedulers are available: `capacity scheduler` and `fair scheduler`.

- **Hive authorization**

  Hive has the following two built-in authorization schemes. You can configure both of them because they are compatible with each other.

  - Storage-based authorization
  - SQL standards-based authorization

- **HBase authorization**

  If authorization is disabled, all accounts can access HBase clusters to perform any action, such as disabling tables, dropping tables, and performing major compaction. For clusters without Kerberos authentication, users can use a forged identity to access cluster services. This is the case even if HBase authorization is enabled. We recommend that you create clusters with Kerberos authentication enabled to ensure high security.

- **Kafka authorization**

If Kafka authorization is enabled but Kafka authentication (or simple username-password cryptography) is disabled, users can use a forged identity to access cluster services. We recommend that you create Kafka clusters with Kerberos authentication enabled to ensure high security.

- **Ranger authorization**

  Apache Ranger is a centralized security framework that manages fine-grained access control across Hadoop components, such as HDFS, Hive, YARN, Kafka, Storm, and Solr. You can manage access policies in the Apache Ranger console.

# 5.29. DataHub

## 5.29.1. Platform security

### 5.29.1.1. Data isolation

DataHub uses symmetric encryption based on an AccessKey (AccessKey ID and AccessKey secret) to verify the identity of requesters. Additionally, each HTTP request is signed and authenticated through signature verification. DataHub isolates the data of different users from one another by storing the data in the Apsara Distributed File System.

DataHub ensures that user data and metadata are stored separately.

### 5.29.1.2. Authentication and Authorization

#### Authentication

You can generate an AccessKey in the console. An AccessKey consists of an AccessKey ID and an AccessKey secret. The AccessKey ID identifies a user. The AccessKey secret must be kept strictly confidential and is used to authenticate user identity.

When you send a request to DataHub, you must first generate a signature string for the request in the format specified by DataHub. Then encrypt the signature string by using the AccessKey secret to obtain the request signature. After receiving a request, DataHub maps the AccessKey ID to the AccessKey secret and extracts the signature string and verification code. If the extracted verification code is the same as the one provided, the request is considered valid. Otherwise, DataHub rejects the request and returns the 403 status code.

#### Access control

You can access DataHub resources by using the Alibaba Cloud account or RAM users. With an Alibaba Cloud account, you can create one or more RAM users. An Alibaba Cloud account can grant RAM users permission to access DataHub resources by using RAM authorization policies.

- When you access a DataHub resource by using an Alibaba Cloud account, DataHub checks whether the account is the owner of the corresponding resource. Only the resource owner has the permission to access the corresponding resource.
- When you access a DataHub resource as a RAM user, DataHub checks whether you are authorized by the corresponding Alibaba Cloud account to access the resource and whether the Alibaba Cloud account owns the resource. For more information about RAM authorization, see RAM and STS authorization.

> Note DataHub does not support resource authorization between Alibaba Cloud accounts.

### RAM and STS authorization

DataHub supports Resource Access Management (RAM) and Security Token Service (STS) authorization.

**RAM** is a resource access control service provided by Alibaba Cloud. With an Alibaba Cloud account, you can create RAM users by using the RAM service. You can grant RAM users the permission to access certain resources owned by the Alibaba Cloud account.

**STS** is a short-term resource access control service provided by Alibaba Cloud. STS generates a temporary access credential for users. You can define authorization rules and set the validity period for the credential. The access credential automatically expires after the specified validity period.

DataHub resources are authorized by using RAM authorization policies. You must specify Action, Resource, and Effect in a policy. An example of policy content is as follows:

```
{
 "Version": "1",
 "Statement": [
  {
   "Action": [ "dhs:GetRecords"],
   "Resource": "acs:dhs:cn-hangzhou:1001:projects/A/topics/B",
   "Effect": "Allow"
  }
 ]
}
```

> Note The format of Resource is `acs:dhs:{Region}:{User}:{DataHubResource}` .

This sample policy indicates the control of access to a DataHub resource that resides in a particular region. Specifically, the policy grants RAM user 1001 the permission to read data from Topic B of Project A that resides in the China (Hangzhou) region.

DataHub supports fine-grained permission control policies. You can control resource access as needed. For more information, see the DataHub documentation on the Alibaba Cloud website.

## 5.29.1.3. Data encryption

Data stored in DataHub is not encrypted and each record can only be retained in DataHub for up to seven days. To ensure secure data transmission, DataHub provides RESTful APIs for data transmission by using HTTPS.

## 5.29.1.4. Data security

In Apsara Stack, DataHub provides a flat memory model that slices linear addresses into shards. A shard is referred to as a chunk. Each chunk has three copies. The copies are stored on different nodes in a cluster according to the distribution policy to ensure the reliability of data.

