

Alibaba Cloud

Apsara Stack Enterprise

Introduction to Alibaba Cloud
Apsara Stack

Product Version: V3.18.1

Document Version: 20231027

Legal disclaimer

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

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

Document conventions

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

Table of Contents

1. Alibaba Cloud Apsara Stack	10
2. Overview	11
3. Security and compliance	13
4. Why Apsara Stack	16
5. Architectures	24
6. Scenarios	30
7. Cloud platform services	32
7.1. Unified Cloud Management Platform	32
7.1.1. Features	32
7.1.2. Benefits	34
7.1.3. Scenarios	35
7.2. Apsara Stack Resilience	36
7.2.1. Product details	36
7.2.2. Benefits	43
7.2.3. Scenarios	45
8. Computing services	48
8.1. ECS	48
8.1.1. ECS	48
8.1.2. Benefits	52
8.1.3. Scenarios	53
8.2. Auto Scaling	54
8.2.1. Features	54
8.2.2. Benefits	56
8.2.3. Scenarios	57
8.3. ROS	58
8.3.1. Features	58

8.3.2. Benefits	59
8.3.3. Scenarios	59
8.4. OOS	59
8.4.1. Service details	60
8.4.2. Benefits	60
8.4.3. Scenarios	61
8.5. What is Container Service?	62
8.5.1. Features	62
8.5.2. Benefits	63
8.5.3. Scenarios	65
8.6. Container Registry	69
8.6.1. Features	69
8.6.2. Benefits	70
8.6.3. Scenarios	70
9.Storage services	72
9.1. CDS	72
9.1.1. Features	72
9.1.2. Benefits	83
9.1.3. Scenarios	87
9.2. Simple Log Service	90
9.2.1. Features	91
9.2.2. Benefits	92
9.2.3. Scenarios	93
10.Network services	95
10.1. What is SLB?	95
10.1.1. Features	96
10.1.2. High availability	97
10.1.3. Scenarios	99

10.2. Virtual private cloud	100
10.2.1. What is a VPC?	100
10.2.2. Benefits	102
10.2.3. Use scenarios	102
10.3. EIP overview	104
10.3.1. Features	105
10.3.2. Benefits	106
10.3.3. Scenarios	106
10.4. Express Connect	106
10.4.1. What is Express Connect?	106
10.4.2. Benefits	107
10.4.3. Scenarios	108
10.5. NAT Gateway	109
10.5.1. What is NAT Gateway?	109
10.5.2. Benefits	110
10.5.3. Use scenarios	111
10.6. IPv6 Gateway	111
10.6.1. What is an IPv6 Gateway?	112
10.6.2. Terms	113
10.6.3. Common scenarios	114
10.7. Cloud Gateway	116
10.7.1. Service description	116
10.7.2. Benefits	116
10.8. Apsara Stack DNS	117
10.8.1. Service description	117
10.8.2. Benefits	119
10.8.3. Scenarios	121
11.Database services	123

11.1. ApsaraDB RDS	123
11.1.1. What is ApsaraDB RDS?	123
11.1.2. Benefits	123
11.1.3. Scenarios	124
11.2. PolarDB	124
11.2.1. Features	124
11.2.2. Benefits	125
11.2.3. Scenarios	125
11.3. ApsaraDB for MongoDB	126
11.3.1. What is ApsaraDB for MongoDB?	126
11.3.2. Benefits	127
11.3.3. Scenarios	129
11.4. KVStore for Redis	130
11.4.1. Product details	130
11.4.2. Benefits	133
11.4.3. Scenarios	134
11.5. What is DTS?	135
11.5.1. Overview	136
11.5.2. Benefits	137
11.5.3. Scenarios	138
11.6. What is DMS?	143
11.6.1. Features	144
11.6.2. Benefits	147
11.6.3. Scenarios	148
11.7. DBS	149
11.7.1. Features	150
11.7.2. Benefits	150
11.7.3. Scenarios	153

12. Middleware services	155
12.1. What is Message Queue for Apache RocketMQ?	155
12.1.1. Features	155
12.1.2. Benefits	156
12.1.3. Scenarios	157
13. Big data services	159
13.1. Apsara Big Data Manager	159
13.1.1. Features	159
13.1.2. Benefits	160
13.1.3. Scenarios	160
13.2. What is MaxCompute?	161
13.2.1. Features	161
13.2.2. Benefits	162
13.2.3. Scenarios	164
13.3. DataWorks	166
13.3.1. Features	166
13.3.2. Benefits	173
13.3.3. Scenarios	174
13.4. DataHub	175
13.4.1. Features	175
13.4.2. Benefits	176
13.4.3. Scenarios	176
14. Security services	178
14.1. Apsara Stack Security	178
14.1.1. What is Apsara Stack Security?	178
14.1.2. Benefits	186
14.1.3. Scenarios	188
14.2. Key Management Service	189

14.2.1. Service description	189
14.2.2. Benefits	190
14.2.3. Scenarios	191
15.Application Services	192
15.1. API Gateway	192
15.1.1. Features	192
15.1.2. Benefits	193
15.1.3. Scenarios	194

1. Alibaba Cloud Apsara Stack

Alibaba Cloud Apsara Stack is an open, unified, and trusted cloud solution that is tailored for enterprise customers. Apsara Stack is developed based on the same distributed architecture as Alibaba Cloud. Apsara Stack allows enterprises to deploy public cloud services in on-premises environments and scale the cloud services to the public cloud with a few clicks.

Apsara Stack provides the following editions: Apsara Stack Enterprise and Apsara Stack Agility. Apsara Stack Enterprise provides stable, secure, all-in-one, lightweight, and easy-to-integrate capabilities that allow enterprises of different business scales to adapt to various business scenarios. Apsara Stack Agility is designed to meet requirements in big data scenarios, database-related scenarios, middleware scenarios, and storage scenarios.

Benefits

Apsara Stack is an extension of the public cloud and brings public cloud technologies to private clouds. Apsara Stack helps enterprises deliver complete and customizable Alibaba Cloud software solutions and allows enterprises to deploy hyper-scale cloud computing services and big data services of the public cloud in on-premises environments. Apsara Stack provides a consistent hybrid cloud experience for enterprises and allows enterprises to obtain IT resources based on their business requirements to ensure business continuity.

Apsara Stack supports on-premises deployment, and can be independently run and managed outside Alibaba Cloud.

Apsara Stack helps public service sectors and enterprises digitally transform their business systems and services based on the combination of diverse services, digitalization practices, and mature solutions of Alibaba Group. Apsara Stack provides the following benefits:

- **On-demand provisioning:** Apsara Stack supports hyper-scale clusters that have more than 10,000 servers in each region and supports the multi-region deployment to meet the management requirements of large business systems and applications. Apsara Stack also provides small-scale cloud platforms for enterprises that are in the early phase of cloud migration to reduce cloud migration costs.
- **Comprehensive capabilities:** Based on the same technologies that are used in the public cloud, Apsara Stack supports more than 80 public cloud services and hot upgrades for the services. This way, you can use the Alibaba Cloud services in on-premises environments. Apsara Stack also allows you to scale the services to the public cloud with a few clicks to provide hybrid cloud solutions in various business scenarios.
- **Rich industry experience:** The Apsara Stack team has rich experience in constructing full-stack enterprise-level clouds for sectors such as public service and finance and is capable of ensuring security during cloud migration of large and medium-sized enterprises.
- **Security and stability:** Apsara Stack uses a native security architecture to provide multi-layered, integrated security protection. Apsara Stack is the first platform in China to pass the MLPS 2.0 Level 4 certification. Apsara Stack has also passed multiple security certifications such as Trusted Cloud Services (TRUCS) certification, ISO 27001, General Data Protection Regulation (GDPR), and Commercial Cryptography Application Security Assessment. Apsara Stack provides disaster recovery solutions like three data centers across two regions disaster recovery, and Multi-Site High Availability (MSHA). Such solutions meet regulatory requirements in the finance industry and helps ensure high reliability and business continuity.

2.Overview

Definition

Private clouds are divided into the following types based on the enterprise scale or business requirements:

- Multi-tenant comprehensive private clouds for industries and large groups: full-stack cloud systems that are created in a top-down manner. The system is designed to drive hyper-scale digital applications and meet IT requirements such as the continuous integration and development of DevOps applications and the operations support for production environments.
- Single-tenant basic private clouds for small and medium-sized enterprises and scenarios: cloud systems that can support local computing tasks and host technical systems such as large-scale Software as a Service (SaaS) applications, industry clouds, and large group clouds.

Industry trends

Cloud computing has become a major driving force for digital transformation. Cloud computing is widely used in various industries and sectors, such as public services, finance, and energy. An increasing number of enterprises are seeking for their opportunities of business transformation in the trend of digital transformation aided by cloud migration. During this process, the following new requirements arise to enterprises: reduce costs and increase efficiency, precisely manage and maintain business to meet complex requirements, and innovate and coordinate applications to drive the data intelligence upgrade of enterprise business.

Digital transformation is accelerated, and IT environments become more complex in recent years. A single public cloud provides similar services, which cannot meet security and compliance requirements in special industries. In addition, a single private cloud requires high costs for IT infrastructure and has the bottleneck of scaling computing resources, which cannot respond to changing business requirements of enterprises. The preceding issues give rise to hybrid clouds. A secure, controllable, elastic, and open hybrid cloud combines a private cloud and a public cloud to accelerate digital transformation.

Gartner predicts that 90% of enterprises will adopt hybrid cloud solutions to manage infrastructure in the future. Enterprises deploy both a public cloud and a private cloud to meet complex business requirements that cannot be met by a single cloud. Hybrid cloud solutions that are developed based on private cloud services allow enterprises to seamlessly integrate multi-cloud systems with Infrastructure as a Service (IaaS) or Platform as a Service (PaaS). The hybrid cloud solutions simplify the migration of core application systems of enterprises to the cloud, and the migration of workloads of traditional enterprises such as those in public service sectors. This way, enterprises can focus on their business development, which brings more benefits to enterprises. International Data Corporation (IDC) also predicts that more than 50% of enterprise applications in the Chinese mainland will be deployed in containerized hybrid cloud environments or multi-cloud environments by 2023 to provide an agile and seamless deployment and management experience.

Hybrid clouds will become the first choice for enterprises to deploy their applications in the future.

Apsara Stack

As the earliest practitioner in the native-cloud field, Apsara Stack has become the largest native hybrid cloud platform that can be used in most industry scenarios. Apsara Stack provides the following core competitiveness:

- **Native and fully self-developed:** Apsara Stack is developed by using the cloud architecture of Alibaba Cloud public cloud. This allows Apsara Stack to leverage the long-term accumulation and innovation in public cloud to provide you with a consistent experience in resource self-service, development, and management.
- **Large-scale commercial use:** Apsara Stack is available for commercial use on a large scale and has been adopted by a large number of enterprises in various sectors such as public service sectors. As a prominent cloud platform in Asia Pacific, the public cloud can ensure optimal service stability for you.
- **Trusted computing power:** Apsara Stack is a full-stack cloud platform whose computing power has been verified by Double 11 for many years. Hybrid Cloud has continuously made improvements in the IaaS layer, PaaS layer, and Desktop as a Service (DaaS) layer. Apsara Stack also optimizes more than 80 cloud services, such as Artificial Intelligence (AI) services, services based on Internet of things (IoT), and DingTalk, to help you achieve rapid business innovation.

Benefits

- **Stable:** Apsara Stack provides a financial-grade disaster recovery solution that is used by millions of users in the public cloud and is continuously optimized to ensure high system availability and business continuity.
- **Secure:** Apsara Stack adopts a cloud native security architecture to provide you with a hierarchical security system. Apsara Stack is the first cloud platform that passed the security assessment of commercial cryptography applications in the Chinese mainland. Alibaba Cloud is the first cloud service provider who received the MLPS 2.0-Level 4 Certification.
- **Open:** Apsara Stack provides standard and compatible APIs, can be deployed by using hardware from different vendors, and allows you to deploy more than 70 cloud services, such as AI, middleware, and big data services, to data centers in a flexible manner.
- **Intelligent:** Apsara Stack provides an enterprise-level cloud management portal, implements centralized cloud asset management to support Apsara Uni-manager Dashboards and automated O&M, and uses mature AI algorithms that are extensively verified.

3. Security and compliance

Apsara Stack devotes itself to data security and user privacy. Alibaba Cloud is committed to building private cloud platforms which are featured by openness and security. Technological innovation, ever-growing computing power, and economies of scale enables transformation of cloud computing into infrastructure.

Apsara Stack is dedicated to providing users with stable, reliable, secure, and compliant cloud computing infrastructure services to help protect the availability, confidentiality and integrity of users' systems and data.

Qualifications

The security protection mechanisms used by Alibaba Cloud are recognized by authorities from in and outside the Chinese mainland. Alibaba Cloud integrates years of experience in protection against Internet security attacks into the security protection features of the cloud platform. Alibaba Cloud also integrates many compliance standards into the internal management and product design of the cloud platform. In addition, Alibaba Cloud also proactively participates in formulating standards related to the cloud platform and shares best practices.

Alibaba Cloud has been qualified by more than ten organizations at home and abroad, and is the most recognized cloud service provider in Asia.

Alibaba Cloud qualifications

Qualification	Description
ISO 27001	An international certification system for information security management. This certifies the commitment of the Alibaba Cloud platform to provide data security, network security, communication security, and operation security.
CSA STAR	An international certification system for cloud security management. Alibaba Cloud is the first cloud service provider that has won the gold CSA STAR certification in the world.
ISO 20000	A certification system for IT service management. This certifies that Alibaba Cloud has established standard service processes and implemented service standardization to enhance efficiency and reduce IT risks.
ISO 22301	A certification system for business continuity management. This certifies that Alibaba Cloud has rolled out business continuity plans, disaster recovery deployment, and regular disaster recovery drills to improve the stability of the cloud platform.
MLPS 2.0 level 4	Apsara Stack delivers security capabilities set forth in Multi-Level Protection Scheme (MLPS) 2.0 which is formulated in accordance with Baseline for Classified Protection of Cybersecurity (GB/T 22239-2019).
Cloud computing service certified by MIIT	The national laboratory accreditation for cloud products is the only product-level certification based on national standards.
Service Organization Controls (SOC) audits	Alibaba Cloud has passed the SOC 1 Type II, SOC 2 Type II and SOC 3 audits.

Chinese qualifications

Qualification	Organization
---------------	--------------

ITSS cloud computing service (level 1 for private cloud IaaS service)	China Electronics Standardization Association
Trusted Cloud Services - User data protection certification (private cloud)	China Academy of Information and Communications Technology
Assessment report for classified protection of information systems (level 4 for private cloud)	Information Security Rating Center of the Ministry of Public Security
Assessment report for classified protection of Apsara Stack V3.0	Information Security Rating Center of the Ministry of Public Security
Assessment report for security of Apsara Stack Agility Insight	China Academy of Information and Communications Technology
Cloud assessment certification - Cloud computing reference architecture - Cloud solutions	China Electronics Standardization Institute
Trusted Cloud Services - Open-source solution (Apsara Stack Agility Edition) /Virtualization and virtualization management software	China Academy of Information and Communications Technology

Compliance

Alibaba Cloud continues improving management and mechanisms based on industrial standards and best practices. Alibaba Cloud has passed a bunch of standard certifications, third-party audits, and self-assessments to demonstrate its compliance practices.

Alibaba Cloud compliance practices are divided into the following parts depending on perspectives, industries, and geological areas:

Management compliance

These compliance certifications exhibit Alibaba Cloud's mature management mechanisms and compliance with industrial best practices:

- ISO 27001: information security management.
- ISO 20000: IT service management
- ISO 22301: business continuity management
- CSA STAR: a maturity model for cloud security management
- MLPS 2.0 level 4
- CNAS cloud computing testing

Systematic compliance reports

These compliance certifications demonstrate the integrity and effectiveness of Apsara Stack management, such as whether system management effectiveness is consistent, whether separation of duties is appropriate, and whether O&M audits are perfect.

SOC 1 Type II and SOC 2 Type II: SOC reports are presented by independent third-parties to confirm that the effectiveness of Alibaba Cloud key compliance controls and objectives is consistent. These reports help users and their auditors understand measures for operation control and compliance. Alibaba Cloud has obtained three types of SOC reports:

- SOC 1 Type II: internal control for financial statements
- SOC 2 Type II: security, availability, and confidentiality
- SOC 3: security, availability, and confidentiality

Apsara Stack MLPS 2.0 Compliance White Paper

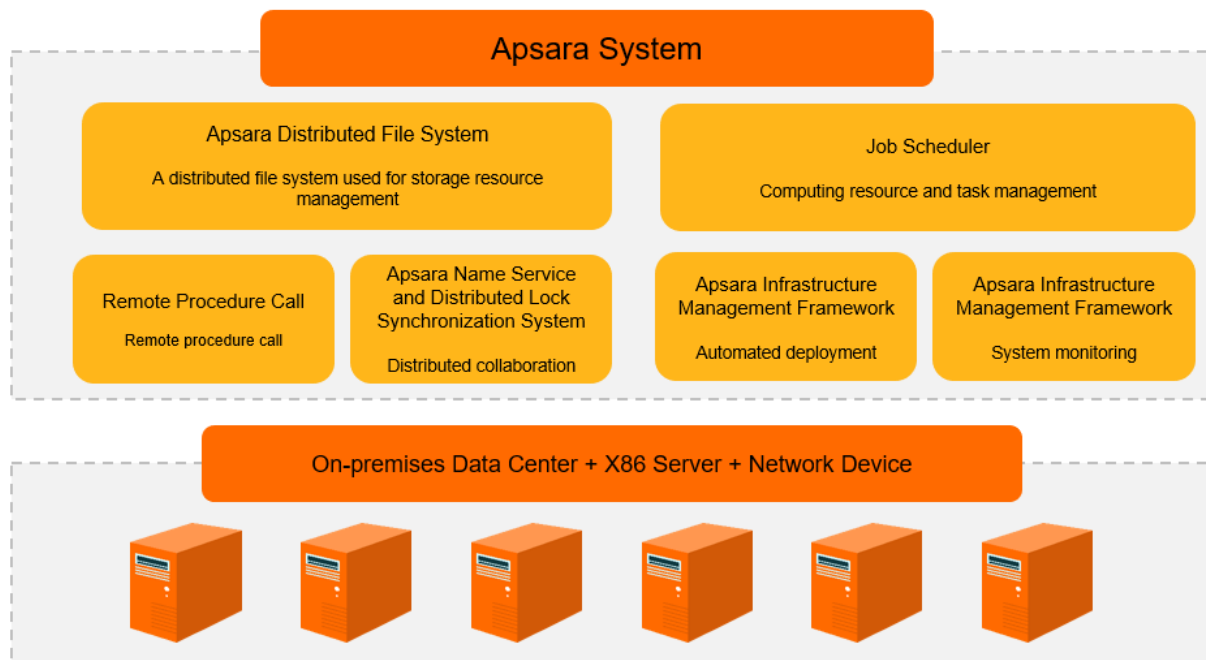
In response to the requirements of MLPS 2.0 and in accordance with Baseline for Classified Protection of Cybersecurity, the Information Security Rating Center of the Ministry of Public Security and Alibaba Cloud Computing Co., Ltd. have jointly formulated and promulgated the Apsara Stack MLPS 2.0 Compliance White Paper. This white paper elaborates on MLPS 2.0 technical architecture, the compliance status of Apsara Stack, and suggestions on white paper usage. The white paper can customers to quickly acquire compliance capabilities for Apsara Stack in multiple delivery scenarios. Such capabilities can be combined with customer protection measures such as applications, security management, and physical environment to build overall security protection systems that meet the requirements of both MLPS 2.0 and customers.

4. Why Apsara Stack

Hyper-scale distributed cloud operating system

Both Apsara Stack and Alibaba Cloud public cloud are based on Apsara Distributed Operating System. Apsara Distributed Operating System provides underlying services such as storage, computing, and scheduling for upper-layer services.

Apsara Distributed Operating System is a hyper-scale universal operating system developed by Alibaba Cloud. Apsara Distributed Operating System connects millions of servers around the world to act as a supercomputer and provides computing capabilities such as online public services. The computing capabilities provided by Apsara Distributed Operating System are powerful, universal, and accessible to everyone.



? Note

- The ARM architecture supports servers that use Kunpeng 920, Phytium FT-2000+, and Phytium 2500 chips.
- The X86 architecture supports servers that use Intel and Hygon chips.

Apsara Distributed Operating System consists of the following modules:

- **Underlying services for distributed systems:** This module provides the coordination, remote procedure call (RPC), security management, and resource management services needed in a distributed environment. These services provide support for upper-layer modules such as the distributed file system and task scheduling module.
- **Distributed file system:** This module provides a reliable and scalable service capable of storing large amounts of data. The distributed file system aggregates the storage capabilities of each node in a cluster and automatically protects against hardware and software faults to provide uninterrupted access to data. This module also supports incremental scaling and automatic data load balancing. An API similar to Portable Operating System Interface of UNIX (POSIX) is provided to access files in the user space. Additionally, the module supports random read/write and append write operations.

- **Task scheduling:** This module schedules tasks in the cluster system and supports both online services that rely on a quick response speed and offline tasks that require high data processing throughput. The module can automatically detect faults and hot spots in the system. The module ensures stable and reliable service operations by using methods such as error retry and concurrent backup for long-tail operations.
- **Cluster monitoring and deployment:** This module monitors the status of clusters as well as the status and performance metrics of upper-layer application services, and generates alerts and records of exception events. Additionally, the module provides O&M engineers with deployment and configuration management of the entire Apsara system and its upper-layer applications. The module supports both the online elastic scaling of clusters and the online upgrade of application services.

Unified deployment and control system

Apsara Infrastructure Management provides cloud services with basic support capabilities such as centralized deployment, authentication, authorization, and control.