The Apsara Stack data storage system consists of three roles: master, chunk server, and client. Each write operation in DataHub is executed by the client. The execution process is as follows:

1. The client determines the chunk of the write operation.
2. The client queries the master for the storage location of the three copies of the chunk.
3. The client sends a write request to the three corresponding chunk servers based on the results returned from the master.
4. If the write operation fails, the client returns an error message to the user.

The distribution policy evaluates the disk usage, the distribution on different switch racks, power supply, and the load of all chunk servers. This ensures that three copies of a chunk are stored in different chunk servers on different racks, effectively preventing the unavailability of data caused by a chunk server or a rack failure.

If a node is damaged or any hard disk on a node fails, some chunks in the cluster will have fewer than three valid copies. In this scenario, the master activates the replication mechanism and copies data between chunk servers to ensure that all chunks in the cluster have three valid copies.

All operations on data in DataHub, including inserting, updating, and deleting data, are synchronized to the three copies. This mechanism guarantees data reliability and consistency.

After a delete operation is performed, the corresponding space in the Apsara Distributed File System is released. To ensure data security, the content will be erased before the space can be used again.

## 5.29.2. Tenant security

### 5.29.2.1. Audit logging

DataHub logs various types of information for security and troubleshooting purposes. DataHub creates a topic in the system to store log data, including queries per second (QPS), request duration, processing duration, source IP address, and status code. DataHub displays some of the statistics in the console to help you analyze issues.

# 5.30. Graph Analytics

## 5.30.1. Platform security

### 5.30.1.1. Security isolation

Graph Analytics supports multi-tenant isolation. Users are walled off from seeing the settings or querying the data created by other users. Tenants can only obtain their own metadata settings. Different metadata corresponds to different business data, so one tenant can only query the business data corresponding to its own metadata.

### 5.30.1.2. Authentication

### Identity authentication

Graph Analytics currently supports two authentication methods in Apsara Stack:

- Authentication in Graph Analytics: Log on with the user password created in the Graph Analytics user system. To prevent password leakage, the password is MD5 encrypted and encrypted when transmitted on the network.
- Authentication of external users: Graph Analytics is connected to the user systems of customers. These systems are responsible for the security of their own users.

### Permission control

Graph Analytics controls the access permissions to functional modules, data rows, and columns based on the permissions of users.

## 5.30.1.3. Data security

Graph Analytics is deployed in distributed clusters. The control nodes and compute nodes are separated, which helps to avoid a single point of failure (SPOF). In addition, data is synchronized among the clusters through a distributed cache to prevent data loss during the failover.

## 5.30.1.4. Transmission encryption

Graph Analytics provides Web services by using the HTTPS protocol. HTTPS is a secure and reliable data transmission protocol. It can effectively prevent security issues that may occur when data is transmitted over networks.

## 5.30.1.5. System security

## 5.30.1.5.1. Vulnerability scanning mechanism

Before launched, Graph Analytics has been scanned by Apsara Stack for security vulnerabilities in the following aspects:

- System security scanning: targets the security vulnerabilities in the operating system of Graph Analytics.
- Middleware dependency scanning: targets the security vulnerabilities in the middleware used by Graph Analytics.
- Code vulnerability scanning: targets the security vulnerabilities in Graph Analytics's own code and the third-party open-source framework.

## 5.30.1.5.2. Update scheme for security vulnerabilities

For security vulnerabilities detected by Alibaba Cloud security groups, Apsara Stack security tests, or other methods, the research and development team of Graph Analytics will perform immediate version updates or iterative updates based on the severity of the vulnerability. This ensures that the new version conforms to the production security regulations of Alibaba Cloud.

## 5.30.1.5.3. System defense mechanism

Graph Analytics is launched in the Apsara Stack environment. The system defense mechanism of Graph Analytics depends on the defense mechanism of Apsara Stack systems.

#### 5.30.1.6. Infrastructure security

Graph Analytics is launched in the Apsara Stack environment and its infrastructure is secured by Apsara Stack security facilities. For more information, see the infrastructure security topic in Alibaba Cloud Security White Paper.

#### 5.30.1.7. Level-based data security

Graph Analytics meets the Apsara Stack V3 requirements for security level 4. Graph Analytics has obtained the following security certificates from the Apsara Infrastructure Management Framework: cacert.pem, privatecloud.pem, privatecloud_key.pem, and privkey.pem.

For Apsara Stack versions later than V3.3, you can connect to Graph Analytics by using the HTTPS protocol.