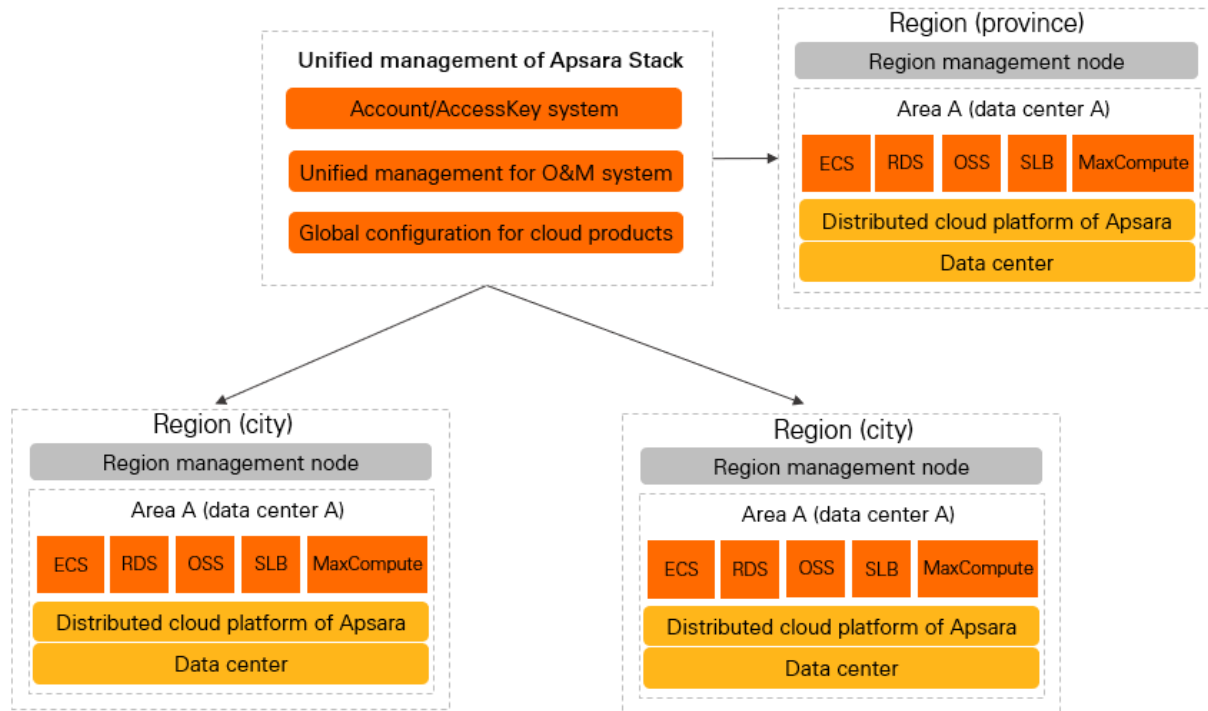
Apsara Infrastructure Management includes modules such as deployment framework, resource library, metadata database, authentication and authorization, API Gateway, Log Service, and control service.

- The deployment framework provides an access platform for centralized service deployment and manages service dependencies.
- The resource library stores the executable files of all cloud services and components on which the cloud services depend.
- The authentication and authorization module provides access control capabilities for cloud services and supports multi-tenant isolation.
- API Gateway provides a centralized API management platform for cloud services.
- Log Service provides log storage, retrieval, and access for cloud services.
- The control service module monitors the basic health status of each cloud service and supports the Apsara Stack O&M system.

Centralized cross-region management of multiple data centers

Apsara Stack implements centralized management for O&M, operations, and metering in each region.

Express Connect supports cross-region data access and sharing between two VPCs. This transforms cloud computing into a public utility such as water, electricity, and coal, and benefits everyone. The Apsara Stack management system that converges resource pools can centralize and monitor the computing, storage, and network resources as well as their usage of multiple data centers. This system provides centralized capabilities, including resource management, resource deployment, O&M management, service management, and self-services.



Highly reliable disaster recovery solutions

Apsara Stack provides a variety of solutions for disaster recovery (DR), such as Apsara Stack Resilience for Zone-disaster Recovery, Apsara Stack Resilience for Geo-disaster Recovery, Apsara Stack Resilience for Backup and Recovery (ASR-BR), hybrid networking of multi-region deployment and DR, and DR within three data centers across two regions.

A DR system includes two or more systems that provide the same features in distant locations. These systems mutually monitor health status and switch features. If one system stops due to an unexpected incident such as a fire, flood, earthquake, or vandalism, services can be failed over to a system in a different location to ensure business continuity.

Apsara Stack DR solutions are designed and developed based on the cloud computing capabilities of Alibaba Cloud. These solutions comply with common international DR standards. When the network conditions meet the design requirements, the Apsara Stack platform implements the active-active mode on the network access and user application layers and the active-standby mode on the data persistence layer.

Apsara Stack Resilience for Geo-disaster Recovery

Apsara Stack Resilience for Geo-disaster Recovery use the active-standby mode, in which the resources of both primary and secondary data centers are available to users. Protected resources such as ApsaraDB RDS instances and Object Storage Service (OSS) buckets and their rules are equally distributed between the primary and secondary data centers. Protection groups are created for applications to implement DR.

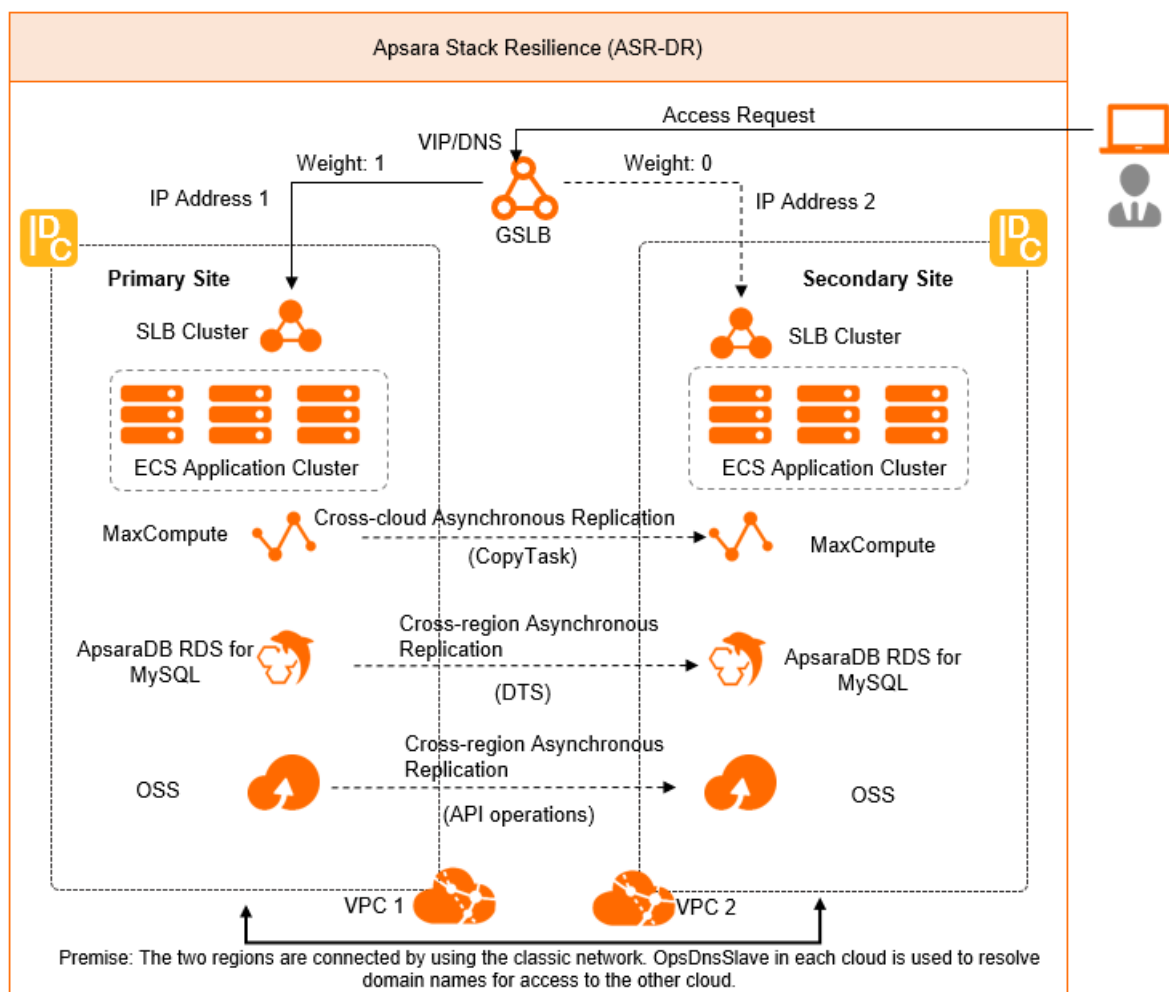
Apsara Stack Resilience for Geo-disaster Recovery supports the following scenarios: cross-cloud DR and cross-region DR.

- **Cross-cloud DR**

Cross-cloud DR is implemented between two cloud instances. The primary and secondary sites are two independent cloud instances deployed in different locations and use independent account systems. Users must use separately authorized accounts to log on to the two cloud instances.

- **Cross-region DR**

Cross-region DR is implemented between two regions. The primary and secondary sites are two regions of a single cloud instance but use the same account system. DR from a central region to a general region and from a general region to a general region are supported in the cross-region DR scenario.

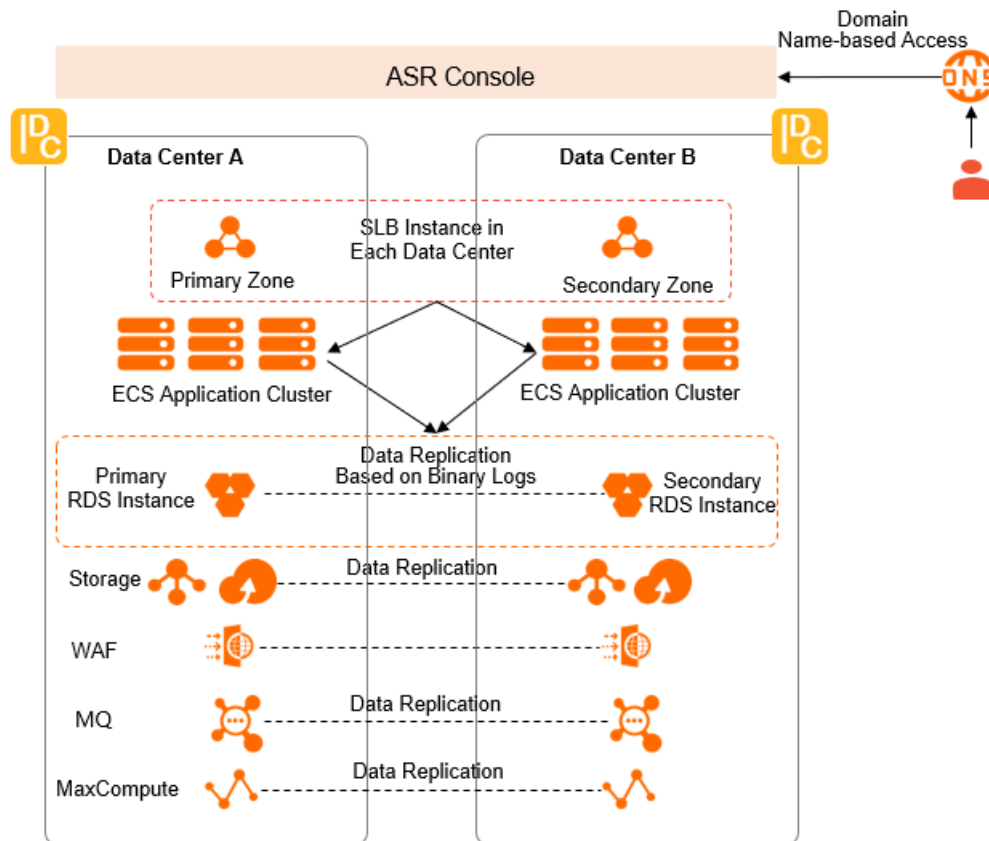


Apsara Stack Resilience for Zone-disaster Recovery

Apsara Stack Resilience for Zone-disaster Recovery refers to two independent, mutually backed-up data centers within the same region. If an exception occurs in the primary data center, the secondary data center takes over services by using Apsara Stack Resilience for Zone-disaster Recovery.

- **Zone-disaster recovery (two data centers)**

You can use domain names to access cloud services deployed in the primary and secondary data centers. The domain names of cloud services do not change if services are failed over to the secondary data center. You do not need to remodel your applications. This simplifies application development, makes cloud services easy to use, and allows you to focus on business development.



Note

Data Center A is the primary data center. Data Center B is the secondary data center.

- **Zone-disaster recovery (three data centers)**

You can add a third data center based on the architecture of the two data centers-based mode and deploy distributed databases to ensure zero data loss and zero recovery point objective (RPO) in the finance industry. In the architecture of the three data centers-based zone-disaster recovery, the DR mechanism for services in the active-standby or active-active mode remains unchanged. Failovers are implemented based on the policies used in the architecture of two data centers-based zone-disaster recovery.

Apsara Stack Resilience for Backup and Recovery

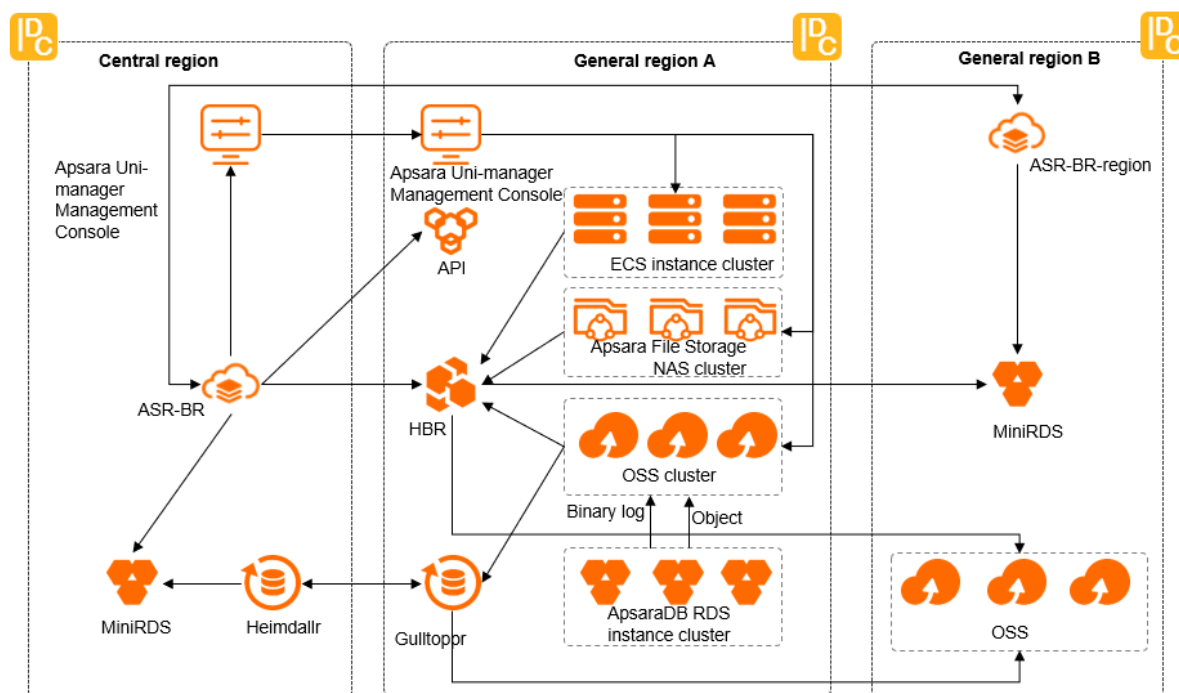
ASR-BR supports the following two architecture modes: local backup and cross-region backup.

- **Local backup**

Data of the primary site is backed up to an OSS-compliant storage system on a regular basis. You must deploy ASR-BR and related services at the primary site. This architecture reduces deployment costs.

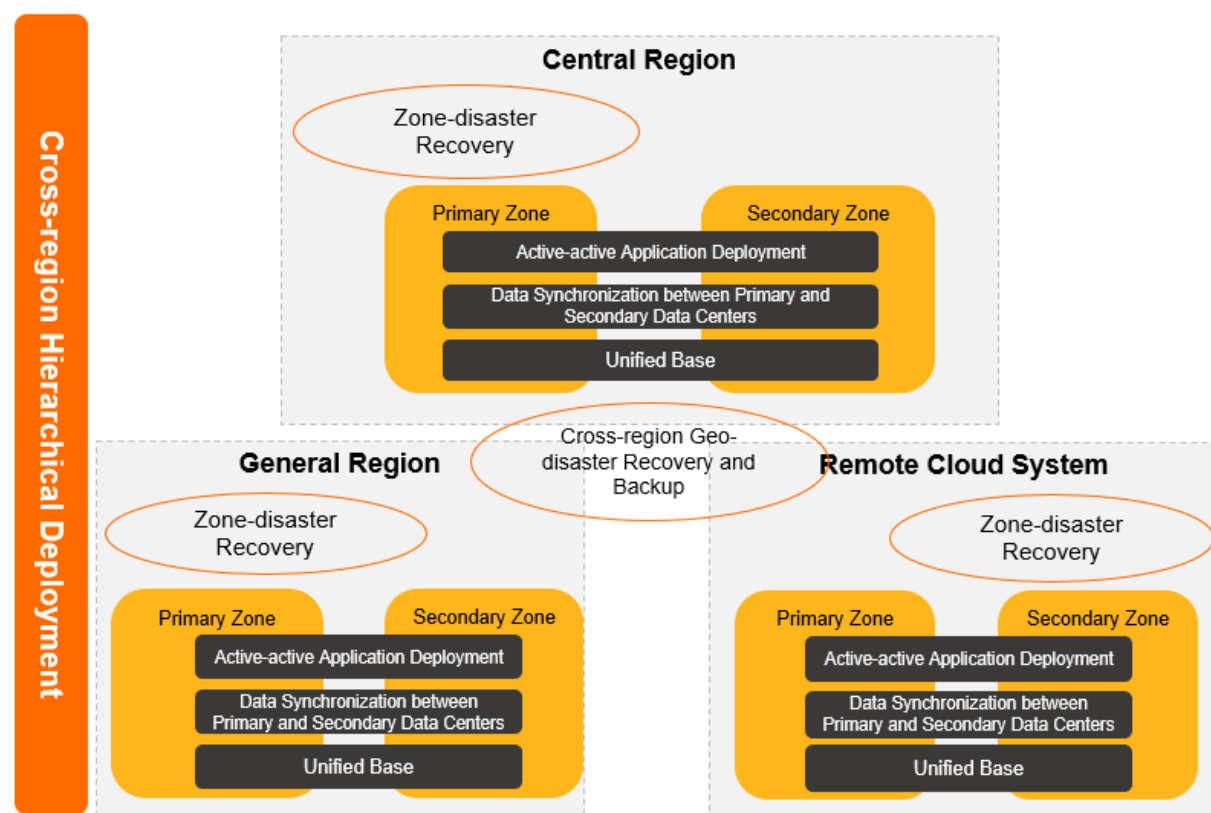
- **Cross-region backup**

In the multi-region architecture, the data of a region is backed up to another region.



Hybrid networking of multi-region deployment and DR

Apsara Stack Enterprise supports Apsara Stack Resilience for Zone-disaster Recovery, cross-region DR, and cross-region backup in the multi-region scenario. This can satisfy the DR requirements of a variety of industries.

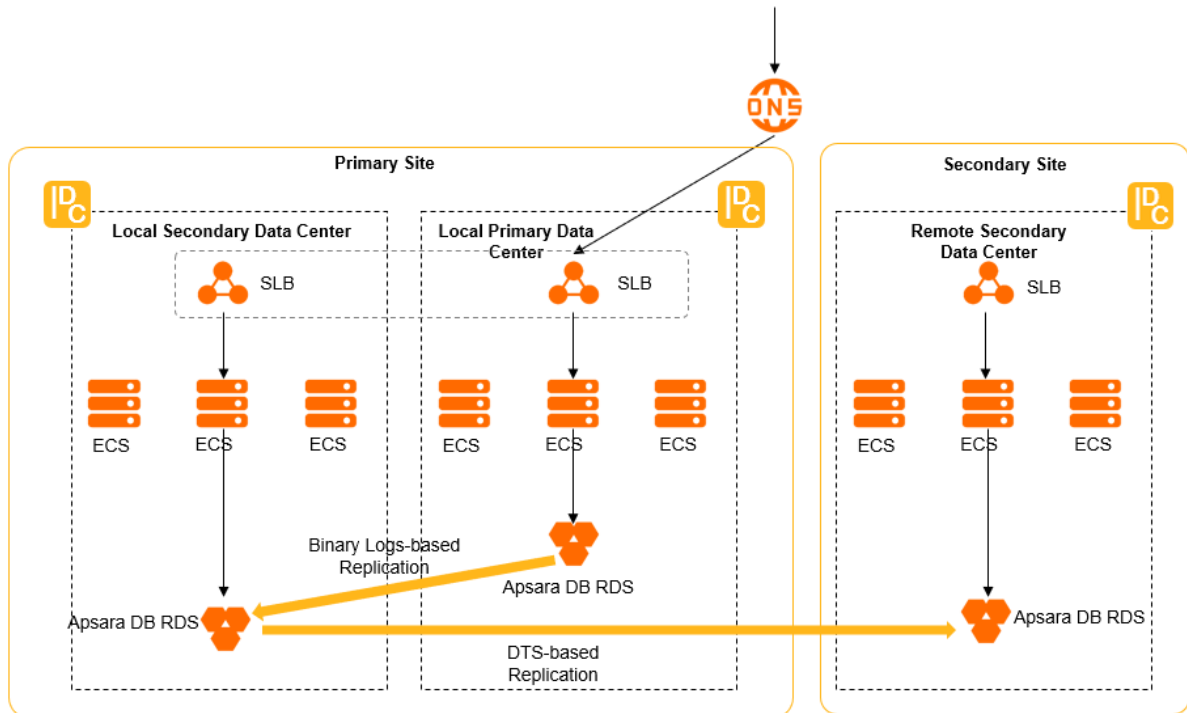


Three data centers across two regions

In the DR solution of three data centers across two regions, the three data centers include the local primary data center, the local secondary data center, and the remote secondary data center. This solution can provide high disaster backup capabilities. Cloud services supported by the DR solution include ApsaraDB RDS and OSS.

- DR solution for ApsaraDB RDS deployed in three data centers across two regions

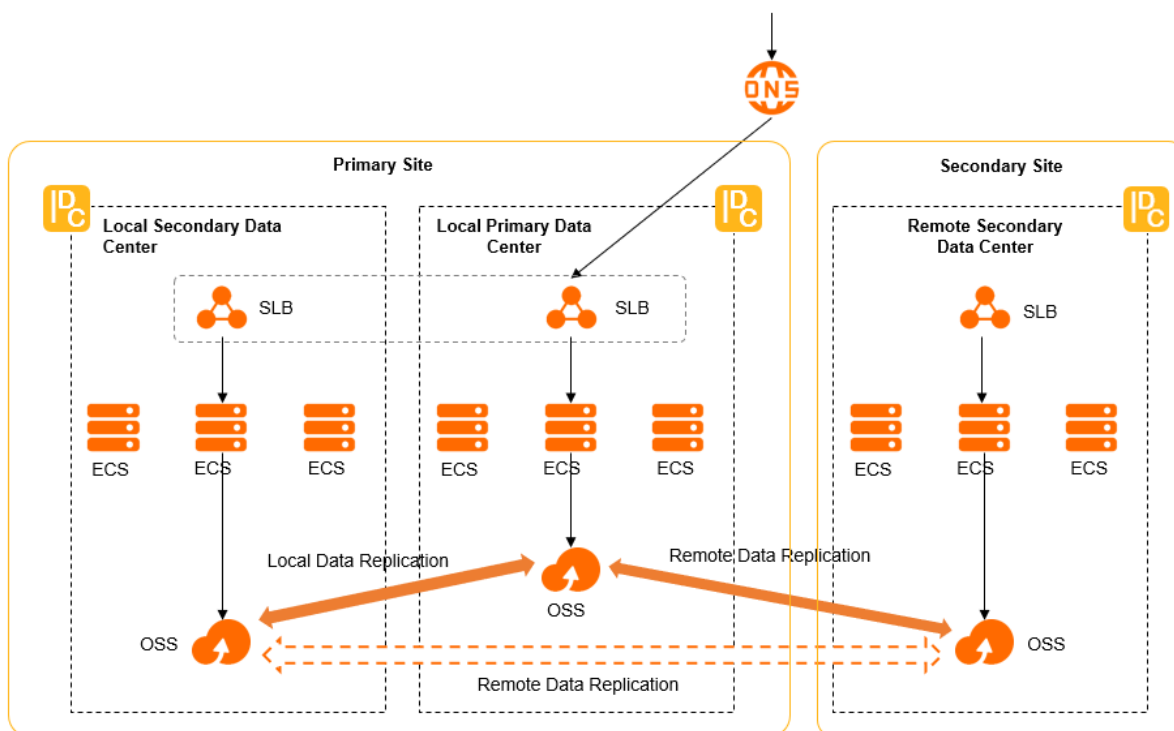
ApsaraDB RDS is independently deployed at each of the primary and secondary sites. Two data centers are deployed at the primary site. ApsaraDB RDS is independently deployed in each of the two data centers. The primary and secondary ApsaraDB RDS instances deployed at the primary site are used to implement Apsara Stack Resilience for Zone-disaster Recovery. The primary and secondary sites are used to implement Apsara Stack Resilience for Geo-disaster Recovery. DR is implemented between two cloud instances.



- DR solution for OSS deployed in three data centers across two regions

OSS is independently deployed at each of the primary and secondary sites. Two data centers are deployed at the primary site. OSS is independently deployed in each of the two data centers. The two data centers at the primary site are used to implement Apsara Stack Resilience for Zone-disaster Recovery. DR is implemented between two data centers. The primary and secondary sites are used to implement Apsara Stack Resilience for Geo-disaster Recovery. DR is implemented between two cloud instances.

If the primary data center at the primary site fails, the BDR administrator initiates a failover plan in the Apsara Stack Resilience for Zone-disaster Recovery console. The primary OSS bucket is failed over to the local secondary data center with a few clicks to ensure business continuity. If both data centers at the primary site fail, the BDR administrator initiates a failover plan in the Apsara Stack Resilience for Geo-disaster Recovery console. Then, OSS protection groups are failed over to the secondary site to ensure business continuity. After the primary site is recovered, a reverse data replication tunnel is created between the primary bucket at the primary site and the secondary bucket at the secondary site. This way, incremental data is synchronized from the secondary OSS bucket at the secondary site to the primary bucket at the primary site. After the data is synchronized, start a failback plan to fail back the traffic to the primary site. The DR solution of three data centers across two regions enhances the business continuity of customer systems.



Unified Hybrid Cloud Management Platform: Apsara Uni-manager

Apsara Uni-manager is an enterprise-level cloud management platform provided by Apsara Stack. Apsara Uni-manager can be used in Apsara Stack and hybrid cloud scenarios. Apsara Uni-manager supports provisioning, operations, and management of cloud resources. Apsara Uni-manager provides core capabilities such as centralized management, intelligent O&M, fine-grained operations, and custom extensions. Apsara Uni-manager simplifies hybrid cloud management, improves user experience, and helps enterprises accelerate digital transformation.

- **Unified portals:** Apsara Uni-manager provides a unified entry that consists of a self-managed portal, an operations portal, an O&M portal, and a data portal. This delivers a comprehensive range of cloud management capabilities for users in different businesses.
- **Unified services:** Apsara Uni-manager provides unified service capabilities, including unified users, unified permissions, unified data, and unified process management.
- **Openness, ease of integration, and scalability:** Apsara Uni-manager furnishes multi-cloud management capabilities and open API gateways. Third-party data and webpages are collected by using northbound APIs and delivered to multi-cloud environments for integration by using southbound APIs.

OpenAPI

Apsara Stack provides a wide range of SDKs and RESTful APIs on the OpenAPI platform. OpenAPI provides flexible access to a variety of Apsara Stack services. You can also use OpenAPI to obtain the basic control information of Apsara Stack and integrate Apsara Stack with your centralized control system.

5. Architectures

Apsara Stack adopts a cloud native architecture and is built on the operating system, distributed technologies, and products that are developed by Alibaba Cloud. The architecture supports all cloud services and allows the complete openness of the cloud platform. Apsara Stack comes with comprehensive service features for enterprises, delivers disaster recovery and backup capabilities, and can be fully self-managed.

System architecture

The system architecture of Apsara Stack consists of the following layers:

- **Physical device layer:** includes hardware devices for cloud computing, such as servers and network devices.
- **Basic service layer:** provides the basic service capabilities, including out-of-band management, system cloning, clock source, YUM source, metadatabase, and platform log service.
- **Converged management and control layer:** manages various cloud services on Apsara Stack.
- **Cloud service and API layer:** provides centralized management and O&M features for VMs and physical machines by using a converged node management mechanism, and an open API platform for centralized API management and custom development.
- **Centralized management layer:** provides a unified layer for centralized O&M.

Apsara Stack is a full-stack solution that ensures the stability of the architecture, the reliability of the cloud platform, and business continuity.

Logical architecture

Apsara Stack virtualizes the computing capabilities and storage capabilities of physical servers and network devices to provide virtual computing, distributed storage, and software-defined networks on which ApsaraDB RDS, distributed middleware services, and big data services can run. This allows Apsara Stack to provide the underlying IT services for the business applications of enterprises, and to be integrated into their existing accounts and monitoring and O&M systems.

The logical architecture of Apsara Stack has the following characteristics:

- The hardware infrastructure consists of servers and network devices. Only x86 servers are supported.
- The Apsara kernel of Apsara Distributed Operating System provides kernel services for all cloud services based on Apsara Distributed Operating System.
- All cloud services use the same API framework, security system, and O&M system. The O&M system is used for account management, authorization, monitoring, and logging.

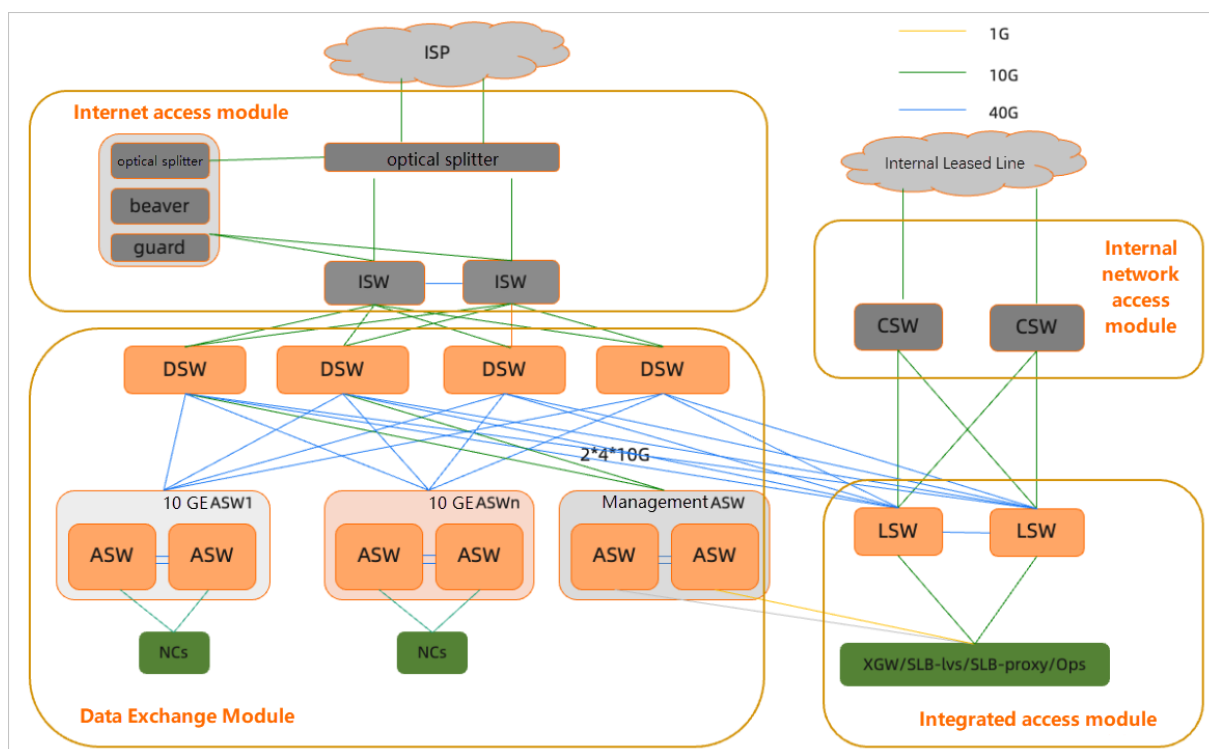
Network architecture

Apsara Stack adopts a flat Layer 2 Clos network architecture. The network architecture isolates the service plane from the out-of-band management plane, and supports linear scaling and load sharing of switches.

The network architecture of Apsara Stack consists of four modules: internal network access module, Internet access module, data exchange module, and integrated access module.

- **Internal network access module:** connects self-managed network resources to cloud resources and allows users to access virtual private clouds (VPC) or regular cloud services.
- **Internet access module:** connects to the networks of Internet service providers (ISPs) or the backbone networks of enterprises. A business service area establishes a connection over the Internet or to other data centers by using this module.

- **Data exchange module:** allows access to all cloud service servers. The internal traffic among the servers is exchanged in this module.
- **Integrated access module:** allows access to various basic services and cloud services, such as Server Load Balancer (SLB) and VPC.



Roles and usage of a vSwitch in an access module

Role name	Section	Description
Inter-connection Switch (ISW)	Internet access module	An ISW is an egress switch that provides access to the networks of ISPs or the backbone networks of enterprises.
Customer Switch (CSW)	Internal network access module	A CSW facilitates access to the internal backbone networks of enterprises, including access to VPCs by using leased lines, and performs route distribution and interactions between the internal network and the Internet.
Distribution Switch (DSW)	Data exchange module	A DSW is a core switch that provides access to access switches.

Access Switch (ASW)	Data exchange module	An ASW provides access to cloud servers and uses an uplink to connect to a DSU.
LVS Switch (LSW)	Integrated access module	An LSW provides access to cloud services such as VPC and SLB.

Internal network access module

The internal network access module consists of two CSWs to allow internal users to access VPCs and common cloud services.

- **Access to VPCs:** The CSWs route traffic of internal users to VPCs by mapping the internal users to the VPCs. User groups on a CSW are isolated from each other.
- **Access to regular cloud services:** The CSWs are connected to the integrated access module over Exterior Border Gateway Protocol (eBGP) and allow direct access to all resources in the business service area.

The leased line access solution of VPC allows enterprises to manage their virtual networks. For example, enterprises can select their own CIDR blocks and configure route tables and gateways. Enterprises can also connect their VPC to a traditional data center by using leased lines or VPN connections to create a custom network environment. This helps ensure the smooth migration of applications to the cloud.

Internet access module

The Internet access module consists of two ISWs. This module allows access to the networks of ISPs or the public backbone networks of enterprises and performs route advertisements and interactions between the internal network and the Internet.

Two ISWs back up routes to each other over the Interior Border Gateway Protocol (IBGP). The ISWs are uplinked with the networks of ISPs or the public backbone networks of enterprises by using static routes or eBGP based on business requirements. The Internet access bandwidth is determined based on the network scale on Apsara Stack and backbone network bandwidth of each enterprise.

To improve network reliability, we recommend that you connect two ISWs to multiple ISPs and establish network connections between ISWs and ISPs over BGP. In addition, make sure that your network is connected to two 10 GE networks of each ISP. The Internet access module uses eBGP to exchange routes with the data exchange module. The Internet access module advertises relevant Internet routes to the data exchange module and receives internal cloud service routes from the data exchange module. This way, the internal network and the Internet can communicate.

The network security module of Apsara Stack Security is mounted in the bypass of the Internet access module. Traffic that is transmitted from the Internet to cloud networks is routed by an optical splitter to Network Detection and Response of Apsara Stack Security. If Network Detection and Response detects malicious traffic, it advertises the relevant routes to route the malicious traffic to the network security module of Apsara Stack Security for scrubbing. Then, the scrubbed traffic is injected back into the Internet access module.

Data exchange module

The data exchange module has a typical Layer 2 Clos architecture that consists of DSUs and ASWs.

Each ASW pair is used to create a stack as a leaf node. This node can select suitable data exchange models based on network scales. All service servers of Apsara Stack are connected to devices on the ASW stacks over uplinks. The ASWs are connected to the DSWs over eBGP. The DSWs are isolated from each other. The data exchange module is interconnected with other modules over eBGP.

The data exchange module receives the Internet routes that are advertised by ISWs in the Internet access module and advertises the CIDR blocks of cloud services to the ISWs.

Integrated access module

The integrated access module consists of the servers of various cloud services, such as extendable gateways (XGWs) on Cloud Network Management, Layer 4 load balancers or Layer 7 load balancers, and OPS servers. These servers are connected to two LSWs and exchange routing information over Open Shortest Path First (OSPF). Two LSWs exchange routing information over iBGP. LSWs exchange routing information with DSWs in the data exchange module and CSWs in the internal network access module over eBGP.

Security architecture

Apsara Stack provides comprehensive security capabilities based on underlying communication protocols for upper-layer applications to ensure the security of user access and user data.

Apsara Stack provides a comprehensive role-based authorization mechanism to allow you to manage access to resources in the multi-tenant mode and ensure the security of resources. Apsara Stack provides multiple security roles including security administrators, system administrators, and security auditors to meet the requirements of multi-level security in O&M scenarios.

Apsara Stack V3.0 and later editions use Apsara Stack Security to provide hierarchical and integrated cloud security protection.

SM algorithms

SSL certificates are required to access all consoles of Apsara Stack. SSL certificates that use internationally accepted algorithms (RSA algorithms) and SM algorithms (Chinese cryptographic algorithms) are installed on servers. Key Management Service (KMS) in Apsara Stack is adapted to support SM algorithms. Hardware security modules (HSM) are used in Apsara Stack.

Key cloud services such as Elastic Block Storage (EBS), Object Storage Service (OSS), MaxCompute, and ApsaraDB RDS support data encryption by using SM algorithms. Data Encryption Service of Apsara Stack Security supports encryption and decryption by using SM algorithms.

Base components

The Apsara Stack base consists of three components that provide support for the deployment and O&M of the cloud platform.

Component classification	Component name	Description
--------------------------	----------------	-------------

OPS components	Yum	<p>The component is the software repository that contains the software packages that you want to install.</p> <p>The software repository is deployed in the initial installation stage to install the operating system and deploy application packages such as the Apsara system and Elastic Compute Service (ECS), and dependent modules of Apsara Stack on physical servers.</p>
	Clone	The machine cloning service.
	NTP	<p>The clock source service.</p> <p>NTP is deployed on Apsara Stack physical servers to synchronize the time from the standard NTP clock source to other hosts.</p>
	DNS	<p>The domain name resolution service.</p> <p>The service supports regular resolution and reverse resolution for domain names in Apsara Stack. You can run a bind instance on each of the two OPS servers and use the keepalived component to provide high-availability services. In this case, if one OPS server is faulty, services can be switched over to the other OPS server.</p>
Base middleware	dubbo	The distributed remote procedure call (RPC) service.
	tair	The caching service.
	mq	The message queuing service.
	ZooKeeper	The distributed coordination service.
	Diamond	The configuration management service.
	SchedulerX	The cron job service.
	Apsara Infrastructure Management	The management service in the data center.
	TianjiMon	The monitoring service in the data center.
	OTS-inner	The table storage service.

Basic components of the base	SLS-inner	The log service on the cloud platform.
	mini-RDS	The metadatabase.
	POP	The Apsara Stack OpenAPI platform.
	OAM	The account system.
	RAM	The authentication and authorization system.
	WebApps	The service that provides support for the Apsara Uni-manager Operations Console.

6.Scenarios

Apsara Stack provides flexible and scalable industrial solutions for customers of different scales and sectors.

Apsara Stack can create customized solutions based on the unique business traits of different sectors such as industry, agriculture, transportation, government, finance, and education to provide users with end-to-end products and services. This topic describes the following scenarios.

City Brain

Urban management is a field that involves one of the largest volumes of data in China. This marks the transition of governmental information from a closed-flow model to an open-flow online model. Urban data has a greater value as it has more time and larger space to flow. Cloud computing becomes urban infrastructure. Data becomes a new means of production and strategic resources. AI technology becomes the nerve center of a smart city. All of these together form the City Data Brain.

City Brain has the following values and features:

- A breakthrough of urban governance mode. City Brain uses urban data as resources to improve government management capabilities, resolve prominent issues of urban governance, and implement an intelligent, intensive, and humane form of governance.
- A breakthrough of urban service mode. City Brain provides more accurate and convenient services for enterprises and individuals, makes urban public services more efficient, and saves more public resources.
- A breakthrough of urban industrial development. City Brain lays down an industrial AI layout, takes open urban data as an important fundamental resource, drives the development of industries, and promotes the transformation and upgrade of traditional industries.

Alibaba Finance Cloud

Alibaba Finance Cloud is an industrial cloud that serves financial organizations, such as banks, security agencies, insurance companies, and finance. It relies on a cluster of independent data centers to provide cloud products that meet the regulatory requirements of the People's Bank of China, China Banking Regulatory Commission (CBRC), China Securities Regulatory Commission (CSRC), and China Insurance Regulatory Commission (CIRC). It also provides more professional and comprehensive services for financial customers. Enterprises can build Alibaba Finance Cloud independently or with Alibaba Cloud. Alibaba Finance Cloud meets the requirements of large and medium-sized financial organizations for independent cloud data centers that are completely physically isolated. It can also provide cloud computing and big data platforms for data centers of customers.

Alibaba Finance Cloud has the following values and features:

- Independent resource clusters
- Stricter data center management
- Better disaster recovery capability
- Stricter requirements for network security isolation
- Stricter access control
- Compliance with the security supervision requirements and compliance requirements of banks
- Dedicated security operation, compliance, and solution teams of the Alibaba Finance Cloud sector
- Dedicated account managers and cloud architects of Alibaba Finance Cloud

- Stricter user access mechanism

7. Cloud platform services

7.1. Unified Cloud Management Platform

The unified cloud management platform Apsara Uni-manager is an enterprise-class cloud management platform for Apsara Stack and hybrid cloud scenarios. Apsara Uni-manager supports provisioning, operations, and management of cloud resources. Apsara Uni-manager provides core capabilities such as centralized management, intelligent O&M, fine-grained operations, and tailored features. Apsara Uni-manager simplifies hybrid cloud management, improves user experience, and helps enterprises accelerate digital transformation.

Apsara Uni-manager consists of the following components:

- **Apsara Uni-manager Management Console:** provides an integrated management portal that delivers capacities such as fine-grained resource governance, intelligent data analysis, and tailored features. This helps enterprises minimize cloud management costs.
- **Apsara Uni-manager Operations Console:** provides a unified operations portal that delivers common O&M capabilities, such as alerting and monitoring, inspection management, and resource management. The portal also provides service O&M capabilities for computing, network, and storage. These capabilities help minimize environment maintenance costs and stabilize environments.
- **Apsara Uni-manager Dashboards:** provides intuitive data dashboards for multi-dimensional and panoramic data presentations, including overall status and resource usage of hybrid clouds. You can design different dashboards on the homepage for different roles.
- **Multi-cloud management platform:** manages cloud resources across heterogeneous cloud platforms such as Apsara Stack and Alibaba Cloud in a unified manner. The multi-cloud management platform supports open standards for cloud service integration and allows you to integrate services from heterogeneous cloud platforms and manage the resources, organizations, users, and permissions of the services. This way, you can manage cloud services from heterogeneous cloud platforms through a unified portal in a standard management platform.