### 5.30.2. Tenant security

#### 5.30.2.1. Log audit

All user requests are logged, so they can be used for audit or specific behavior analysis. The logs include information such as the user name, IP addresses, operation content, and operation statuses.

## 5.31. Elasticsearch (on k8s)

### 5.31.1. Security isolation

Each Elasticsearch cluster uses independent storage. You can configure a password for a cluster and use this password to access the cluster. Clusters are independent of each other. The Elasticsearch process runs in a container and does not support the root account. Elasticsearch uses Java virtual machines (JVMs) to isolate resources. You can access only the directories of data and configuration files even if third-party plug-ins are used.

### 5.31.2. Authentication and authorization

#### 5.31.2.1. Identity authentication

Elasticsearch provides the security authentication feature to authenticate user identity. You can access Elasticsearch clusters by using the HTTP basic access authentication method. If the account or password that is configured on a client is invalid, you cannot use the client to access your cluster. Elasticsearch does not store passwords in plaintext. It uses the bcrypt algorithm to check whether passwords are valid.

#### 5.31.2.2. Access control

After you configure security settings for an Elasticsearch cluster, you can use its built-in role-based access control (RBAC) system to perform access control for the cluster at the field level. In addition, you can use the Kibana console to manage users, roles, and permissions.

### 5.31.3. Data security

The underlying layer of Elasticsearch uses shards to store data and dynamically modifies the number of replicas for the shards. Data is replicated from shards to replicas to ensure data security. If a data node in a cluster fails or a hard disk on a data node in the cluster fails, the cluster can detect damage to shards. The health states of related indexes and the entire cluster change from green to yellow. In addition, the replication mechanism is automatically triggered to clone a shard based on existing data. The cluster then returns to the green health state after all data is restored.

## 5.31.4. Transmission encryption

Nodes in an Elasticsearch cluster communicate with each other over SSL. This way, attackers cannot intercept network traffic or launch attacks to the system.

# 6.Apsara Stack Security

Apsara Stack Security, integrated with the strong data analysis ability of the Alibaba Cloud computing platform, is an achievement of Alibaba Group after many years of research on security technologies. It provides you with a variety of one-stop security services, including DDoS Traffic Scrubbing, Server Intrusion Detection, Web Application Firewall (WAF), and Threat Detection Service (TDS).

For more information about Apsara Stack Security, see the *Apsara Stack Security Technical Whitepaper* .

# 6.1. Overview

Apsara Stack Security, integrated with the strong data analysis ability of the Alibaba Cloud computing platform, is an achievement of Alibaba Group after many years of research on security technologies.

### Apsara Stack Security Standard Edition

Apsara Stack Security Standard Edition provides various security modules and on-premises security operations services, such as Network Traffic Monitoring System, Server Guard, Server Intrusion Detection, Web Application Firewall (WAF), and Threat Detection Service (TDS). It ensures the availability, confidentiality, and integrity of business systems and data in Apsara Stack.

Apsara Stack Security Standard Edition makes full use of the security features of Apsara Stack services and Apsara Stack Security to provide Apsara Stack users with one-stop security assurance that covers intrusion protection, threat detection, and centralized management. Apsara Stack Security Standard Edition enhances the security of the Apsara Stack environment in terms of technology and management, effectively protecting your Apsara Stack platform.

### Optional security services

Apsara Stack Security provides a wide range of optional services for you to choose from based on the business scenarios on the Apsara Stack platform.

For more information about Apsara Stack Security, see the *Apsara Stack Security Technical White Paper*.

# 6.2. Apsara Stack Security Standard Edition

## 6.2.1. Threat Detection Service

Threat Detection Service (TDS) detects possible intrusions or attacks by means of machine learning and data modeling. It detects zero-day vulnerabilities and potential virus attacks. TDS also provides you with up-to-date information about ongoing attacks, so that you can better monitor the security of your business systems. This prevents data leaks caused by network attacks and allows you to track hacker identities.

The following table describes features of TDS.