7.1.1. Features

Apsara Uni-manager provides various features, such as resource management, personnel management, permission management, operations center, security center, and application center. The features help simplify the management and deployment of physical and virtual resources, improve resource usage, and reduce operations costs. Apsara Uni-manager provides features such as general O&M, service O&M, security compliance, and system configuration to simplify daily O&M operations and improve the efficiency of O&M. Apsara Uni-manager also provides predefined dashboards to adapt to typical business scenarios, allows you to create custom dashboards, and displays business data from different dimensions.

Operations management

- **Resource management**

You can activate and monitor resources. After resource sets are managed and resource metrics are configured, operations personnel can easily understand the usage of each resource and prevent risks at the earliest opportunity when exceptions occur.

- **Personnel management**

You can manage organizations, users, and user groups in your enterprise to control organizations, permissions, and user groups in a centralized manner. This helps improve the management efficiency by administering the full lifecycle of users in different scenarios and satisfying the needs of different users to access systems and resources.

- **Permission management**

You can configure general roles, RAM roles, data permissions, and access control policies for the system and services to improve system security.

- **Operations center**

You can manage quotas, metering and billing rules, statistical analysis methods, and bills to control resource usage and resource billing in a centralized manner. The operations center helps you quickly understand resource usage and billing information. The operations center also provides flexible control features to meet requirements during daily operations.

- **Security center**

You can use operation logs, multi-factor authentication (MFA), and AccessKey pairs to reinforce security when users access the system and resources.

O&M management

- **General O&M**

You can manage alerts, inspections, resources, capacities, and changes, and archive backup files to complete general O&M.

- **Service O&M**

The following service O&M features are provided: computing O&M, network O&M, storage O&M, database O&M, middleware O&M, big data O&M, cloud platform O&M, security services, and application services.

- **Security compliance**

You can audit operation logs and access matrix of cloud services and manage server passwords, AccessKey pairs, and platform encryption.

- **System configuration**

You can configure the system. For example, you can manage user permissions, configure the platform, and manage O&M APIs.

Dashboards for management

- **Predefined dashboards**

The following predefined dashboards are provided to display the status of the platform and resource usage in real time: resource dashboard, organization dashboard, security dashboard, network dashboard, and alerting dashboard.

- **Custom dashboards**

You can use the online dashboard editor to create custom dashboards based on templates, specify styles for the dashboards, and specify data that you want to view on the dashboards.

Multi-cloud management

- **Multi-cloud access**

You can create cloud platforms. You can create cloud platform instances, including Apsara Stack platforms of different editions and public cloud platforms. Apsara Stack editions such as Apsara Stack Enterprise, Apsara Stack Agility, Apsara Stack Insight, DBStack, CNStack, and ZStack are available. You can also create instances for heterogeneous cloud platforms. In addition, you can go to the page for managing third-party adapters that are connected to Apsara Uni-manager in a centralized manner.

- **Cloud service management**

You can classify, define, and manage cloud services on different cloud platforms in a centralized manner. This way, you can implement unified product classification. Service catalogs of cloud services on each cloud platform are also provided. This helps integrate service capabilities on different cloud platforms.

- **Centralized management of multi-cloud resources**

You can use the resource collection feature to collect resources and display the resources in a list on the multi-cloud management platform to improve the efficiency of centralized management.

- **User and organization management**

You can manage organizations, users, and user groups.

- **Permission management**

Role-based access control is provided. You manage roles, manage permissions of roles, and query management operations on the permissions of roles. Management operations include create, modify, delete, and query operations.

You can grant permissions based on business features, cloud platforms, cloud services, organizations, and projects. If you grant permissions on a project, permissions are independently granted to the users and user groups within the project. You can also grant permissions to an organization, user group, project, or user.

- **Project management**

A project mainly involves user groups, cloud instances, and cloud resources. You can use projects to manage users and cloud services that belong to multiple cloud platforms. You can also use projects to manage various resources on cloud platforms, and grant and manage fine-grained permissions to users by project.

- **Usage statistics on multiple cloud platforms**

You can manage usage statistics of resources on multiple cloud platforms in a centralized manner. After the usage statistics of resources on multiple cloud platforms are collected and displayed, you can view and search for the statistics of resources.

- **Audit log**

You can use the security log feature to view the operation logs of the multi-cloud management platform, manage the logs of cloud services on the cloud platform, and implement quick log retrieval.

- **Message management**

You can create a custom message template to define the message content for different management statuses. You can also configure a variety of collaborative software to send messages to users.

7.1.2. Benefits

Apsara Uni-manager helps you quickly build business systems, improve resource usage, and reduce management costs. Apsara Uni-manager also delivers various capabilities, such as proactive alerting and monitoring, root cause locating, and automatic troubleshooting by using automated O&M processes. The capabilities help minimize environment maintenance costs and ensure environment stability.

Unified portals to maximize user experience

- Unified portals are provided to enable centralized management and flexible scheduling of hybrid cloud and multi-cloud resources.
- The simple self-service experience is consistent with that on the public cloud.

- The integrated management capability is achieved from application perspectives.

Flexible permissions to facilitate control

- Multiple predefined roles can be used to perform operations such as operations management, resource usage viewing, resource monitoring, and security management.
- Custom roles can be created to define shared permissions, managed resources, application permissions, and menu permissions in a flexible manner.
- RAM authentication and permission management methods that are consistent with the permissions in the public cloud are provided.

Intelligent analysis to facilitate management

- Real-time update of global data and unified resource scheduling are implemented.
- Resource usage trends are monitored and analyzed, resource configurations are optimized, and resource usage is improved.
- Detailed resource usage statistics and bills are provided, the value of resources is intuitive, and the stability of service operations is ensured.

Open and simple integration

- Intuitive API portals are provided to reduce learning costs and improve development efficiency.
- Standard northbound APIs and SDKs of multiple languages are provided.
- Page-level integration and custom configurations can be used to facilitate interconnections with third-party platforms.

Centralized O&M to ensure stability

- Centralized monitoring on hybrid cloud resources, inventory, and alerts is supported to check operating conditions and identify risks.
- Automatic and real-time updates of the inter-resource dependency topology and predefined algorithms can be used to locate and analyze root causes, narrow down the scope for troubleshooting, and accelerate error demarcation.
- An O&M script platform and visualized orchestration capabilities are provided. Automatic troubleshooting methods for a large number of scenarios are also provided to reduce manual intervention.
- Inventory AI algorithms and dynamic analysis are shared with the public cloud to automatically calculate global optimal scaling policies, implement cost optimization, and reduce resource waste.

Dashboards for comprehensive management

- Standard preset dashboards and flexible custom dashboards are provided to meet requirements in different business scenarios.
- Various data sources of JSON, ASAPI, ASAPI data pool, HTTP, JSONP, and Excel are supported for comprehensive O&M data.
- Rich visualization components and data sources allow you to create custom dashboards.
- Data presentations in bar charts, pie charts, dashboards, and maps improve the efficiency of O&M and help you quickly obtain information.

7.1.3. Scenarios

Apsara Uni-manager is suitable for various scenarios, such as hybrid cloud management, multi-level cloud management, industry cloud operations, and cloud O&M management.

Hybrid cloud management

If you want to deploy Apsara Stack, Apsara Uni-manager can provide tenant-side capabilities such as resource distribution, permission management, and metering and billing management based on enterprise organization models. This improves resource usage, helps control access to resources to ensure security, and improves the efficiency of resource management. If you have deployed Apsara Stack and have purchased or plan to purchase a large number of public cloud resources of Alibaba Cloud, you must centrally manage hybrid cloud resources to provision resources to users in a centralized manner.

Multi-level cloud management

Apsara Stack is deployed at multiple levels and generally includes the headquarters level and regional level. The regional level uses the autonomous management method and is managed by the headquarters level in a centralized manner.

O&M personnel can perform O&M operations and manage multi-level clouds at the headquarters in a centralized manner. Regular users can deploy their services on different Apsara Stack platforms and allocate resources to the platforms based on their business requirements to support business growth.

Industry cloud operations

If you want to build cloud operations platforms, you can perform personalized configurations and process provisioning on the Apsara Uni-manager Management Console to deliver cloud services and operations to your users.

Cloud O&M management

You can manage your cloud resources in a comprehensive and multi-dimensional manner in the following scenarios: daily O&M, automated O&M, security O&M, and remote O&M.

7.2. Apsara Stack Resilience

As informatization proliferates in industries, the stability and security of IT infrastructure becomes more crucial to enterprises. The enforcement of MLPS 2.0 drives rapid growth of new infrastructure. Cloud disaster recovery has a great future. To ensure business continuity and data security for enterprises, Apsara Stack provides multiple disaster recovery solutions such as Apsara Stack Resilience for Zone-disaster Recovery, Apsara Stack Resilience for Geo-disaster Recovery, and Apsara Stack Resilience for Backup and Recovery (ASR-BR). Combinations of these solutions are also supported. Disaster recovery solutions are designed and developed based on Alibaba Cloud computing capabilities and international disaster recovery standards.

Apsara Stack Resilience (ASR) is a tool with a graphical interface that is used to implement fast failover in the event of a disaster and to minimize RTO. One-click disaster recovery control from application system perspectives makes ASR an ideal choice for enterprises to protect business continuity and data security.

7.2.1. Product details

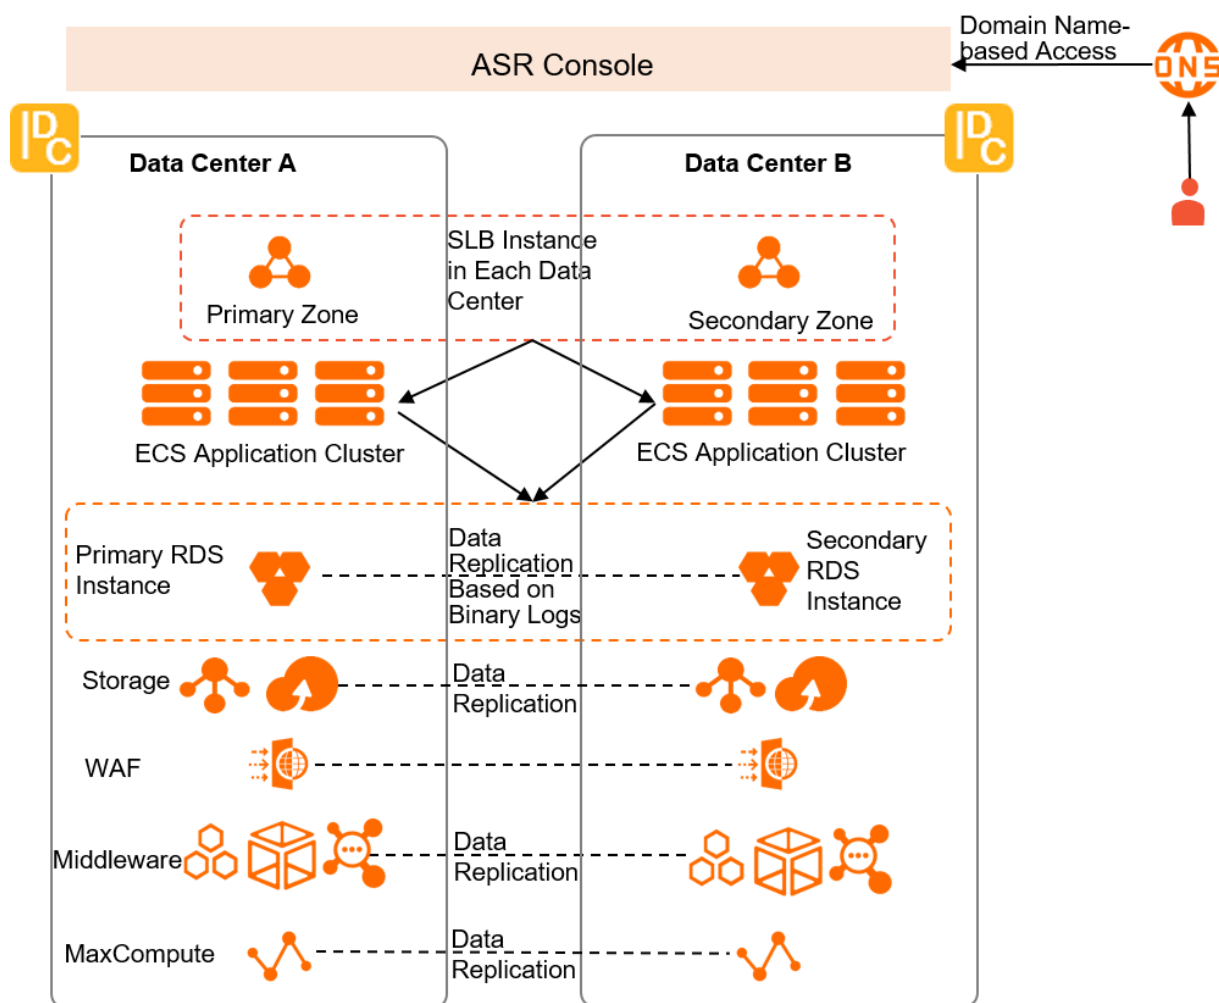
Apsara Stack Resilience is dedicated to full-stack cloud backup and recovery and is used to ensure data security and business continuity for enterprises. Apsara Stack Resilience provides the following disaster recovery solutions: Apsara Stack Resilience for Zone-disaster Recovery, Apsara Stack Resilience for Geo-disaster Recovery, and Apsara Stack Resilience for Backup and Recovery (ASR-BR). You can combine these solutions to fit your business scenarios. For example, the three data centers across two regions architecture combines Apsara Stack Resilience for Zone-disaster Recovery and Apsara Stack Resilience for Geo-disaster Recovery. It aims to quickly implement switchover and recovery in multiple scenarios.

Apsara Stack Resilience for Zone-disaster Recovery

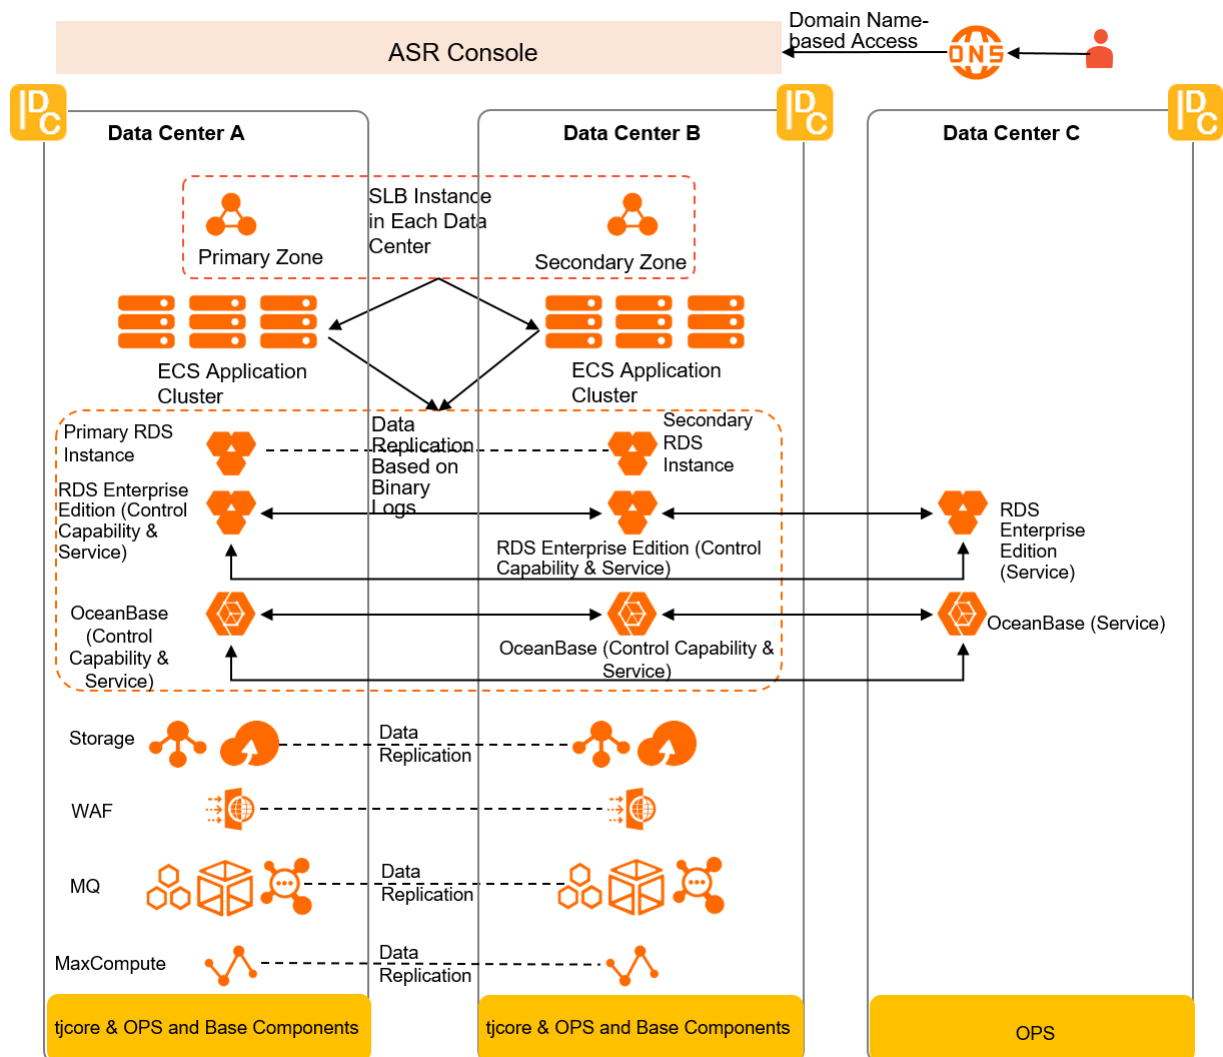
Apsara Stack Resilience for Zone-disaster Recovery consists of a primary data center and a secondary data center. The secondary data center is in the same region as the primary center, but in different zones. This deployment method is used to combat local disasters with a small impact. Apsara Stack Resilience for Zone-disaster Recovery is adopted on all core cloud products such as networking, cloud computing, databases, storage, middleware, and big data products. This can be used in cross-zone deployment for high availability. In the event of a disaster, business can be quickly switched over to the secondary zone. The endpoint of the cloud product remains unchanged.

Apsara Stack Resilience for Zone-disaster Recovery can be used in two data centers and three data centers. On the foundation of the two data centers-based mode, a third data center can be added and the databases are deployed in a distributed manner.

The architecture of Apsara Stack Resilience for Zone-disaster Recovery with two data centers:



The architecture of Apsara Stack Resilience for Zone-disaster Recovery with three data centers:



Apsara Stack Resilience for Zone-disaster Recovery is mainly used for monitoring, drills, protection group disaster recovery, data center-level fault recovery, and single-product fault recovery.

- Cloud product monitoring: checks the status of synchronization tasks between data centers, business status of each data center, and synchronization latency on a regular basis, and monitors disaster recovery in both data centers in real time on the dashboard.
- Disaster recovery drills: allows you to customize hot switchover for cloud products to meet your requirements of routine disaster recovery drills.
- Data center-level fault recovery: handles data center-level faults, such as power outages in the primary and secondary data centers, network outages in the primary and secondary data centers, and split-brain scenarios.
- Single-product fault recovery: helps the cloud product recover from faults and ensures business continuity when the high-reliability feature becomes unavailable.