| Feature | Description |
| --- | --- |
| Overview | Provides a comprehensive security overview with statistics on security score, asset status, unhandled alerts, and handled alerts. |
| Security Dashboard | Displays the security data on dashboards, including assets, vulnerabilities, baselines, attack sources, and attack distribution. |
| Security Alerts | Allows you to view and handle security events, including suspicious process, webshells, unusual logons, sensitive file tampering, malicious processes, suspicious network connections, and web application threat detection. |
| Attack Analysis | Displays the attack trends and attack type distribution in the last 7 days and 30 days.<br>Displays the attack information such as the attack time, attack source, attacked assets, number of attacks, risk level, and attack type. |
| Cloud Service Check | Checks the security configurations of cloud services from the aspects of network access control and data security. It supports periodic checks that run automatically and manual checks. You can verify the check results or configure whitelist policies for the check results. |
| Application Whitelists | Allows you to add servers to a whitelist based on intelligent learning and identifies programs as trusted, suspicious, or malicious based on the whitelist. Unauthorized processes will be terminated. |
| Assets | • Server: displays the security statuses for servers. You can view the numbers of all servers, risky servers, unprotected servers, inactive servers, and new servers.<br>• Cloud Product: provides security status information for cloud services and supports SLB and NAT. |
| Security Reports | Allows you to query reports. For example, you can retrieve historical reports by report name. |

## Best practices

TDS provides features such as asset management, security monitoring, intrusion backtracking, attacker tracking, and intelligence warning. We recommend that you use TDS in the following scenarios to help visualize your cloud security:

- Security status awareness

  You can be fully aware of your cloud security, such as attacks, vulnerabilities, intrusions, protection effectiveness, business weaknesses, and the security status of services available for external use. TDS detects network- and server-layer attacks, advanced persistent threat (APT) attacks, and business security threats. It also identifies abnormal network connections and generates daily security reports.

- Countermeasures against intrusions

Your business systems in the cloud may be attacked. For example, during a sudden surge of the server load, you receive an SMS message indicating that your Elastic Compute Service (ECS) instances are being attacked. Your instances may be controlled to launch attacks, malicious advertisement links may appear on your web pages, or your data may be encrypted by an attacker by using ransomware. You can use the following features of TDS to defend against these types of intrusions:

  - Intrusion detection: TDS detects dozens of intrusion types, such as WannaCry ransomware attacks, intrusions through webshells, one-line Trojans, software viruses, and connections between internal servers and command-and-control (C&C) servers.
  - Intrusion behavior analysis: TDS analyzes the causes and processes of intrusions and collects evidence of hacker behavior during the intrusions.
  - Details of security events: You can view webshell addresses, protocol analysis results of DDoS attacks, process addresses, and attack prevention effectiveness.

- Real-time monitoring on dashboards

  TDS displays your cloud security status on dashboards. This improves teamwork efficiency and allows you to monitor the cloud security in real time.

# 6.2.2. Traffic Security Monitoring

Developed by Alibaba Cloud Security, the Traffic Security Monitoring module is able to monitor attacks within milliseconds. By performing in-depth analysis on the traffic packets mirrored from the Apsara Stack portal, this module can detect various attacks and abnormal behaviors in real time and coordinate with other protection modules to implement defenses. Throughout the Apsara Stack Security defense process, the Traffic Security Monitoring module provides a wealth of information and basic data support.

The following table describes the features provided by the Traffic Security Monitoring module.

| Feature | Description |
| --- | --- |
| Traffic data collection and analysis | Uses a bypass in traffic mirroring mode to collect inbound and outbound traffic through the interconnection switch (ISW) and generates a traffic diagram. |
| Unusual traffic detection | Uses a bypass in traffic mirroring mode to detect the unusual traffic that has exceeded the scrubbing threshold and reroutes the traffic to the DDoS Traffic Scrubbing module. The traffic rate (Unit: Mbit/s), packet rate (Unit: PPS), HTTP request rate (Unit: QPS), or number of new connections can be set as the threshold. |
| Malicious server detection | Detects attacks that are launched by malicious servers within the Apsara Stack network. |
| Web application protection | Uses a bypass to block common Web attacks at the network layer based on default Web attack detection rules. The attacks that can be blocked include Structured Query Language (SQL) injections, code and command execution, Trojan scripts, file inclusion attacks, and exploitation of upload vulnerabilities and common content management system (CMS) vulnerabilities. |
| Suspicious TCP connection blocking | Uses a bypass to send TCP RST packets to the server and the client to block layer-4 TCP connections. |

| Feature | Description |
| --- | --- |
| Network log recording | Records UDP and TCP traffic logs and the Request and Response logs of HTTP queries. These logs are used by Threat Detection Service (TDS) for big data analysis. |

**Best practices**

By checking the traffic at different times, in different regions, and from each IP address, you can identify the traffic peak and trough periods and view traffic distribution by rate or region. You can also check the top five IP addresses that generate the most traffic to block access from malicious IP addresses.

# 6.2.3. Server Guard

Server Guard provides security protection measures such as vulnerability management, baseline check, intrusion detection, and asset management for Elastic Compute Service (ECS) instances by means of log monitoring, file analysis, and feature scanning.

Server Guard uses the client-server model. To protect the security of ECS instances in real time, Server Guard clients work with the Server Guard server to monitor attacks, vulnerabilities, and baseline configurations at the system layer and the application layer on the ECS instances.

The following table describes the features provided by Server Guard.

| Category | Feature | Description |
| --- | --- | --- |
| Overview | Overview | Displays assets, vulnerabilities, exceptions, configuration defects, and events that require attention. |
| Servers | Server Fingerprints | Provides the following modules:<br>• Port: checks and displays the listening port information, including the listening port, protocol, process, IP address, and update time.<br>• Software: checks and displays the software installation information on servers, including the software name, software version, software installation directory, and update time.<br>• Process: checks and displays the process information, including the process name, process path, startup parameter, startup time, user, permission, process ID (PID), parent process, and update time.<br>• Account: checks and displays the host account information, including the account name, logon permission, root permission, user group, expiration time, last logon time, and update time.<br>• Scheduled Tasks: checks and displays the scheduled tasks of the host, including the task path, execution command, task cycle, account name, and update time. |

| Category | Feature | Description |
| --- | --- | --- |
| Threat Prevention | Baseline Check | Automatically detects configuration risks related to the system, account, database, weak password, and security compliance on your servers, and provides security hardening suggestions. This feature also checks database, system, and middleware assets. |
| | Vulnerabilities | Detects four types of vulnerabilities: Linux, Windows, Web CMS, and emergency vulnerabilities and provides vulnerability fix solutions. You can verify vulnerability fixes, view vulnerability details, and identify all vulnerabilities at one click. |
| Intrusion Prevention | Intrusions | Displays the alert information of affected host assets, including the number of alerting servers, the total number of unhandled alerts, and the number of urgent alerts. |
| | File Tamper Protection | Supports web page tamper-proofing and provides the blacklist and whitelist prevention modes. |
| | Virus Removal | Detects and removes virus and webshell. The system automatically detects and removes common trojan viruses, ransomware, mining viruses, and DDoS trojans. |
| Log Retrieval | Log Retrieval | Allows you to query logs for logon, brute-force attack, process snapshot, network connection, listening port snapshot, account snapshot, and process startup. |
| Server Settings | Client Installation | Allows you to view offline servers. You can install clients for the servers again based on the Client Installation Guide. You can uninstall the Server Guard client from the specified server. |
| | Protection Mode | Provides business first and protection first modes for different scenarios. |

**Best practices**

You can use the features of Server Guard to periodically perform baseline check for ECS instances, detect security threats and vulnerabilities on ECS instances, and fix them promptly for higher server security.

# 6.2.4. WAF

Based on the intelligent semantic analysis engine, Web Application Firewall (WAF) defends against common attacks defined by Open Web Application Security Project (OWASP), including Structured Query Language (SQL) injections, cross-site scripting (XSS) attacks, common web server plug-in vulnerabilities, Trojan uploads, and unauthorized access to core resources. WAF filters out massive malicious access attempts to prevent leakage of website assets and data, and to safeguard website security and availability.

The following table describes the features provided by WAF.

| Category | Feature | Description |
| --- | --- | --- |
| Detection Overview | Detection Overview | Provides statistics on protection for the last 24 hours and the last 30 days. |
| | Access Status Monitor | Displays the top 100 access requests in real time. |
| | Export Detection Report | Allows you to export daily reports, weekly reports, and scheduled task reports. |
| | Attack Detection Statistics | Provides statistics on attack detection. |
| Detection Logs | Attack Detection Logs | Provides attack detection logs. The log list displays the processing results, attacked addresses, attack types, attacker IP addresses, and attack time. You can view log details for each attack. |
| | HTTP Flood Detection Logs | Provides HTTP flood protection logs. The log list displays logs for matched HTTP flood protection rules, including the request URL, the name of the matched rule, and the match time. You can filter logs based on the event generation time and the name of the HTTP flood protection rule. |
| | System operation log | Provides system operations logs, including usernames, operations, and IP addresses. |
| | Access Log | Provides access logs, including the access address, destination IP address, source IP address, request method, and response code. |
| | Protection site management | Allows you to create, delete, modify, enable, and disable function forwarding proxies of a protected site. |
| | Customized Rules | Allows you to create, delete, enable, and disable custom rules. This implements fine-grained HTTP access control for websites. |