Compared with other solutions, Apsara Stack Resilience for Zone-disaster Recovery has the following advantages:

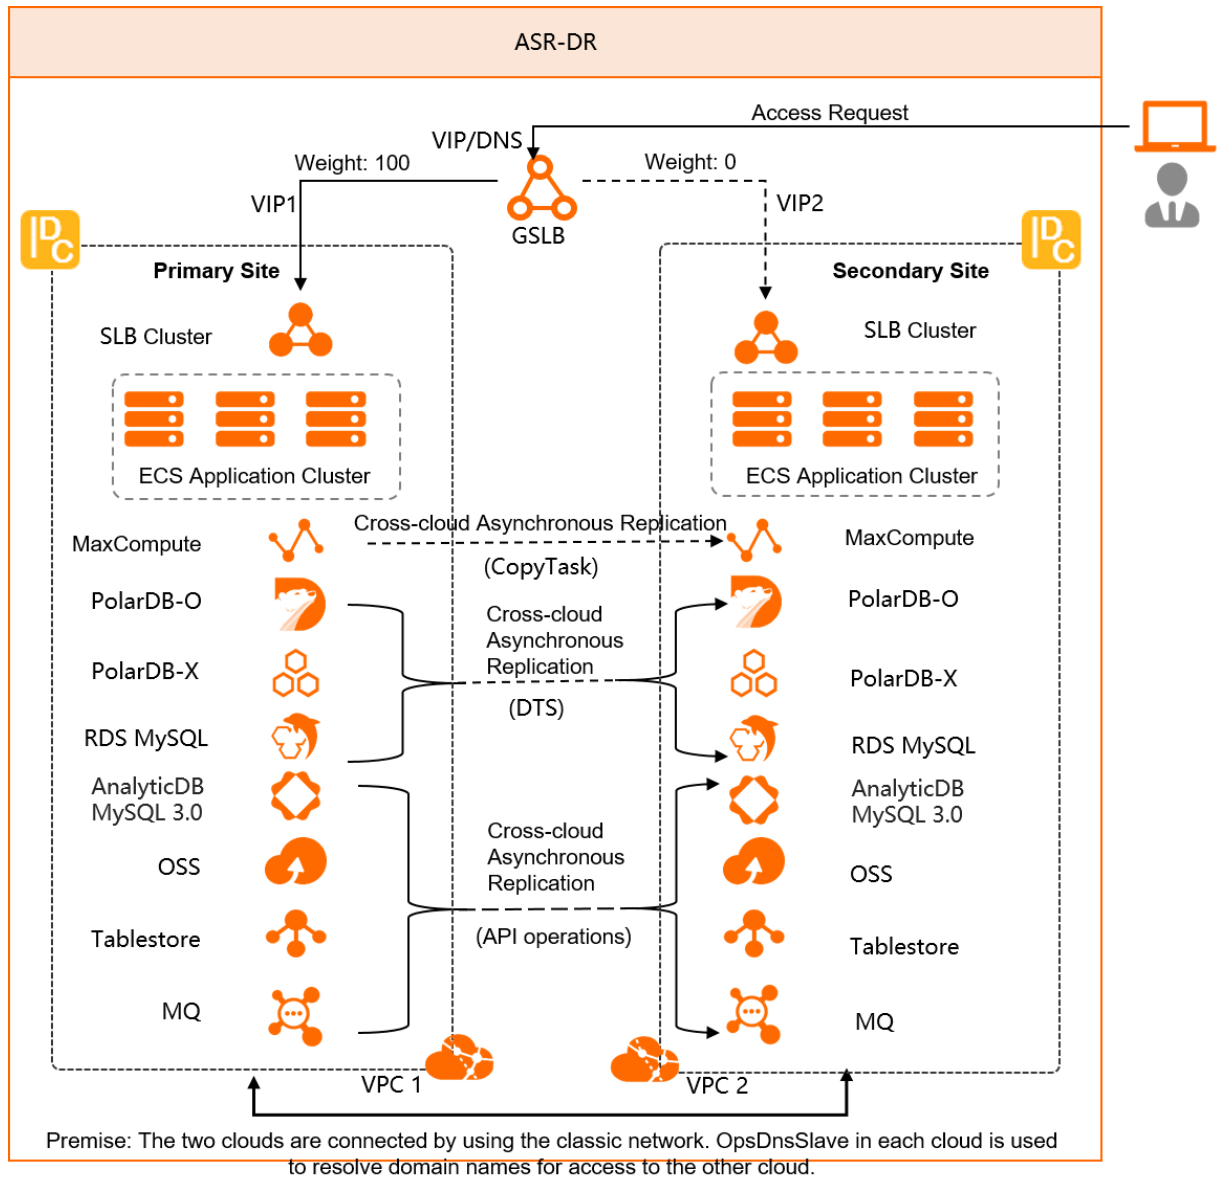
- It is cost-effective to access applications because Apsara Stack Resilience for Zone-disaster Recovery is oriented to cloud-native applications and no business transformation is required for application migration to the cloud.

- Full-stack disaster recovery provides integrated disaster recovery architecture to cover Apsara Uni-manager and cloud products.
- Apsara Stack Resilience for Zone-disaster Recovery with three data centers supports 0 synchronization latency for some cloud products (ApsaraDB RDS for MySQL, ApsaraDB for OceanBase, and Elasticsearch). This meets the strong consistency requirements of business data in the financial cloud industry.

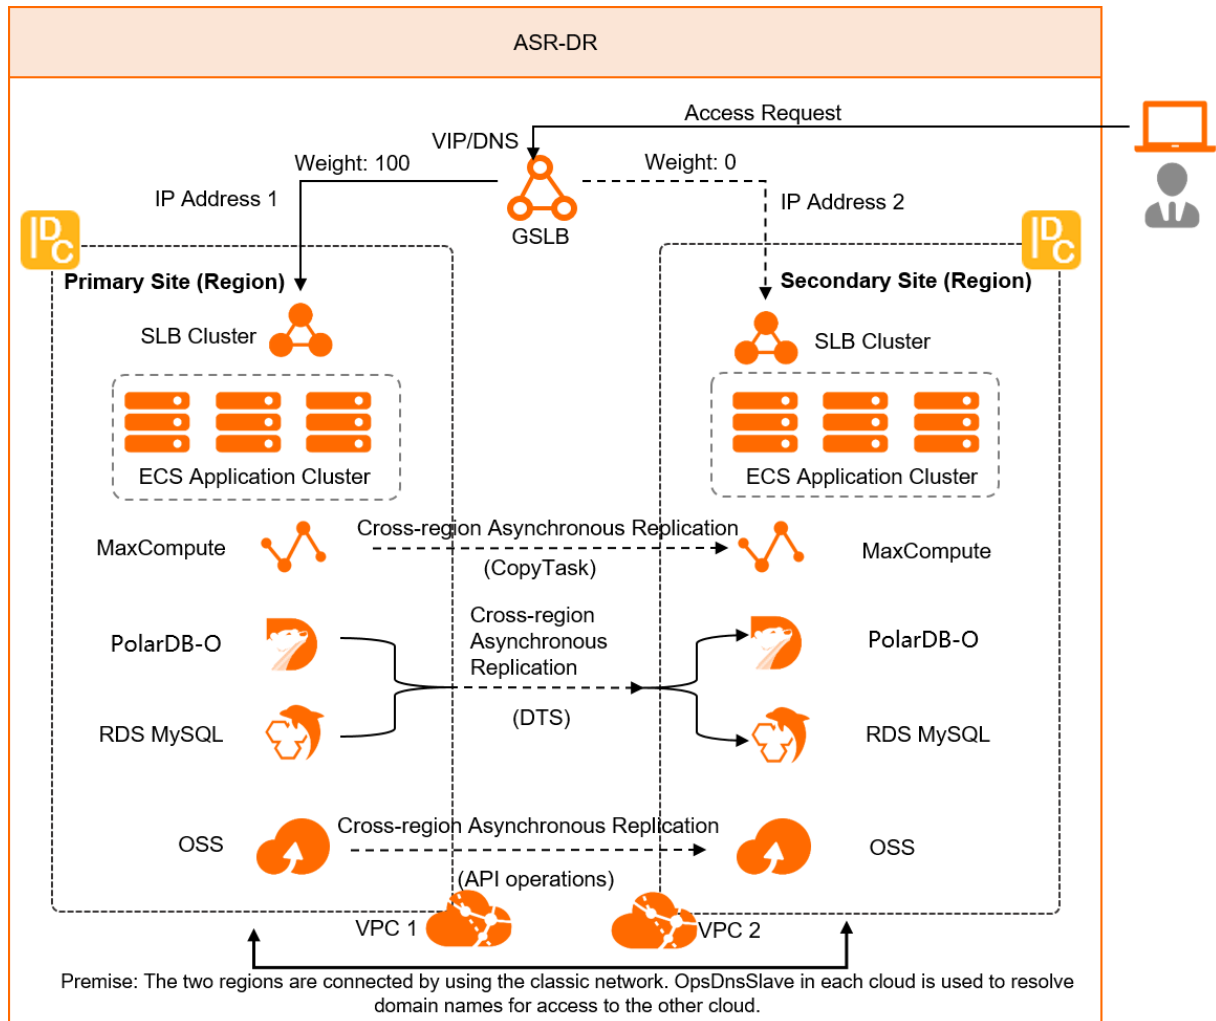
Apsara Stack Resilience for Geo-disaster Recovery

Apsara Stack Resilience for Zone-disaster Recovery consists of a primary data center and a secondary data center. The secondary data center is in a different region from the primary data center to cope with possible disasters at the primary data center. After you create a protection group, you can establish a hot backup channel between the resources of the primary and secondary sites. Changes to the resources at the primary site are asynchronously replicated to the corresponding resources at the secondary site in hot backup mode. If the primary data center fails, applications or services are failed over to the secondary data center by using a failover plan. After the primary data center is recovered, business is failed back to the primary cloud by using a failback plan. This way, business continuity is ensured. Apsara Stack Resilience for Geo-disaster Recovery includes the cross-cloud mode and cross-region mode. The following figures show architectures of the two modes.

Cross-cloud mode



Cross-region mode



Apsara Stack Resilience for Geo-disaster Recovery offers the following main features: DR drills, system management, and DR Dashboard.

- DR drills: include DR for Business and Global DR. You can set protection groups, drill plans, and recovery plans. You can also associate plans and then start the associated plans in batches based on execution sequence number. Plans can be divided into switchover, switchback, failover, and failback.
- System management: includes user management, primary cloud instance configuration, secondary cloud instance configuration, log management, capacity management, self-disaster recovery, and historical messages.
- DR dashboard: displays metrics such as the RTO duration, synchronization latency satisfaction of both the drill plan and recovery plan, the numbers of successful and failed drill plans and recovery plans within a specified time period, and the overall progress of the switchover and switchback processes.

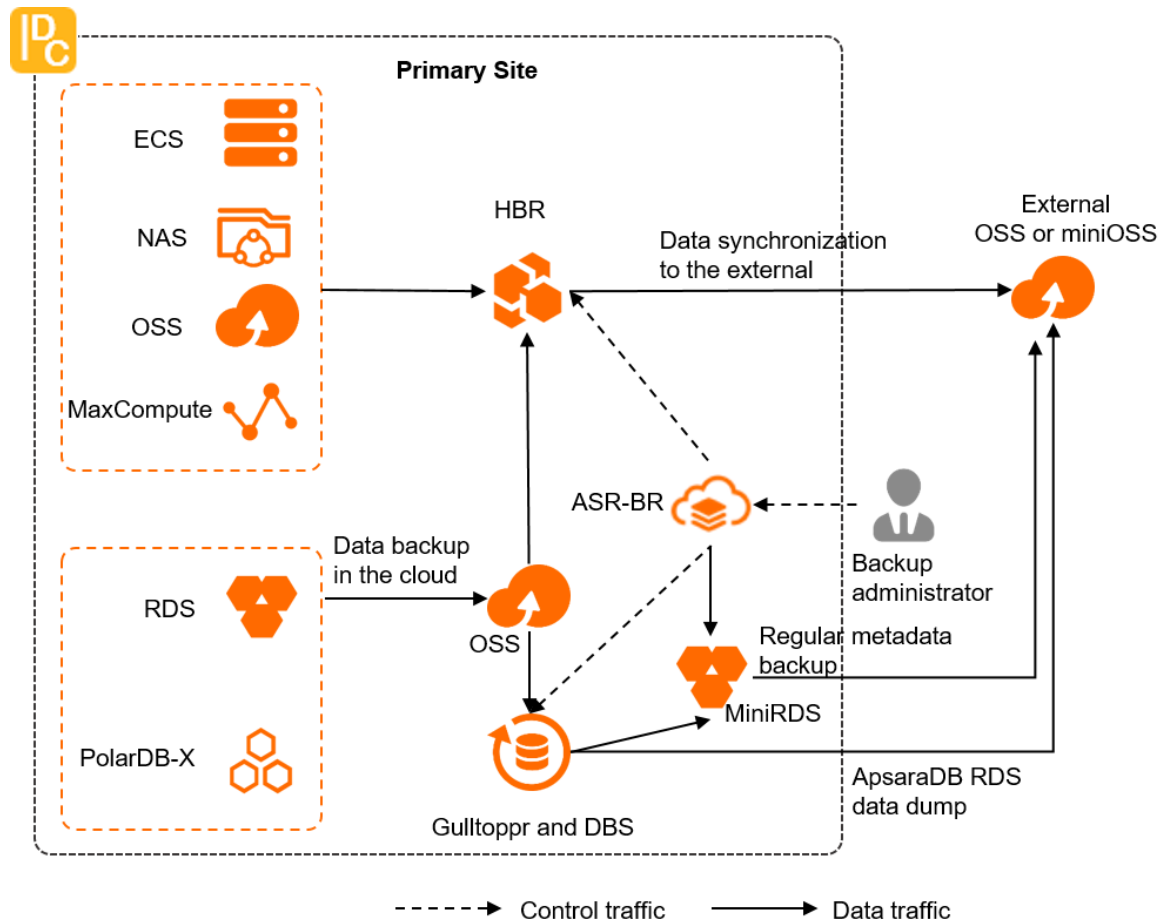
Apsara Stack Resilience for Geo-disaster Recovery is an application-based protection group disaster recovery mode and has the following advantages:

- Multiple disaster recovery scenarios are supported, such as many-to-one, primary/secondary, and cross-region.
- It is from the intuitive business perspective and supports the following features: disaster recovery drills for protection groups, protection groups switchover, and backward protection.
- Switchover of a cloud is not required if a single service is faulty.

Apsara Stack Resilience for Backup and Recovery

One or more secondary data centers are created to back up data of the primary site. Apsara Stack Resilience for Backup and Recovery provides the cloud-native backup capability and is required in MLPS-compliant projects. Apsara Stack Resilience for Backup and Recovery uses the unified backup management platform and supports virtual machine backup (including cloud disk), storage backup (excluding cloud disk), VMware backup, database backup, big data backup, and local backup of cloud platform metadata. Apsara Stack Resilience for Backup and Recovery supports the following services: ApsaraDB RDS, OSS, ECS, Apsara File Storage NAS, PolarDB-X 1.0, VMware virtual machine, and MaxCompute.

Backup architecture:



Apsara Stack Resilience for Backup and Recovery offers the following main features: backup, migration, and recovery.

- **Backup:** includes cross-region backup and local backup. You can create database backup vaults and create and delete storage backup vaults. All backup vaults are stored in OSS.
- **Recovery:** includes routine recovery, cloud rebuilding recovery, recovery on cloud of VMware VMs, and local recovery of VMware VMs.
- **Migration:** Migration plans are created to migrate data from on-premises VMware VMs to ECS instances.

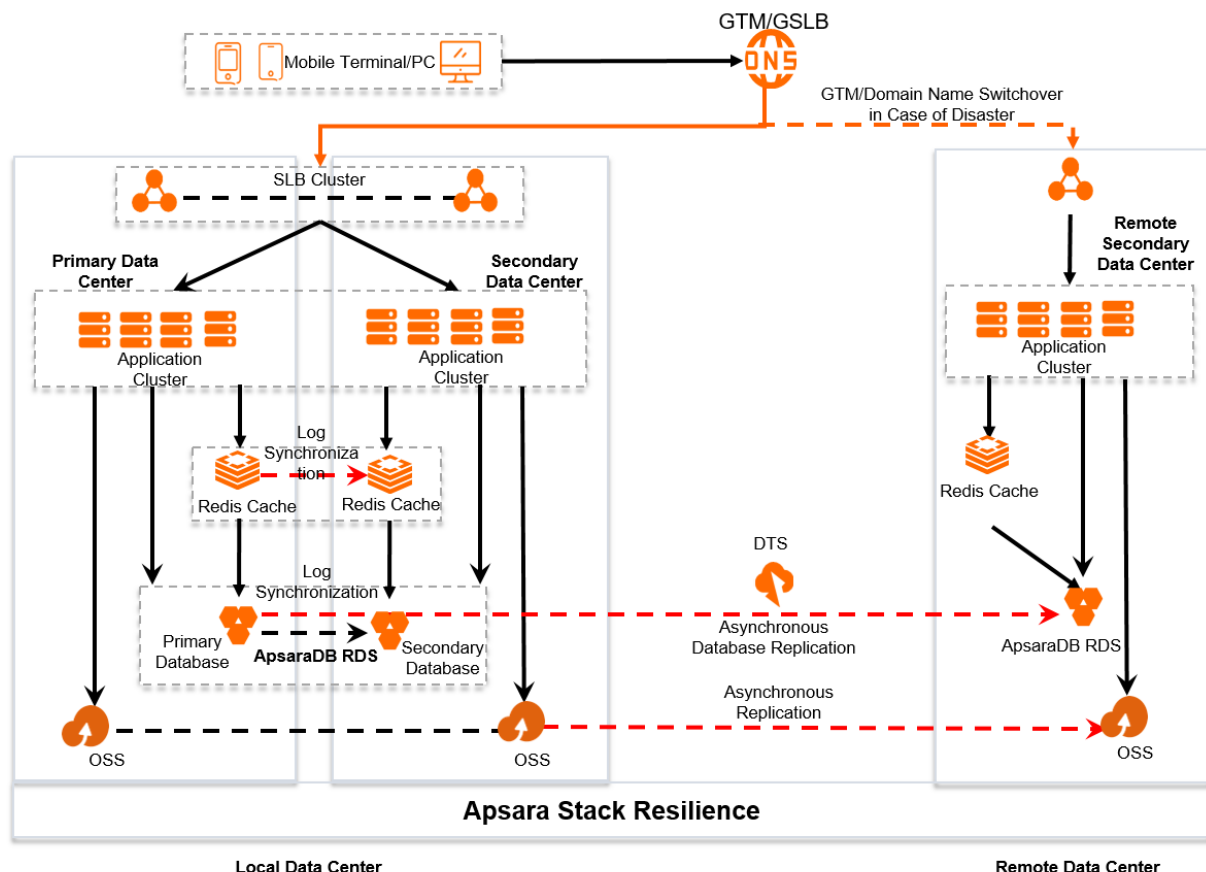
Apsara Stack Resilience for Backup and Recovery has the following advantages:

- A unified backup management platform is provided and can be used to back up databases, storage services, and MaxCompute.
- Physical backup is used for databases. The performance of the primary cloud instance is not affected because the backup time is very short.

Three data centers across two regions disaster recovery

Three data centers across two regions disaster recovery combines zone-disaster recovery and geo-disaster recovery. This solution is used to combat regional natural disasters and prevent regional failures. Two regions include the local and remote regions. Three data centers include the local primary data center, the local secondary data center, and the remote secondary data center. This solution provides active-active zone-disaster recovery. Data is synchronously replicated between two data centers. Data is asynchronously replicated between data centers in different regions that are in the active-standby mode. This mode has low business intrusion and can be quickly set up. It provides financial-grade reliability and unified disaster recovery management for full-stack cloud services.

Architecture of three data centers across two regions disaster recovery



Compared with traditional disaster recovery, the three data centers across two regions mode has the following advantages:

- The combination of zone-disaster recovery and geo-disaster recovery achieves lower RPO and RTO.
- This mode can be used to address regional failures and meet the high MLPS requirements in the finance industry.

7.2.2. Benefits

Unlike the disaster recovery that covers only ECS instances, ASR provides multiple disaster recovery solutions such as Apsara Stack Resilience for Zone-disaster Recovery, Apsara Stack Resilience for Geo-disaster Recovery, and Apsara Stack Resilience for Backup and Recovery (ASR-BR). A combination of these solutions is also supported. These make ASR an ideal choice for enterprises to ensure business continuity and data security. Apsara Stack Resilience

Apsara Stack Resilience for Zone-disaster Recovery

• Application-specific disaster recovery solutions

You can select appropriate cloud-based disaster recovery solutions based on application features. It is cost-effective to access applications because Apsara Stack Resilience for Zone-disaster Recovery is oriented toward cloud-native applications and no cloud migration is required. It can cover all distributed applications. For applications that use traditional architectures, you can create protection groups to migrate them to the cloud without the need for business transformation.

• Strong data consistency

The three data centers-based mode of Apsara Stack Resilience for Zone-disaster Recovery supports 0 synchronization latency for ApsaraDB RDS for MySQL, ApsaraDB for OceanBase, and Elasticsearch.

• Full-stack disaster recovery

Apsara Stack Resilience for Zone-disaster Recovery delivers full-stack disaster recovery schemes that cover base, cloud products, and Apsara Uni-manager. It can be used in cross-zone deployment for high availability. In the event of a disaster, business can be quickly switched over to the secondary zone. The endpoint of the cloud product remains unchanged.

• Full disaster recovery scenarios

Apsara Stack Resilience for Zone-disaster Recovery covers all disaster recovery scenarios. It allows you to customize hot switchover for cloud products to meet your routine disaster recovery drill requirements. It also supports data center-level fault recovery and single-product fault recovery.

Apsara Stack Resilience for Geo-disaster Recovery

• Intuitive business perspective

Apsara Stack Resilience for Geo-disaster Recovery is a business-based protection group disaster recovery. It supports the following features in an intuitive and visualized manner: disaster recovery drills for protection groups, push-button switchover of protection groups, and backward protection. Switchover of a cloud is not required if a single service is faulty.

• Abundant disaster recovery scenarios

Many-to-one, primary/secondary, and cross-region disaster recovery scenarios.

• Efficient disaster recovery drills

Supports flexible and custom disaster recovery drill plans and custom cloud product switchover policies. The drill plans are executed in a standardized process. You do not need to manually write drill scripts. The drill results are visualized.

• Controlled disaster recovery management

Supports multitenancy. Disaster recovery resources are isolated by organization to ensure data security. Each service can be switched independently to avoid mutual interference.

Apsara Stack Resilience for Backup and Recovery

• Backup of multiple services

Apsara Stack Resilience for Backup and Recovery provides the cloud-native backup capability and supports the following services: ApsaraDB RDS, Object Storage Service (OSS), Elastic Compute Service (ECS) instances and files, Apsara File Storage NAS, PolarDB-X 1.0, VMware virtual machine, and MaxCompute.

- **High backup efficiency for databases**

Physical backup is used for databases. The backup time is very short. All backup jobs can complete in time.

- **Multiple backup scenarios**

You can create backup plans as a tenant. ASR-BR supports periodic backup, manual backup, and cloud rebuilding recovery of the business data for cloud services.

- **Automated plug-in installation and connections**

The Tunnel mechanism is used to implement automatic one-way connections between VPC and storage network. Cloud Assistant is used to install the proxy plug-in and no manual installation is required.



- **High security**

The integration of the management page of Apsara Stack Resilience for Backup and Recovery with the cloud platform ensures that the account permission system is consistent. Security is enhanced because the cloud security mechanism is also integrated.

7.2.3. Scenarios

Selection of disaster recovery solutions is determined by a variety of factors, such as investment, acceptable RTO and RPO, and disaster types. The following table describes the scenarios of different disaster recovery solutions.