| Category | Feature | Description |
| --- | --- | --- |
| Protection Configuration | Website Protection Policies | • Supports decoding methods, such as URL decoding, JSON parsing, Base64 decoding, hexadecimal conversion, backslash unescape, XML parsing, PHP deserialization, and UTF-7 decoding.<br>• Detects SQL injections, cross-site scripting (XSS), intelligence, cross-site request forgery (CSRF), server-side request forgery (SSRF), Hypertext Preprocessor (PHP) deserialization, Java deserialization, Active Server Pages (ASP) code injections, file inclusion attacks, file upload attacks, PHP code injections, command injections, crawlers, and server responses.<br>• Provides five built-in protection templates, including the template with default protection policies, monitoring mode template, anti-DDoS template, template for financial customers, and template for Internet customers. WAF allows you to customize the decoding algorithms in the templates, enable or disable each attack detection module separately, and configure the detection granularity. WAF also allows you to specify the Block Status Code parameter.<br>• Allows you to enable HTTP response detection and configure the length of the response body in detection rules.<br>• Allows you to configure the length of the request body in detection rules.<br>• Allows you to enable or disable detection timeout settings. |
| | HTTP Flood Protection | Allows you to configure access frequency control rules for domain names and URLs. This restricts the access frequency of IP addresses or sessions that meet the criteria, or blocks these IP addresses or sessions. Restricts the access frequency of known IP addresses or sessions or blocks these IP addresses or sessions. Supports the HTTP flood protection whitelist function. HTTP flood protection rules are not applicable to IP addresses or sessions in a whitelist. |
| | SSL Certificate Management | Allows you to upload certificate files and SSL private keys to manage SSL certificates. |

| Category | Feature | Description |
| --- | --- | --- |
| System Management | Node status | • Payload Status: displays the CPU utilization and memory usage.<br>• Node Network Status: displays the read throughput and write throughput.<br>• Detection Status: displays the queries per second (QPS) and the average detection time consumed by WAF nodes.<br>• Forward Status: displays the number of new connections per second and the average latency.<br>• Disk Status: displays the disk usage and total disk size. |
| | Syslog Configuration | Configures syslog to send logs and also configures the service- and system-related alert thresholds. |

Best practices

- Use WAF to prevent the leakage of sensitive information

  WAF effectively defends against security threats such as accesses from unauthorized URLs, accesses to the unauthorized content, and malicious crawling of sensitive information on websites.

- Use WAF to prevent WordPress reflection attacks

  WAF effectively prevents WordPress reflection attacks with custom HTTP-based access control list (ACL) rules.

# 6.2.5. Security Operations Center (SOC)

Security Operations Center (SOC) provides security administrators with centralized management of all users and the platform and analysis functions of Apsara Stack logs.

SOC provides the following features:

| Feature | Description |
| --- | --- |
| Dashboard | Allows you to view the overall security statistics and perform operations. |
| Security monitoring | Allows you to view the security events of all users and the platform. |
| Asset management | Allows you to view the security status of user assets and platform assets. |
| Log analysis | Analyzes logs from multiple data sources, detects unexpected alerts, and improves alert detection of Apsara Stack. |
| Report management | Allows you to quickly export reports for various purposes. |
| System configurations | Allows you to configure system features such as alerts, updates, global policies, and account management. |

## Best practices

- Scenario 1: routine monitoring

SOC regularly inspects system security. Currently, SOC focuses on security issues on the users. The following features are provided:

- Urgent risk detection: checks for urgent security risks on a daily basis. Security risks include user security alerts, vulnerabilities, and server configuration risks.
- Risk management: identifies and handles high-risk security alerts, vulnerabilities, and server configuration risks.
- Attack data collection: shows the number of attacks and attack protection information.
- Security reports: sends daily, weekly, or monthly security reports to users.

- Scenario 2: security evaluation for new assets

  Monitors asset changes, detects new assets, and evaluates asset security. Generates security evaluation reports on new assets to help you determine whether to add these assets to your network. The following features are provided:

  - Scans vulnerabilities on servers and web applications.
  - Verifies server configurations.
  - Performs baseline check on cloud services.

- Scenario 3: urgent event handling and cause tracking

  After an urgent event is detected, SOC handles the event and tracks the event cause.

### 6.2.6. On-premises security operations services

On-premises security operations services help you make better use of the features of Apsara Stack products and Apsara Stack Security to ensure your application security.

Security operations services include pre-release security assessment, access control policy management, periodic security check, routine security inspection, and urgent event handling. These services cover the entire lifecycle of your businesses in Apsara Stack. On-premises security operations services help you create a security operations system for cloud businesses. This system enhances the security of application systems and ensures the security and stability of your businesses.

## 6.3. Optional security services

### 6.3.1. DDoS Traffic Scrubbing

Backed by its large-scale and distributed operating system and more than a decade of experience in defending against security attacks, Alibaba Cloud has designed and developed the DDoS Traffic Scrubbing module based on the cloud computing architecture to protect the Apsara Stack platform against large amounts of distributed denial of service (DDoS) attacks.

The following table describes the features provided by the DDoS Traffic Scrubbing module.

| Feature | Description |
| --- | --- |
| Traffic scrubbing against DDoS attacks | Detects and prevents attacks such as SYN flood, ACK flood, ICMP flood, UDP flood, NTP flood, DNS flood, and HTTP flood. |

| Feature | Description |
| --- | --- |
| DDoS attack display | Allows you to view DDoS attacks in the console and search for DDoS attacks by IP address, status, and event information. |
| DDoS traffic analysis | Allows you to monitor and analyze the traffic of a DDoS attack, and view the attack traffic protocol and the top 10 IP addresses that have launched most attacks. |

Note Apsara Stack Security cannot scrub the traffic between internal networks.

**Best practices**

The DDoS Traffic Scrubbing module automatically detects and protects against DDoS attacks targeted at public IP addresses on the Apsara Stack platform. When the platform is subjected to a DDoS attack, the DDoS Traffic Scrubbing module redirects, scrubs, and reinjects network traffic that is detected and scheduled by the Traffic Security Monitoring module. In addition, you can view the detailed information about a DDoS attack event to understand the traffic elements of the attack and analyze the attack source.

# 6.3.2. Cloud Firewall

Cloud Firewall manages north-south traffic in a centralized manner and provides access control and traffic analysis functions. This better protects your network.

Cloud Firewall supports the following features:

| Category | Feature | Description |
| --- | --- | --- |
| Internet Firewall | Access Control | Supports Internet firewalls. You can configure outbound and inbound policies, including the access source, destination type, destination, protocol, port type, port, application, and policy action. |
| | Firewall Switch Policy | Allows you to search for target assets by asset type, region, instance ID, and IP address. You can enable firewall policies, including Internet Firewall, VPC-VPC, and IDC-VPC policies. |
| | Intrusion Prevention Policies | Allows you to set the threat engine mode to the monitoring mode or traffic control mode, to configure an IP address whitelist for outbound and inbound policies, and to customize basic policies for basic protection. You can use the Virtual Patches function and turn on the function at one click.<br>This feature protects your network against abnormal connections, command execution, brute-force cracking, scanning, information leakage, distributed denial of service (DDoS) attacks, overflow attacks, web attacks, backdoors, trojans, worms, mining, and reverse shells. |

| Category | Feature | Description |
| --- | --- | --- |
| | Event Log | Allows you to search for event logs by source IP address, destination IP address, event type, action, and time range. |
| | Traffic Log | Allows you to search for traffic logs by different conditions. |
| VPC Firewall | VPC Firewall | Detects and controls communication traffic between two VPCs. You can configure VPC firewall policies, including the access source, destination type, protocol type, port type, application, and policy action. |
| | Internal Firewall | Controls inbound and outbound traffic between ECS instances. You can configure internal firewall policies, including the access source, destination, protocol type, port range, and action. |
| | IDC-VPC Firewall | You can configure IDC-VPC firewall policies, including the access source, destination type, protocol type, port type, application, and action. |
| | Firewall Switch Policy | Allows you to search for target assets by asset type, region, instance ID, and IP address. You can enable the firewall policies, including Internet Firewall, VPC-VPC, and IDC-VPC policies. |
| | Intrusion Prevention Policies | Allows you to set the threat engine mode to the monitoring mode or traffic control mode, to configure an IP address whitelist for outbound and inbound policies, and to customize basic policies for basic protection. You can use the Virtual Patches function and turn on the function at one click.<br>This feature protects your network against abnormal connections, command execution, brute force cracking, scanning, information leakage, distributed denial of service (DDoS) attacks, overflow attacks, web attacks, backdoors, trojans, worms, mining, and reverse shells. |
| | Event Log | Allows you to search for event logs by source IP address, destination IP address, event type, action, and time range. |
| | Traffic Log | Allows you to search for traffic logs by different conditions. |

## Best practices

Cloud Firewall is applicable to the following scenarios:

- Control the access traffic from the Internet to ECS instances: For example, a financial company on Alibaba Cloud uses IPS to protect their HTTP and other businesses exposed on the Internet.
- Prevent command-and-control activities: For example, a governmental organization on

Alibaba Cloud analyzes not only the access traffic from the Internet to ECS instances but also the command-and-control traffic from ECS instances to the Internet. Based on the analysis, the organization can determine which ECS instances are at risk and block anomalous access in real time to avoid potential risks.

# 6.3.3. Sensitive Data Discovery and Protection

Sensitive Data Discovery and Protection (SDDP) is a data security service used to detect and protect sensitive data in Apsara Stack services.