Disaster recovery solution	Scenarios	Industry
Two data centers-based mode of Apsara Stack Resilience for Zone-disaster Recovery	<ul style="list-style-type: none">• High business continuity is required.• Strong data consistency is not required.• The two data centers have independent power supplies, independent networks, and optical fiber length of no more than 50 kilometers in between.	Industries such as finance, healthcare, government, energy, and transportation

Three data centers-based mode of Apsara Stack Resilience for Zone-disaster Recovery	<ul style="list-style-type: none"> High business continuity is required. Strong data consistency is required to ensure RPO of 0 for instances in the same region. <div> <p> Note</p> <p>Three data centers-based mode of Apsara Stack Resilience for Zone-disaster Recovery supports ApsaraDB RDS for MySQL, ApsaraDB for OceanBase, Elasticsearch, Elastic Compute Service (ECS), Elastic Block Storage (EBS), Server Load Balancer (SLB) and Virtual Private Cloud (VPC).</p> </div> <ul style="list-style-type: none"> The three data centers have independent power supplies, independent networks, and optical fiber length of no more than 50 kilometers in between. 	Finance
Apsara Stack Resilience for Geo-disaster Recovery	<ul style="list-style-type: none"> High business continuity is required. Strong data consistency is not required. The capability to cope with regional failures is required. Distance between regions is more than 100 km. 	Industries such as government, energy, transportation, and healthcare
Apsara Stack Resilience for Backup and Recovery	<ul style="list-style-type: none"> Need to meet security requirements for core business data. Virtual machine backup (including cloud disks), storage backup (excluding cloud disks), VMware backup and recovery, database backup, big data backup, and cloud platform metadata backup are required. <div> <p> Note</p> <p>The following products are supported: ApsaraDB RDS, Object Storage Service (OSS), Elastic Compute Service (ECS) instances and files, Apsara File Storage NAS (NAS), PolarDB-X 1.0, VMware, and MaxCompute.</p> </div>	All industries, especially those with classified protection requirements.

Three data centers across two regions mode	<ul style="list-style-type: none">• A solution is needed to combat regional natural disasters.• A solution is needed to prevent regional faults.• Near 0 synchronization latency for business-critical data is required.• The budget is adequate.	Finance
--	--	---------

8. Computing services

8.1. ECS

Elastic Compute Service (ECS) is an IaaS offering on Alibaba Cloud that provides high-performance, stable, reliable, and scalable compute capacity in the cloud. You can create as many or as few ECS instances as needed to respond to changes in the requirements or popularity of your workloads.

8.1.1. ECS

Elastic Compute Service (ECS) is a computing service that provides scalable processing capabilities. It is easier to manage and more user-friendly than physical servers. You can create or release ECS instances and resize disks on demand anytime.

An ECS instance is a virtual computing environment that contains the basic components of computers such as the CPU, memory, and storage. Operations can be performed on ECS instances. Instances are core components of ECS. You can perform operations on instances in the ECS console. Other resources such as block storage devices, images, and snapshots can only be used after they are integrated into ECS instances.

Instances

An instance is the smallest computing service unit of ECS. Instance types comprise varying combinations of compute, memory, and storage capacity. An ECS instance is a virtual server that includes basic components such as CPUs, memory, an operating system (OS), network configurations, and disks. You can use management tools provided by Alibaba Cloud such as the ECS console and ECS API to create and manage ECS instances. You can manage the status of ECS instances and their deployed applications in the same manner as you would do with local servers. You can also upgrade compute, network, or storage resource specifications of your ECS instances as your requirements increase.

ECS provides a variety of instance families, such as shared instance families, dedicated instance families, instance families equipped with local HDDs, instance families equipped with local SSDs, and heterogeneous computing instance families. An instance family consists of one or more instance types that have similar attributes. ECS also provides innovative instance families that are developed based on state-of-the-art virtualization technology of Alibaba Cloud, such as ECS Bare Metal Instance families and super Computing Cluster (SCC) instance families.

- **ECS Bare Metal Instance**

ECS Bare Metal Instance is an innovative cloud computing service developed by Alibaba Cloud on top of SHENLONG architecture that integrates hardware and software. ECS Bare Metal Instance combines virtual machine features (such as elasticity of resources, resource delivery within minutes, and automated O&M) and physical machine features (such as performance consistency, hardware feature sets, and strong hardware-level isolation). ECS Bare Metal Instance is compatible with services in the Alibaba Cloud ecosystem and can help enterprises migrate their critical and high-load applications to the cloud.

- **SCC**

On top of ECS Bare Metal Instance families, SCC instance families use CPU-based devices and heterogeneously accelerated devices (such as GPU-accelerated devices) that are connected by high-speed InfiniBand (IB) networks to provide high-performance and highly parallel computing cluster services. SCC instance families are suitable for scenarios such as high-performance computing, artificial intelligence, machine learning, scientific and engineering computing, data analysis, and audio and video processing.

Block storage

Block storage is a random block-level storage service that features low latency, high durability, and high reliability. Alibaba Cloud provides a variety of block storage devices, such as Elastic Block Storage (EBS) devices based on a distributed storage architecture and local storage located on the same physical machine that hosts the associated ECS instances.

- **EBS**

EBS devices, also known as cloud disks, provide block-level random storage that features low latency and high performance, durability, and reliability for ECS instances. EBS devices use the triplicate distributed storage mechanism to ensure data durability for ECS instances. EBS devices can be created, released, and resized anytime.

- **Local storage**

Local storage, also known as local disks, are temporary disks attached to physical machines that host ECS instances. Local storage is designed for business scenarios that require high storage I/O performance. Local storage provides block-level data access capabilities for instances and provide low latency, high random IOPS, and high throughput.

- **Disk resizing**

You can resize disks to meet increasing storage requirements as your business and application data grow. You can resize the system disk or data disks of an instance online or offline. After you resize a disk of an instance offline, you must restart the instance for the resize operation to take effect.

- **Disk encryption**

Disk encryption is a simple and secure method to encrypt new disks. Disk encryption eliminates the need to create or maintain your own key management infrastructure, to modify existing applications and maintenance procedures, and to perform additional encryption operations. Disk encryption does not have negative impacts on your business.

Disk encryption can be used to encrypt the following types of data:

- Data on disks.
- Data transmitted between disks and instances. Data within instance operating systems is not encrypted.
- All snapshots created from an encrypted disk.

Images

An image is a running environment template of ECS instances. An image includes an operating system and pre-installed applications.

An image is a copy of data from one or more disks. An instance image can contain data from only the system disk or from both system and data disks. You can use an image to create an ECS instance or replace the system disk of an ECS instance.

Snapshots

A snapshot is a point-in-time backup of a disk. The snapshot service is typically applied to scenarios such as environment replication and disaster recovery and backup.

- **Scenarios**

- Environment replication

When you write and store data on a disk, you can create a snapshot for the disk and use the snapshot to create a disk. The new disk contains all data that was present on the original disk when the snapshot was created.

- Disaster recovery and backup

Disks are a secure storage method that ensures no data is lost. If incorrect data is stored on a disk due to specific reasons (such as application errors or malicious read and write operations performed by hackers exploiting application vulnerabilities), you can create a snapshot for the disk on a regular basis to restore the disk to a desired state.

- **Snapshot-consistent groups**

You can create a snapshot-consistent group to simultaneously create snapshots for one or more disks attached to an ECS instance. When a business system spans multiple disks, you can create a snapshot-consistent group to ensure a consistent write order and crash consistency of business system data.

- You can use snapshot-consistent groups to simultaneously create snapshots for multiple disks on an ECS instance.
- Snapshot-consistent groups are applicable to cluster services.
- After a snapshot-consistent group is created, you can use it to roll back one or more disks in the event of system failures or data errors caused by accidental operations.

- **Automatic snapshot policies**

Automatic snapshot policies can be applied to system disks and data disks to create snapshots for the disks on a periodic basis. You can use automatic snapshot policies to improve data security and tolerance against accidental operations. Automatic snapshot policies can effectively eliminate risks associated with manual snapshots:

Deployment sets

Deployment Set is a service provided by ECS that allows you to view the physical topology of hosts, racks, and vSwitches. You can select deployment policies based on your business requirements to improve the reliability and performance of your business.

When you use multiple ECS instances in the same zone, you can use a deployment set to distribute the instances differently to improve business reliability or network performance.

- Improve business reliability

To prevent the impacts caused by the failure of physical hosts, racks, or vSwitches, you can use deployment policies of deployment sets to distribute ECS instances that host identical applications across different physical hosts, racks, or vSwitches.

- Improve network performance

In scenarios that involve frequent network interactions between ECS instances, you can use deployment policies of deployment sets to associate the instances with the same vSwitch to achieve low-latency and high-bandwidth communication between the instances.

Networks

- **VPCs**

A virtual private cloud (VPC) is a custom private network that is established on Alibaba Cloud. VPCs are logically isolated from each other based on tunnels. You can create and manage cloud service instances in VPCs, such as ECS instances, and ApsaraDB RDS instances.

Each new ECS instance in a VPC is assigned a private IP address based on the CIDR block of the VPC and the CIDR block of the vSwitch to which the instance is connected.

- **EIPs**

An elastic IP address (EIP) is a public IP address that you can purchase and use as an independent resource. After an EIP is associated with an ECS instance, the instance can use the EIP to communicate with the Internet. EIPs can be associated with or disassociated from ECS instances. EIPs can also be deleted or have their public bandwidth modified.

- **ENIs**

An elastic network interface (ENI) is a virtual network interface controller (NIC) that can be bound to an ECS instance. You can use ENIs to deploy high-availability clusters and perform low-cost failover and fine-grained network management.

Security groups

A security group acts as a virtual firewall to control the inbound and outbound traffic of ECS instances to improve security. Security groups provide Stateful Packet Inspection (SPI) and packet filtering capabilities. You can use security groups and security group rules to define security domains in the cloud.

- You can add or modify security group rules based on your business needs to implement finer-grained access control. Security group rules can be used to control access to or from specific IP addresses, CIDR blocks, security groups, or prefix lists.
- New and modified rules in each security group are automatically applied to all instances that are associated with the security group.
- If an instance belongs to multiple security groups, the rules of all the security groups are applied to the instance. When an access request destined for the instance is detected, the request is matched against applied security group rules one by one based on the rule attributes such as protocol, port range, and priority. No sessions are established until an Allow rule matches the request.

Tags

Tags are used to identify resources. Each tag consists of a key and a value. You can use tags to categorize ECS instances for easy search and management. When the number of ECS instances increases, you can use tags to manage, group, and categorize resources so that you can quickly search for and perform batch operations on the resources. For example, you can replace images to deploy applications, upgrade patches, and add security group rules to control network access based on tags.

Dedicated hosts

Dedicated Host (DDH) is a specialized solution that Alibaba Cloud provides for enterprise users. DDH offers highly cost-effective features, such as dedicated physical resources, flexible deployment, and rich configurations.

A dedicated host is a cloud host whose physical resources are exclusively reserved for a single tenant. As the only tenant of a host, you do not need to share the physical resources of the host with other tenants. You can obtain the physical attributes of the host, including the number of CPUs, the number of sockets, the number of physical CPU cores, and the memory size. You can also create ECS instances of an instance family on a dedicated host if only the instance family is compatible with the host type.

Cloud Assistant

Cloud Assistant is a native automated O&M tool developed for ECS. It allows you to batch maintain ECS instances and batch run commands on and send files to ECS instances in a password-free, logon-free manner without the use of jump servers. These commands can consist of shell, PowerShell, or batch scripts.

Typically, you can use Cloud Assistant to install and uninstall software, start and stop services, distribute configuration files, and run commonly used commands (scripts).

Key pairs

Key pairs are a secure and convenient authentication method provided by Alibaba Cloud for ECS instance logons. A key pair consists of a public key and a private key. Only Linux instances support logon based on key pairs.

A key pair is a pair of public and private keys that are generated based on an encryption algorithm. By default, 2048-bit RSA key pairs are used. If you want to log on to a Linux instance by using a key pair, you must first create the key pair. You can specify a key pair when you create an instance, or bind a key pair to an instance after the instance is created. Then, you can use the private key to connect to the instance.

8.1.2. Benefits

Compared with data centers or traditional servers, Elastic Compute Service (ECS) is easy to use and features computing capacity that is scalable, secure, and reliable.

High availability

Compared with traditional servers that are limited by hardware, ECS instances offer a higher standard of O&M and integrate the features of various cloud services to provide more efficient backup, disaster recovery, and failover capabilities.

In addition, ECS provides supports in the following areas:

- Industry and ecosystem partners to help you build a more advanced and stable architecture.
- Diverse training services to help you implement high availability from the business end to the underlying basic service end.

Security

Security and stability are two primary concerns for cloud service users. Alibaba Cloud has received a host of international information security certifications that demand strict confidentiality of user data and user privacy protection, including Multi-Tier Cloud Security (MTCS) and ISO 27001.

- **Apsara Stack Virtual Private Cloud (VPC) uses simple configurations** to increase the flexibility, scalability, and stability of your business.
- **You can connect your own data centers to Apsara Stack VPC** by using leased lines to build a hybrid cloud. You can use a variety of hybrid cloud architectures and network services to provide robust networking.
- **VPCs are more stable and secure.**
 - VPCs allow you to divide, configure, and manage your network.
 - VPCs provide traffic isolation and attack isolation to protect your services against cyberattacks. You can establish a first line of defense against malicious attacks and traffic by deploying your business in VPCs.
- **Comprehensive security protection** is provided for ECS and includes security policies, security hardening, data security, monitoring, and alerts to improve the security of your business and defend against external attacks and unauthorized access.

VPCs provide a stable, secure, controllable, and fast-deliverable network environment. The capability and architecture of the VPC hybrid cloud bring the technical advantages of cloud computing to enterprises in traditional industries not engaged in cloud computing.

Elasticity

Elasticity is a key benefit of cloud computing.

- **Elastic computing**
 - **Vertical scaling**

Vertical scaling is the process of changing the configurations of ECS instances. In a traditional data center, it can be difficult to change the configurations of individual servers. However, in Apsara Stack, you can change the configurations of your ECS instances based on the volume of your business.

- **Horizontal scaling**

Horizontal scaling allows you to scale the quantity of resources for applications. A traditional data center may not be able to immediately provide sufficient resources for online gaming or live video streaming applications during peak hours. The elasticity of cloud computing makes it possible to provide as many resources as required during peak hours. When the load returns to normal levels, you can release redundant resources to reduce operation costs.

The combination of ECS vertical and horizontal elasticity and Auto Scaling enables you to scale resources up and down by specifying quantities as scheduled or against business loads.

- **Elastic storage**

Apsara Stack provides elastic storage. In a traditional data center, you must upgrade server configurations or replace servers to increase the storage space. In Apsara Stack, you can resize attached disks or attach more disks to instances to increase the storage space.

- **Elastic network**

Apsara Stack provides network elasticity. When you purchase Apsara Stack VPCs, you can configure the VPCs in the same way as you would configure data centers. In addition, VPCs have the following benefits: interconnection between data centers, separate secure domains in data centers, and flexible network configurations and planning within a VPC.

In conclusion, Apsara Stack provides elastic computing, storage, networking as well as business architecture planning. You can use Apsara Stack to build your business portfolio in any way.

Ease of use

Apsara Stack provides an easy-to-use console for centralized management. You can use the console to perform operations on ECS instances and ECS-related services to have ECS instances created, related services activated, and the configurations of the instances and services modified. Apsara Stack provides the following resources to help you use it in an easy way:

- Diversified instance families such as shared instance families, dedicated instance families, and ECS Bare Metal Instance families.
- VPCs that serve various purposes. You can choose VPCs based on your business requirements and organizations.
- Diversified storage types.
- Easy-to-use security group policies.

Apsara Stack also allows you to use images, deploy clusters, use custom tags, migrate virtual machines, and modify configurations.

8.1.3. Scenarios

ECS instances can be used either independently as simple web servers or with other Apsara Stack services such as OSS and CDN to provide advanced multimedia solutions. The following sections describe the typical application scenarios of ECS instances:

Official websites for enterprises and simple web applications

Initially, official websites for enterprises do not have high volumes of traffic and only require low-configuration ECS instances to run applications and databases and store files. As your website develops, you can upgrade the configurations and increase the number of ECS instances at any time without the need to worry about insufficient resources during traffic spikes.

Multimedia and high-traffic applications or websites

When you use ECS instances together with OSS, you can store static images, videos, and downloaded packages in OSS to reduce storage costs. You can also use ECS in combination with CDN or SLB to shorten user response time, reduce bandwidth fees, and improve availability.

Applications or websites that have large traffic fluctuations

Some applications and websites may encounter large fluctuations in traffic within a short period of time. ECS provides elastic processing capabilities. The number of ECS instances automatically increases or decreases in response to changes in traffic to meet resource requirements and preserve cost efficiency. ECS can be used in combination with SLB to implement a high availability architecture.

Databases

Databases with high I/O requirements are supported. High-configuration I/O optimized ECS instances can be used together with standard SSDs to support high I/O concurrency and higher data reliability. Alternatively, multiple low-configuration I/O optimized ECS instances can be used in combination with SLB to implement a high availability architecture.

8.2. Auto Scaling

Auto Scaling is a web service that automatically adjusts the number of instances available to handle your business loads based on your configurations. You can use Auto Scaling to scale out Elastic Compute Service (ECS) instances or elastic container instances during peak hours to ensure sufficient computing power. You can also use Auto Scaling to scale in ECS instances or elastic container instances during off-peak hours to minimize resource costs.

You are not charged when you use Auto Scaling. Auto Scaling provides the following benefits: automation, cost-effectiveness, high availability, easy audit, and intelligence. Auto Scaling helps you obtain a sufficient number of instances to handle stable or fluctuating loads for your application.

8.2.1. Features

Auto Scaling is a web service that automatically adjusts the number of instances based on your business requirements and configurations.

Auto Scaling can scale ECS instances and elastic container instances. During peak hours, Auto Scaling creates ECS instances or elastic container instances and adds the instances to your scaling group to ensure sufficient computing power. During off-peak hours, Auto Scaling removes ECS instances or elastic container instances from your scaling group to minimize resource costs.

Auto Scaling provides the following features: automatic scale-out and scale-in and elastic recovery. The following table describes the features.

Feature	Description
---------	-------------

Automatic scale-out	<p>To prevent issues such as access latency and resource overload, Auto Scaling automatically scales out underlying resources during peak hours.</p> <p>For example, when the average CPU utilization of ECS instances or elastic container instances in your scaling group exceeds 80%, Auto Scaling automatically creates new ECS instances or elastic container instances and adds the instances to your scaling group. Auto Scaling adds the ECS instances or elastic container instances to the backend server groups of the Server Load Balancer (SLB) instances that are associated with your scaling group. Auto Scaling also adds the private IP addresses of ECS instances to the whitelists of the ApsaraDB RDS instances that are associated with your scaling group.</p>
Automatic scale-in	<p>To prevent resource wastes, Auto Scaling automatically scales in underlying resources during off-peak hours.</p> <p>For example, if the average CPU utilization of ECS instances or elastic container instances in your scaling group drops below 30%, Auto Scaling automatically removes ECS instances or elastic container instances from your scaling group. Auto Scaling removes the ECS instances or elastic container instances from the backend server groups of the SLB instances that are associated with your scaling group. Auto Scaling also removes the private IP addresses of the ECS instances from the whitelists of the ApsaraDB RDS instances that are associated with your scaling group.</p>
Elastic recovery	<p>If Auto Scaling detects that an ECS instance or elastic container instance in your scaling group does not run as expected, Auto Scaling considers the instance unhealthy.</p> <p>Auto Scaling automatically releases unhealthy ECS instances or elastic container instances and creates new instances. The elastic recovery feature ensures that unhealthy instances are detected at the earliest opportunity and the number of healthy instances is always greater than or equal to the minimum number of instances that must be included in your scaling group.</p>

Scaling group

Scaling groups are the core unit of Auto Scaling. Instances in a scaling group are of the same type and can be used in similar business scenarios. If you have different business scenarios, you can create multiple scaling groups. Auto Scaling adjusts the number of instances in each scaling group based on your configurations and business requirements. You must specify the minimum number of instances that must be included and the maximum number of instances that is allowed for your scaling groups. You must also associate your scaling groups with SLB instances and ApsaraDB RDS instances.

Scaling configuration source

Scaling configuration source defines the instance configuration information. Auto Scaling scales out ECS instances or elastic container instances based on your scaling configuration source. That is, Auto Scaling automatically creates ECS instances or elastic container instances based on your scaling configuration source and adds the instances to your scaling group.

You can specify a launch template or scaling configuration as your scaling configuration source. You can specify a launch template only for scaling groups that contain ECS instances. You can have only one active scaling configuration source in your scaling group at a time. If you apply a new scaling configuration source to your scaling group, the current scaling configuration source becomes invalid.

Scaling rule

Scaling rules define specific scaling operations. For example, you can execute a scaling rule to add N ECS instances or elastic container instances to or remove N ECS instances or elastic container instances from your scaling group. The purpose of a scaling rule varies based on the rule type. You can use a scaling rule to trigger a scaling activity or adjust the maximum and minimum numbers of instances for a scaling group.

Auto Scaling supports step scaling rules, predictive scaling rules, simple scaling rules, and target tracking scaling rules. Predictive scaling rules can be used to predict future metric values based on historical monitoring data. These rules can automatically adjust the maximum and minimum numbers of instances in scaling groups. Step scaling rules, target tracking scaling rules, and simple scaling rules can be used to add or remove instances.

Scaling activity

When you execute a scaling rule or you manually add instances to or remove instances from a scaling group, a scaling activity is triggered. A scaling activity describes an event in which the number of ECS instances or elastic container instances in a scaling group changes, the maximum or minimum number of instances that are allowed in a scaling group changes, or the expected number of instances in a scaling group changes. After a scaling activity is triggered, Auto Scaling automatically scales instances based on your business requirements.

Scaling task

In a scaling task, Auto Scaling executes a specific scaling rule. Auto Scaling supports scheduled and event-triggered tasks.