SDDP uses big data analytics capabilities and artificial intelligence (AI) technologies of Alibaba Cloud to detect and classify sensitive data based on your business requirements. It can also mask sensitive data both in transit and at rest, monitor dataflows, and detect abnormal activities. It provides visible, controllable, and industry-compliant security protection for your sensitive data by using precise detection and analysis. SDDP can detect and protect sensitive data in a variety of Apsara Stack services, such as MaxCompute, Object Storage Service (OSS), and Tablestore.

The following table describes features of SDDP.

| Feature | | Description |
| --- | --- | --- |
| Classification and detection of sensitive data | Detection of new data | A department administrator can authorize SDDP to scan and protect data assets based on business requirements. SDDP only scans and monitors authorized data assets. |
| | Sensitive data classification | SDDP can classify sensitive data detected in big data services, such as MaxCompute, OSS, AnalyticDB, Table Store, and ApsaraDB for RDS. You can define classification rules for sensitive data by using keywords, regular expressions, or other methods. |
| | Sensitive data detection | SDDP has built-in algorithms for detecting sensitive data, and uses file clustering, deep neural networks, and machine learning to detect sensitive images, text, and fields. |
| Management of sensitive data permissions | Asset permissions detection | SDDP can redirect you to pages that display the permissions of data assets and allows you to view the accounts that have permissions to access those assets. The data assets include MaxCompute projects, MaxCompute tables, MaxCompute columns, MaxCompute packages, AnalyticDB databases, AnalyticDB tables, OSS buckets, Table Store instances, and Table Store tables. |
| | Account permissions detection | SDDP allows you to view all accounts in a department and search for departments or accounts in fuzzy search mode. SDDP displays relationships between departments and accounts in a hierarchical and visible layout. |
| | Abnormal permissions usage detection | SDDP automatically detects abnormal permissions usage in big data services, such as MaxCompute, OSS, AnalyticDB, and Table Store. |

| Feature | | Description |
| --- | --- | --- |
| Monitoring of dataflows and operations | Dataflow monitoring | SDDP monitors dataflows among entities, including data storage services (such as MaxCompute, OSS, AnalyticDB, and Table Store), data transmission services (such as Datahub and CDP), the data stream processing service Blink, external databases, and external files. It displays dataflows and abnormal activities on dynamic graphs. This way, you can click an abnormal activity on a graph to redirect to the page for handling the abnormal activity. |
| | Abnormal data operation detection | SDDP detects abnormal operations in big data services, such as MaxCompute, OSS, AnalyticDB, and Table Store. |
| | Abnormal dataflow detection | SDDP detects abnormal dataflows (including abnormal downloads) in big data services, such as MaxCompute, OSS, AnalyticDB, and Table Store. |
| | Detection rule customization | SDDP allows you to customize rules for detecting abnormal dataflows and operations based on algorithms. |
| Abnormal activity processing | Configuration for abnormal activity detection | SDDP allows you to configure thresholds and rules for detecting abnormal activities, such as abnormal dataflows, permissions usage, and data operations. |
| | Abnormal activity processing | SDDP processes abnormal activities with a built-in console. You can search for abnormal activities by department, event type, account, processing status, or time of occurrence. |
| | Abnormal activity statistics | SDDP collects statistics on the processing status of abnormal activities, including abnormal dataflows, permissions usage, and data operations, and then dynamically displays these statistics. |
| Static data masking | Static data masking | SDDP statically masks sensitive data in big data services, such as MaxCompute, OSS, AnalyticDB, Table Store, and ApsaraDB for RDS.<br>It supports the following masking algorithms: hash masking, shield masking, substitution masking, conversion masking, encryption masking, and shuffle masking. |
| Intelligent audit | Intelligent audit | SDDP collects and audits the operation logs of big data services, such as MaxCompute, OSS, and ApsaraDB for RDS. |

## Best practices

- Complies with laws and regulations on personal information protection.

  SDDP detects personal information in large amounts of data, automatically marks risk levels for personal information, and effectively detects data leaks. Enterprises can use SDDP to ensure that their systems comply with laws and regulations on personal information protection.

- Classifies and protects sensitive data of enterprises.

  SDDP classifies and detects sensitive data, manages data permissions, and identifies abnormal activities (such as abnormal dataflows, permission usage, and data operations) based on specified rules. This way, enterprises can properly protect their sensitive data of diverse classifications.

- Handles data leaks.

  SDDP detects abnormal activities based on specific rules and allows you to centrally summarize and handle these activities. This helps enterprises process data leaks online and provides effective support for security O&M.