- **Scheduled task:** If your business loads have specific patterns, such as periodic peak hours and off-peak hours, you can create scheduled tasks to manage your business loads in the Auto Scaling console. Before peak hours, Auto Scaling executes scheduled tasks to prepare sufficient computing resources. After peak hours, Auto Scaling executes scheduled tasks to release idle computing resources.
- **Event-triggered task:** If your business loads do not have specific patterns, you can create event-triggered tasks to manage your business loads in the Auto Scaling console. Event-triggered tasks monitor metrics of scaling groups. Auto Scaling dynamically scales instances based on the monitoring data. To prevent resource overload during peak hours, Auto Scaling executes event-triggered tasks to add instances to scaling groups. To reduce resource costs during off-peak hours, Auto Scaling executes event-triggered tasks to remove excess instances from scaling groups.

8.2.2. Benefits

If you manually manage your Elastic Compute Service (ECS) instances or elastic container instances, your costs are high. If you use Auto Scaling to manage your ECS instances or elastic container instances, the costs of infrastructure and O&M are minimized. Auto Scaling provides the following benefits: automation, cost-effectiveness, high availability, intelligence, and easy audit.

Automation

Auto Scaling automatically scales out or scales in instances in your scaling group based on your configurations. This prevents the issues that are caused by manual operations. Auto Scaling integrates with Server Load Balancer (SLB) and ApsaraDB RDS. This allows Auto Scaling to automatically add instances to or remove instances from the backend servers of the SLB instances that are associated with your scaling group. Auto Scaling can also automatically add the private IP addresses of ECS instances to the whitelists of the ApsaraDB RDS instances that are associated with your scaling group.

Cost-effectiveness

Auto Scaling creates and adds resources to scaling groups when the demand for the resources increases, and removes resources from scaling groups when the demand for the resources decreases. This helps minimize resource costs and increase resource utilization. For example, Auto Scaling automatically creates and adds ECS instances or elastic container instances to scaling groups during peak hours and automatically removes excess instances from scaling groups during off-peak hours. You do not need to invest a large amount of time and manpower to adjust the number of instances. Auto Scaling automatically scales instances, which reduces infrastructure and manpower costs and saves time.

High availability

Auto Scaling monitors the instances in your scaling group in real time. Auto Scaling can detect unhealthy ECS instances or elastic container instances at the earliest opportunity and automatically replaces unhealthy instances with new instances. This ensures that your instances can provide services in a consistent manner.

Intelligence

Auto Scaling supports various scaling modes, such as fixed-number mode, health mode, scheduled mode, dynamic mode, and custom mode. You can combine the scaling modes to configure scaling policies that meet your business requirements. Compared with manual operation, Auto Scaling allows you to scale instances in a more intelligent and efficient manner.

Easy audit

Auto Scaling logs the details of each scaling activity that is triggered. This facilitates troubleshooting. Auto Scaling also integrates with CloudMonitor. You can use CloudMonitor to monitor instances in scaling groups. You can also use CloudMonitor to obtain instance status in a simplified manner. You do not need to repeatedly check the status of multiple ECS instances or elastic container instances.

8.2.3. Scenarios

This topic describes the scenarios in which you can use Auto Scaling.

Business loads that do not have specific patterns

For example, you have a news website. When breaking news is reported, the number of page views of your news website immediately increases. As the news becomes less popular, the number of page views gradually decreases. In this case, you can use Auto Scaling to automatically scale your computing resources. The preceding scenario shows that the loads on your news website do not have specific patterns and the increase and decrease of access traffic to your news website is unpredictable. Therefore, you may find it difficult to determine the appropriate number of instances and adjust the number of instances in a timely manner. To address this issue, you can create event-triggered tasks in the Auto Scaling console. Auto Scaling automatically executes the event-triggered tasks to scale instances based on the performance of metrics such as CPU utilization.

Business loads that have specific patterns

For example, you have a game company. The number of game players increases at 18:00 and decreases at 22:00 every day. In this case, you can use Auto Scaling to scale out instances during peak hours and scale in instances during off-peak hours. Although your business loads have specific patterns, you may still need to invest in manpower and time to manually adjust the number of instances. To address this issue, you can create scheduled tasks in the Auto Scaling console. Auto Scaling executes the scheduled tasks to scale instances at the specified point of time.

Business loads that have small fluctuations

For example, you have a telecommunications company. Your business loads have no obvious fluctuations during a specific period of time. If one of your instances suddenly fails, your service is interrupted. In this case, troubleshooting or replacing the faulty instance is difficult. To address this issue, you can use the health check feature of Auto Scaling.

Business loads that have the preceding characteristics

For example, you have a video production company. Your daily business loads have no obvious patterns. However, your business loads may fluctuate during a specific period of time.

If you purchased a specific number of subscription ECS instances, and you want to only adjust the number of ECS instances whose loads have obvious fluctuations, you can manually add the subscription ECS instances to your scaling group. Then, you can create and execute event-triggered tasks in the Auto Scaling console to enable automatic scaling of instances based on the performance of metrics such as CPU utilization. This ensures that your instances can provide services in a consistent manner and minimizes your resource costs. You can also use the scheduled task feature and health check feature. You can combine the features of Auto Scaling based on your business requirements to ensure better user experience.

8.3. ROS

Resource Orchestration Service (ROS) is a service provided by Apsara Stack to allow you to manage cloud computing resources in a simplified manner.

You can author a stack template based on the template syntax that is defined in ROS. In the template, you can define the required cloud computing resources, such as Elastic Compute Service (ECS) and ApsaraDB RDS instances, and the dependencies between resources. The ROS engine automatically creates and configures all resources in a stack based on the template. This helps achieve automatic deployment and O&M. ROS helps you build infrastructures in the cloud to implement Infrastructure as Code (IaC). Compared with the calls for API operations of cloud services, ROS significantly improves the efficiency of your business.

8.3.1. Features

ROS is a service provided by Apsara Stack to allow you to manage cloud computing resources in a simplified manner. You can use stacks and templates to achieve automatic deployment and O&M of the resources.

Stack

ROS manages a group of Apsara Stack resources in a centralized manner based on the logic of a stack. You can create, update, recreate, and delete Apsara Stack resources by stack.

Template

A template is a JSON text file that is encoded in UTF-8. Templates are used to create stacks and serve as the blueprints for infrastructures and architectures. You can define the configurations of Apsara Stack resources in a template to specify the dependencies between the resources.

An ROS template is a standardized method that is used to deliver resources and applications. Independent software vendors (ISV) can use ROS templates to deliver a holistic system or solution that encompasses cloud resources and applications. This way, Apsara Stack resources can be integrated with the software systems for centralized delivery.

Developers and administrators can define the required Apsara Stack resources, such as ECS and ApsaraDB RDS instances, and the dependencies between resources in a template. The ROS engine automatically creates and configures all resources in a stack based on the template. This helps achieve automatic deployment and O&M.

8.3.2. Benefits

ROS can help you model and configure your Apsara Stack resources. After you create a template in which the required resources such as ECS and ApsaraDB RDS instances are defined, ROS creates and configures the resources based on the template. This way, you can manage resources in a simplified manner.

Improved deployment efficiency

You can use a template in ROS to provision your entire cloud environment. As your business grows, you can use the template to deploy cloud resources in new zones. You can also use the template to deploy the development, testing, and production environments. This helps improve the deployment efficiency and reduce the risk of human errors.

Reduced cost

If you use a template in ROS to provision your cloud environment, ROS can automatically deploy or delete a large number of stacks based on your business requirements. ROS provides elastic cloud resources to help reduce costs.

Compliance control

ROS is an IaC service. You can use templates to define infrastructures. When you want to create or update a template, you can review the code. Continuous integration (CI) and continuous delivery (CD) are integrated into templates to ensure that the templates can meet the management regulations of your organization and improve the security compliance of your cloud environment.

8.3.3. Scenarios

ROS can be used in a wide range of scenarios. You can use ROS to migrate your business to the cloud, deploy multiple resources at a time, and distribute environments based on your business requirements. ROS uses only approved templates to deploy cloud environments. This helps meet IT compliance requirements and minimize financial risks.

Migration of business to the cloud

Based on the best practices accumulated from Apsara Stack, ROS provides a solution to help you deploy all resources in a few steps and optimize the cloud architecture. No professional IT skills or experience in cloud architecture design are required.

On-demand deployment of resources

You can use a template to deploy multiple application runtime environments at a time in business scaling or DevOps scenarios.

Business environment distribution

In centralized IT management scenarios, ROS distributes a standardized environment across regions and accounts to meet the business requirements of each organization and team.

Cloud environment management

ROS uses only approved templates to deploy cloud environments. This helps meet IT compliance requirements and minimize financial risks.

8.4. OOS

Operation Orchestration Service (OOS) is an automated O&M service provided by Alibaba Cloud to help automatically manage and execute O&M tasks.

8.4.1. Service details

You can use templates to define O&M tasks, including task execution sequence, task input, and task output. Operation Orchestration Service (OOS) can automatically execute the O&M tasks after you define the tasks in the templates. OOS can be used to manage other Alibaba Cloud services, such as Elastic Compute Service (ECS), ApsaraDB RDS, Server Load Balancer (SLB), and Virtual Private Cloud (VPC).

Template

OOS uses templates to define the O&M operations that require orchestration. The template content is specified in the YAML or JSON format. The templates are classified into public templates and custom templates.

Execution

OOS supports a variety of execution modes, including automatic execution, semi-automatic execution, and manual execution. This helps you complete diverse O&M tasks.

- Automatic execution: We recommend that you execute a template in this mode in the test environment first. This way, you can have a better understanding of the O&M operations performed by using the template. If the test result meets the expectations, you can execute the template in this mode in the production environment.
- Manual execution: This execution mode is similar to debugging. If you want to know more about the execution of each task, we recommend that you use the manual execution mode.

8.4.2. Benefits

Operation Orchestration Service (OOS) can help enterprises better standardize, manage, and perform automated O&M operations. You can define required operations by using templates and execute the operations in the system. This improves the overall efficiency of O&M operations, enhances the security of O&M operations, and reduces manual O&M errors.

Visualized execution processes and results

You can view the execution process and the result of a template in the OOS console, including:

- The input, output, and details of each O&M task
- The procedure and sequence for executing tasks and redirection errors of failed tasks

Fully managed service

OOS provides a fully managed and serverless O&M service that enables automatic execution of O&M tasks. O&M tasks can be automatically executed without consuming your computing resources, such as Elastic Compute Service (ECS) instances. OOS can meet the O&M requirements of startups, small and medium-sized enterprises, and large enterprises. OOS minimizes the effort of enterprises in O&M and allows enterprises to focus on business growth.

Efficiency in managing multiple O&M tasks at a time

Managing multiple O&M tasks at a time is more complex than managing a single task. To enhance O&M efficiency, you can use OOS to manage task progress, monitor task status, and locate execution errors.

Superior authentication and verification systems

You can use Resource Access Management (RAM) to manage OOS. You can manage and audit permissions on OOS configurations and OOS-based O&M operations on other Alibaba Cloud services. You do not need to worry about the security of configurations and operations.

Easy-to-use templates

OOS provides a variety of templates that are easy to use to minimize your effort in creating templates. OOS integrates common O&M operations for popular Alibaba Cloud services and provides easy-to-use templates. This helps improve O&M efficiency. OOS provides public templates for common O&M tasks. You can select a public template and modify the template based on your business requirements.

Operations as code

OOS provides templates to standardize common O&M tasks. In OOS, you can manage templates by using code. After you create, review, and publish a template to the production environment, you can use the template to generate O&M tasks. This ensures the security and unification of O&M tasks and helps deliver best practices of operations as code.

Delegated authorization

Authorization must be appropriately delegated to O&M engineers. O&M engineers must be granted the required permissions to complete the corresponding O&M tasks. To prevent risks, they cannot be granted higher permissions than required to execute unexpected O&M tasks. OOS provides a delegated authorization solution to meet the preceding O&M requirements. When a high-privileged O&M engineer creates a template, a RAM role can be specified and granted the required permissions to use the template. The high-privileged O&M engineers can also grant a lower-privileged O&M engineer the permissions to execute the template. If the duty of executing the template is transferred to the lower-privileged O&M engineer, accidental operations that require high privileges may be reduced.

8.4.3. Scenarios

Common O&M scenarios of Operation Orchestration Service (OOS) include executing multiple O&M tasks at a time, updating images, managing approval requests, and managing scheduled O&M tasks. You can create custom templates based on your business requirements.

Execute multiple O&M tasks at a time

In some scenarios, you need to perform O&M operations on multiple targets, such as Elastic Compute Service (ECS) instances, to ensure continual business operation and growth. For example, to check the available space of hard disks in multiple ECS instances at a time, you need to perform the following operations: select the ECS instances whose hard disks you want to check, perform the hard disk check by running a Cloud Assistant command, and then view the check results in a centralized manner. OOS provides multiple methods to select instances, including name matching, tag grouping, and resource group classifying.

Update images

To ensure a safe runtime environment for ECS instances, you must install the latest patches or update related components. OOS allows you to update ECS instance images. You can generate a new image from a source image, and use the new image for testing and production.

Manage approval requests

In many scenarios, you need to manage approval requests to ensure that only expected O&M operations are performed. You can define an ACS::Approve action in a template. This action ensures that an O&M task is approved before the task is executed and that expected O&M tasks are executed without mistakes.

Manage scheduled tasks

Some O&M tasks must be executed as scheduled. For example, each morning, you need to delete all Object Storage Service (OSS) files generated by the previous tests to prepare a clean runtime environment for the next test. In this case, you can create a template and schedule delete operations in the morning.

8.5. What is Container Service?

Container Service provides high-performance, enterprise-class management for scalable Kubernetes-based containerized applications throughout the application lifecycle.

Container Service simplifies the creation and scaling of container management clusters. It integrates Apsara Stack virtualization, storage, network, and security capabilities, providing the optimal environment to run Kubernetes-based containerized applications in the cloud. Alibaba Cloud is a Kubernetes certified service provider, with Container Service being among the first services to pass the Certified Kubernetes Conformance Program. Container Service provides professional container support and services.

8.5.1. Features

Container Service provides powerful management capabilities, including Kubernetes cluster management, one-stop container lifecycle management, and scheduling policies that are used to ensure high availability. These capabilities can help enterprises efficiently manage and run containerized applications in the cloud.

Cluster management

- Allows you to create a dedicated cluster that contains GPU servers within 10 minutes in the Container Service console.
- Provides OS images that are optimized for containerized applications and supports Kubernetes versions and Docker versions with high stability and enhanced security.
- Supports multi-cluster management, upgrades, and scaling.

One-stop container lifecycle management

- Networking
 - Provides high-performance virtual private clouds (VPCs) and elastic network interfaces (ENIs) that are optimized by Alibaba Cloud, boasting 20% increased performance compared with common network solutions.
 - Supports access control and traffic throttling for containers.
- Storage
 - Supports Alibaba Cloud disks and Object Storage Service (OSS) buckets, and provides the standard FlexVolume driver.
 - Supports dynamic volume creation and migration.
- Logging
 - Provides high-performance log collection based on Simple Log Service.
 - Supports integration with open source logging solutions provided by third parties.
- Monitoring
 - ACK supports container-level and VM-level monitoring.
 - Supports integration with open source monitoring solutions provided by third parties.

- **Permission control**
 - Supports cluster-level Resource Access Management (RAM) authorization.
 - Supports application-level permission control.
- **Application management**
 - Supports canary releases and blue-green releases.
 - Supports application monitoring and scaling.
- **Component management**

Provides various types of components. You can install, upgrade, or uninstall components based on your business requirements.

Scheduling policies for high availability

- Supports service-level affinity policies and scale-out policies.
- Provides cross-zone high availability and disaster recovery.
- Provides cluster and application management APIs to support continuous integration and integration with external systems.

Performance metrics

- Performance of an individual cluster: supports up to 1,000 nodes and management of 100,000 containers.
- Cluster management: supports management of up to 200 clusters.
- Container startup time: 100 containers, each with the size of 20 MB, can be started within 10 seconds at a time.

8.5.2. Benefits

Overview

Easy to use

- You can easily create Kubernetes clusters in the Container Service console.
- You can easily upgrade Kubernetes clusters in the Container Service console.

When you use custom Kubernetes clusters, you may need to handle clusters of different versions. Currently, each time you upgrade the clusters, you need to make major adjustments and high operation and maintenance costs are incurred. Container Service allows you to perform rolling upgrades based on images and supports full metadata backups. You can easily roll back clusters to previous versions.

- Allows you to easily scale Kubernetes clusters in the Container Service console.

Kubernetes clusters enable you to quickly scale up or down applications to handle traffic fluctuations in a timely manner.

Features

Feature	Description
Network	Supports continuous network integration to optimize network performance.

Load balancing	<p>Allows you to create public and internal SLB instances.</p> <p>If you use an Ingress to control access to your Kubernetes cluster, frequent service releases may negatively affect the performance of the Ingress and increase the error rate. Container Service allows you to create SLB instances, which provide high availability load balancing and can automatically modify network configurations to suit your business needs. This solution is adopted by a large number of users and has been proven to be a more stable and reliable alternative to Ingresses.</p>
Storage	<p>Supports Apsara Stack cloud disks, Network Attached Storage (NAS), and Block Storage, and provides FlexVolume drivers.</p> <p>Supports seamless integration with cloud storage services for custom Kubernetes clusters that cannot use cloud storage resources.</p>
O&M	<ul style="list-style-type: none"> • Supports integration with Apsara Stack Log Service. • Supports automatic scaling.
Image repository	<ul style="list-style-type: none"> • Provides high availability and high concurrency. • Supports accelerated image retrieval. • Supports peer-to-peer image distribution. <p>Custom image repositories may stop responding when millions of clients attempt to pull images at the same time. Container Service provides an image repository system that offers enhanced reliability and reduces O&M and upgrade costs.</p>
Stability	<ul style="list-style-type: none"> • Dedicated support teams guarantee the stability of containers. • All Linux and Kubernetes versions must pass rigorous testing before they are available to the public. <p>Container Service supports Docker CE and provides a Docker community to help you communicate with other Docker enthusiasts and solve problems. Best practices are provided to help you address issues, such as network interruptions, kernel incompatibilities, or Docker crashes.</p>
Technical support	<ul style="list-style-type: none"> • Allows you to quickly upgrade Kubernetes clusters to the latest version. • Provides professional technical support services to help you solve the issues that may occur when you use containers.

8.5.3. Scenarios

DevOps continuous delivery

Optimized continuous delivery pipeline

Container Service works with Jenkins to automate the DevOps pipeline, from code submission to application deployments. The service ensures that code is only submitted for deployment after passing automated testing, and provides a better alternative to traditional delivery models that involve complex deployments and slow iterations.

Benefits

- DevOps pipeline automation

Automates the DevOps pipeline, from code updates to code builds, image builds, and application deployments.

- Consistent environment

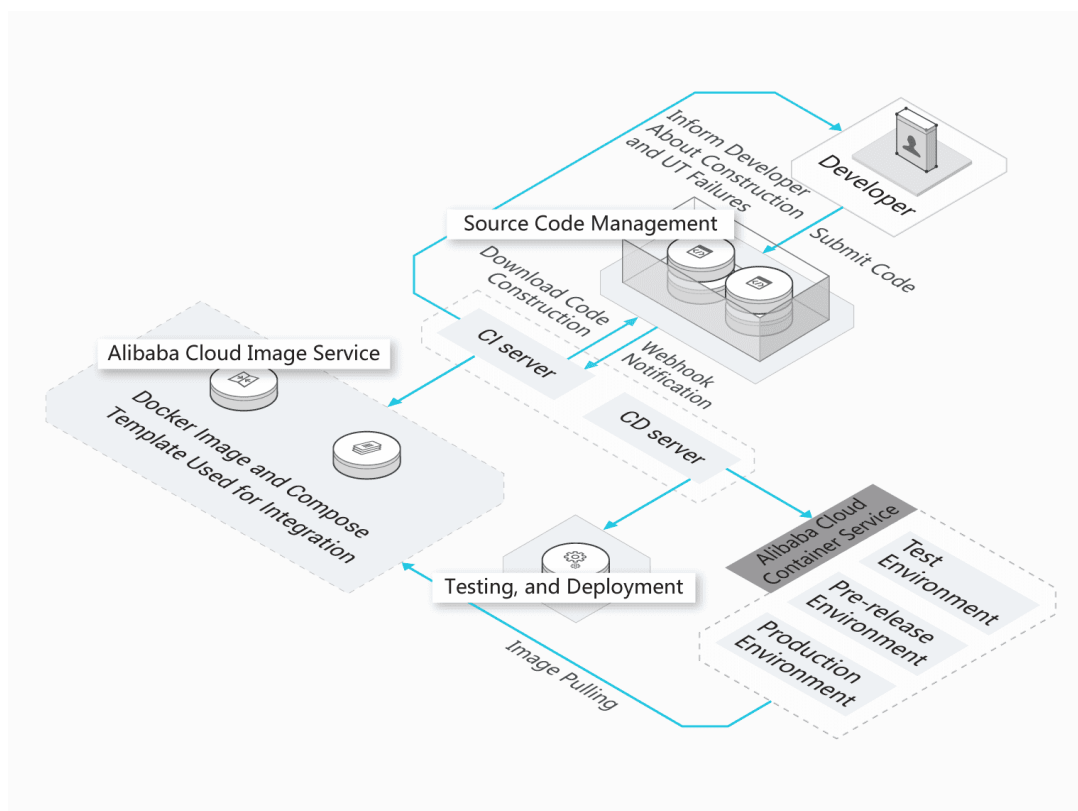
Allows you to deliver code and runtime environments based on the same architecture.

- Continuous feedback

Provides immediate feedback on each integration or delivery.

Related products and services

ECS + Container Service



Machine learning based on cloud-native technology

Enables rapid application developments with a focus on machine learning

Container Service allows data engineers to easily develop and deploy machine learning applications in heterogeneous computing clusters. Integrated with multiple distributed storage systems, the service supports faster read and write speeds to facilitate the testing, training, and release of data models. You can focus on your core business operations instead of worrying about the deployment and maintenance process.

Benefits

- **Ecosystem support**
Supports mainstream deep learning frameworks, such as TensorFlow, Caffe, MXNet, and Pytorch, and offers optimized features of these frameworks.
- **Quick start and elastic scaling**
Provides machine learning services for development, training, and inference. Supports the startup of training and inference tasks within seconds, and elastic scaling of GPU resources.
- **Easy to use**
Allows you to easily create and manage large-scale GPU clusters and monitor core metrics, such as GPU utilization.
- **Deep integration**
Seamless integration with Apsara Stack storage, logging and monitoring, and security infrastructure capabilities.

Related products and services

ECS/EGS/HPC + Container Service + OSS/NAS/CPFS

Microservices architecture

Agile development and deployment to speed up the evolution of business models

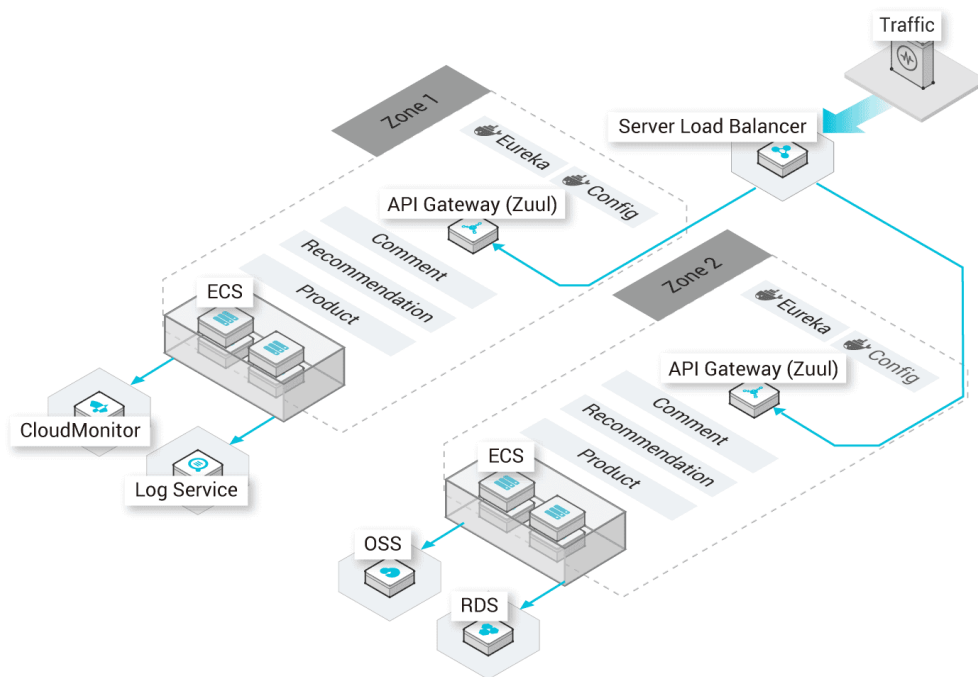
In the production environment, you can split your system into microservices and use Apsara Stack image repositories to store these microservice applications. Apsara Stack can schedule, orchestrate, deploy, and implement phased releases of microservice applications while you focus on feature updates.

Benefits

- **Load balancing and service discovery**
Forwards layer 4 and layer 7 requests and binds the requests to backend containers.
- **Multiple scheduling and disaster recovery policies**
Supports different levels of affinity scheduling policies, and cross-zone high availability and disaster recovery.
- **Microservices monitoring and auto scaling**
Supports microservice and container monitoring, and microservice auto scaling.

Related products and services

ECS + ApsaraDB RDS + OSS + Container Service



Hybrid cloud architecture

Unified O&M of cloud resources

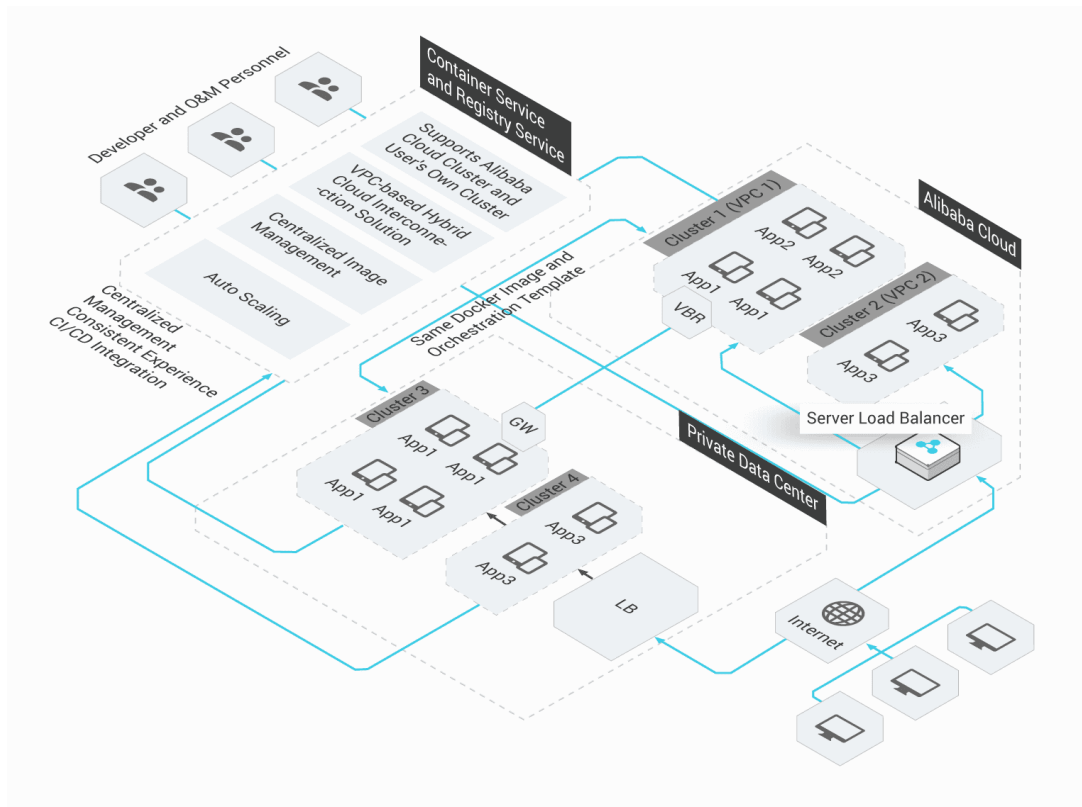
You can centrally manage cloud and on-premises resources in the Container Service console. Containers hide the differences between infrastructures. This enables you to use the same images and orchestration templates to deploy applications in the cloud and on premises.

Benefits

- Application scaling in the cloud
During peak hours, Container Service can scale up applications in the cloud and forward traffic to the scaled-up resources.
- Disaster recovery in the cloud
Business systems can be deployed on premises for service provisioning and in the cloud for disaster recovery.
- On-premises development and testing
Applications that are developed and tested on premises can be seamlessly released to the cloud.

Related products and services

ECS + VPC + Express Connect



Automatic scaling architecture

Traffic-based scalability

Container Service enables businesses to auto-scale their resources based on traffic. This prevents traffic spikes from bringing down your system and eliminates idle resources during off-peak hours.

Benefits

- Quick response

Container scale-out can be triggered within seconds when traffic reaches the scale-out threshold.

- Auto scaling

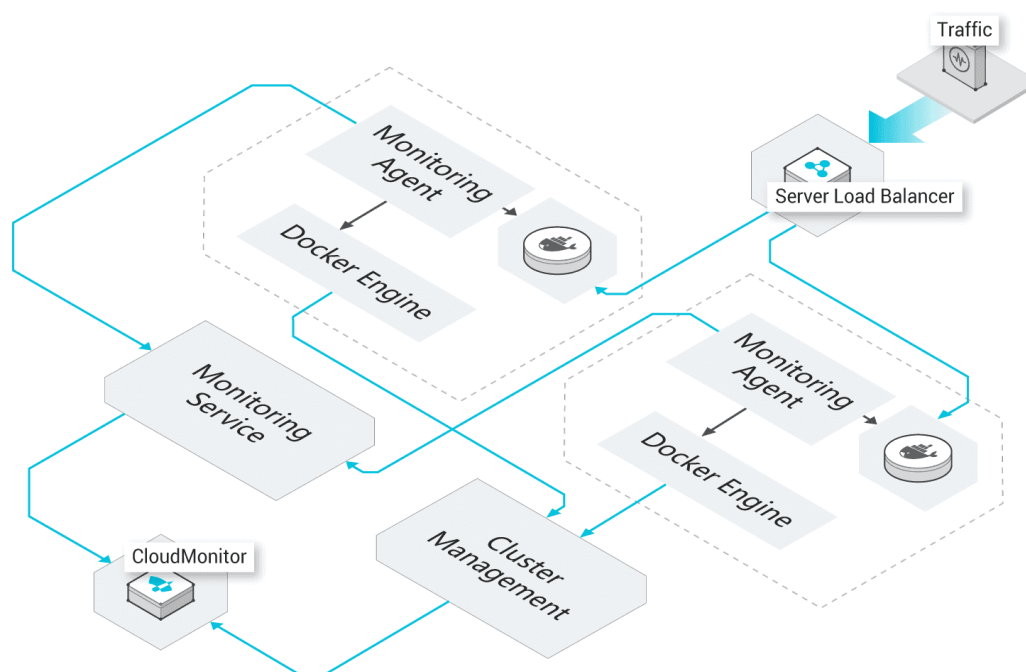
The scaling process is fully automated without human interference.

- Low cost

Containers are automatically scaled in when traffic decreases to avoid resource waste.

Related products and services

ECS + CloudMonitor



8.6. Container Registry

Container Registry is a platform that allows you to manage and distribute cloud-native artifacts in a secure and efficient manner. Cloud-native artifacts include container images and Helm charts that meet the standards of Open Container Initiative (OCI). Container Registry provides the following features: image permission management, synchronous image distribution, and content signing. The features allow you to manage the entire lifecycle of container images. Container Registry simplifies the setup and O&M of container registry. Container Registry is integrated with Alibaba Cloud services such as Container Service for Kubernetes (ACK) to easily create and deliver a one-stop solution for cloud-native applications.

8.6.1. Features

Container Registry is a platform that allows you to manage and distribute cloud-native artifacts in a secure and efficient manner. Cloud-native artifacts include container images and Helm charts that meet the standards of Open Container Initiative (OCI). Container Registry provides the following features: artifact management, image replication, artifact security, and deployment integration.

Artifact management

- Secure management: You can securely manage container images by namespace.
- Lifecycle management: You can query artifacts and image tags. You can also delete artifacts and image repositories.
- Fine-grained permission management: You can manage user permissions, Apsara Stack Cloud Management (ASCM) departments, and resource sets.
- Version immutability: You can configure version immutability for OCI artifacts.

Image replication

- Trigger: If a container image is updated, the corresponding event is automatically triggered.
- Image replication: You can manually trigger the replication of a container image of a specific version to implement geo-disaster recovery for container images. Container Registry can also automatically replicate a container image across multiple accounts after the image is pushed to an image repository.

Artifact security

- Encrypted image distribution: You can configure secure HTTPS protocol to distribute container images.
- Container image signing: This feature prevents man-in-the-middle (MITM) attacks and unauthorized image updates or deployments. This ensures image consistency and security from distribution to deployment.
- Image scanning: This feature allows you to scan container images to identify vulnerabilities.

Deployment integration

- Image pulls without a secret: You can configure image pulls without a secret in the Alibaba Cloud Container Service for Kubernetes (ACK) console. This way, you do not need to specify a secret for each image pull.
- Image selection: You can select image repositories and image tags when you configure a Deployment in the ACK console.

8.6.2. Benefits

Container Registry provides you with the following benefits: ease of use, security, and integrability.

Ease of use

- Allows you to create an image repository without the need to manually build and maintain the images.
- Allows you to pull and push images across multiple regions in a fast and stable manner.

Security and controllability

- Provides an all-in-one management system for image permissions. This ensures secure and convenient image sharing.
- Provides the image scanning feature to identify vulnerabilities in images and prompt vulnerability levels.

Efficient distribution

- Supports large-scale image distribution in a single region and concurrent image pulls on 500 nodes.
- Supports cross-region image replication and cross-cloud synchronous image distribution. Supports manual and automatic image replication.

Seamless integration with other Alibaba Cloud services

Integrated with Alibaba Cloud services such as Container Service for Kubernetes (ACK) to implement continuous deployment after images are updated.

8.6.3. Scenarios

Container Registry is suitable for scenarios, such as DevOps, continuous delivery, and automatic image replication.

DevOps and continuous delivery

Container Registry is integrated with Jenkins to automate the DevOps pipeline from code committing to application deployments and ensure that code is committed for deployment only after the code passes automated testing. This simplifies application deployments and accelerates application iterations.

- DevOps automation

Automates the DevOps pipeline, from code updates to code builds, image builds, and application deployments.

- Environment consistency

Allows you to deliver code, and deliver runtime environments that are built based on immutable architectures.

- Continuous feedback

Returns results in real time after integration or delivery.

Automatic image replication

If the container business of an enterprise is deployed in multiple regions and multiple clouds, container applications of the enterprise must be deployed across multiple regions and clouds after submission. Container Registry provides the image replication feature to improve automated distribution efficiency and disaster recovery capabilities and reduce manual O&M costs.

- Multi-scenario replication

- Supports cross-region, cross-cloud, and cross-account image replication.
- Supports manual replication. Supports automatic replication after images are updated.

- Optimized scheduling

Optimizes the scheduling of replication to increase the success rate of replication.

- Security compliance

Supports encryption of replication links to ensure the security of replicated data.

9.Storage services

9.1. CDS

Cloud Defined Storage (CDS) is a distributed file system that provides storage based on cloud services. CDS is secure, cost-efficient, and highly reliable. You can centrally manage resources in CDS and flexibly scale resources on and off the cloud. This facilitates local data storage and retrieval.

You can deploy cloud services, such as Object Storage Service (OSS), Simple Log Service, Elastic Block Storage (EBS), and Tablestore in CDS. Then, you can use CDS to store unstructured data such as files, images, and videos, and store, query, and analyze log data. This helps meet the requirements in different industries. CDS is applicable to big data scenarios such as mobile applications and large websites where you need to store unstructured data and process massive logs. CDS can provide one-stop storage solutions for enterprises in different industries in a cost-efficient, secure, and reliable manner.

9.1.1. Features

You can deploy services such as Object Storage Service (OSS) and Simple Log Service in storage clusters of Cloud Defined Storage (CDS). The features supported by a storage cluster vary based on the service that you deployed in the cluster.

The following table describes the differences among different storage services.

Service	Description	Scenario
Elastic Block Storage (EBS)	A high-performance and low-latency block-level storage service provided by Alibaba Cloud for Elastic Compute Service (ECS) instances. EBS supports random read and write operations. You can format an EBS device and create file systems on the device in the same way as you do with a physical disk.	EBS can meet the data storage requirements of most business scenarios.
Object Storage Service (OSS)	A storage service designed to store large amounts of unstructured data such as images, audio, and videos on the Internet. You can call API operations to access data in OSS.	OSS is suitable for scenarios such as website construction and separation of dynamic and static resources.

OSS

OSS is a secure, cost-effective, and highly reliable cloud storage service of Alibaba Cloud.

Feature	Description
Bucket and object management	Before you upload objects to OSS, you must create a bucket in OSS to store objects. After you create a bucket, you can manage the objects in the bucket based on bucket configurations, such as hotlink protection and lifecycle rules.

Object upload and download	You can upload objects of any formats to a bucket. After you upload objects to a bucket, you can share and download the uploaded objects based on the URLs of the objects. You can obtain the URLs of the uploaded objects one by one or at a time.
Access control	OSS provides access control lists (ACLs) for access control. An ACL is an access policy that is used to grant access permissions on buckets and objects. You can configure an ACL when you create a bucket or upload an object, or modify the ACL at any time after you create a bucket or upload an object.
Static website hosting	Static websites are websites in which all web pages consist of only static content, including scripts such as JavaScript code that can be run on a client. You can use the static website hosting feature to host your static website on an OSS bucket and use the endpoint of the bucket to access the website.
Hotlink protection	The hotlink protection feature allows you to configure a referer whitelist for a bucket to protect your resources in the bucket from unauthorized access. After the hotlink protection feature is configured, only requests from the domain names that are added to the referer whitelist are allowed.
Log management	<p>When you access OSS, large numbers of access logs are generated. You can enable and configure the logging feature for a bucket to manage logs. This way, OSS generates logs every hour based on the defined naming rule and stores the logs as objects in the specified bucket.</p> <p>The real-time log query feature integrates OSS with Simple Log Service and allows you to query and collect statistics on OSS access logs. You can use this feature to audit access to OSS, track anomalous events, and identify issues. This helps improve work efficiency and make informed decisions.</p>
Cross-origin resource sharing (CORS)	CORS is a standard cross-origin solution that is provided by HTML5 to allow web application servers to manage cross-origin access. This ensures the security of data transmission.
Lifecycle management	You can configure lifecycle rules to periodically delete expired objects. This reduces storage costs.
Retention policies	You can configure time-based retention policies for buckets. After you configure and lock a retention policy for a bucket, you can read objects from or upload objects to the bucket. However, you cannot delete the objects in the bucket or the retention policy within the retention period that is specified by the retention policy. You can delete the objects only after the retention period ends. You can configure the retention policies based on the Write Once Read Many (WORM) strategy to prevent objects from being deleted or overwritten within a specific period of time.

Image management	<p>You can use the image style feature to add multiple image processing parameters to an image style to perform complex operations on images.</p> <p>To prevent images that allow anonymous access in an OSS bucket from unauthorized use, you can enable source image protection for the bucket. After you enable source image protection for your bucket, anonymous users can access images in the bucket only by adding style parameters in the requests or by using signed URLs.</p>
Cloud Storage Gateway (CSG)	<p>You can read and write all objects in a specific OSS bucket over standard Network File System (NFS) and Server Message Block (SMB) protocols. CSG also uses on-premises storage to cache hot data and provides high-performance data access in addition to the large storage capacity of OSS buckets.</p>

EBS

EBS is a persistent random block storage service designed for ECS. EBS provides low latency and high reliability. EBS provides devices based on a distributed storage architecture.

EBS provides ECS instances with block-level storage that features low latency and high performance, durability, and reliability. EBS devices use a triplicate distributed storage mechanism to ensure data durability for ECS instances. You can create, release, and resize EBS devices at any time.

EBS supports cloud disks. Cloud disks are block-level storage devices designed for ECS instances. Cloud disks are classified by performance into the following types: premium performance disks, standard performance disks, ultra disks, and standard SSDs. A cloud disk can be attached to a single ECS instance that resides in the same zone as the cloud disk.

Feature	Description
Disk management	<p>After you create a cloud disk, you can attach the disk to an ECS instance to provide storage space for the ECS instance. Before you use a cloud disk to store data, you must format the disk.</p>
Storage cluster	<p>After you create a storage cluster, you can classify EBS clusters by metrics such as business type, and bind the EBS clusters to different partitions of the storage cluster. This way, EBS clusters of different business types are separated from each other.</p>
Disk snapshot	<p>A disk snapshot is a copy of data stored in a cloud disk at a specific point in time. You can schedule the system to periodically create disk snapshots to ensure business continuity. Snapshots are suitable for scenarios such as environment replication, disaster recovery, and data backup.</p>
Disk encryption	<p>You can encrypt a new disk to improve security in a simple and secure manner.</p>

Snapshot-consistent group	You can create a snapshot-consistent group to simultaneously create snapshots for one or more disks attached to an ECS instance. When a business system spans multiple disks, you can create a snapshot-consistent group to ensure a consistent write order and crash consistency of business system data.
Automatic snapshot policy	Automatic snapshot policies can take effect on system disks and data disks to periodically create snapshots for the disks. You can use automatic snapshot policies to improve data security and tolerance against operation faults.
Asynchronous replication	The asynchronous replication feature protects data across zones in a region based on the data replication capability of EBS. This feature can asynchronously replicate data from a disk in a zone to a disk in another zone in the same region to implement disaster recovery and backup. You can use this feature to implement disaster recovery for critical business to protect data in your databases and improve business continuity.
Replication pair-consistent group	You can create a replication pair-consistent group to manage replication pairs between the primary site and the secondary site in a centralized manner.
Multi-attach	After multi-attach is enabled for a premium performance disk, the disk can be attached to multiple ECS instances that support the Non-Volatile Memory Express (NVMe) protocol within the same zone to allow concurrent read and write access from the instances.

Tablestore

Tablestore is a NoSQL data storage service independently developed by Alibaba Cloud with Computer Software Copyright registered in China. Tablestore is built on the Apsara distributed operating system of Alibaba Cloud, and can store large amounts of structured data and allow real-time access to the data. Tablestore allows you to perform table operations and data operations, and provides features such as data versions and time to live (TTL), auto-increment primary key column, conditional update, Tunnel Service, secondary index, and search index.

- General-purpose features
 - Offers schema-free data storage. You do not need to define attribute columns before you use them. Table-level changes are not required to add or delete attribute columns. You can configure the TTL parameter for a table to manage the lifecycle of data. Expired data is automatically deleted from the table.
 - Adopts a multi-node cluster architecture. The management nodes in the platform support a high-availability mechanism. Faults on daily O&M management nodes do not affect business operations.
 - Adopts the triplicate technology to store three copies of data on different racks. A cluster can support single storage type instances (SSD only) or mixed storage type instances (SSD and HDD) to meet different budget and performance requirements.
 - Adopts a fully redundant architecture that prevents single points of failure (SPOFs). Tablestore supports smooth online upgrades, hot cluster upgrades, and automatic data migration, which enable you to dynamically add or remove nodes for maintenance without incurring service interruptions. The concurrent read and write throughput and storage capacity can be linearly scaled. Each cluster can have at least 500 servers.

- Supports highly concurrent read and write operations. Concurrent read and write capabilities can be scaled up as the number of servers increases. The read and write performance is indirectly related to the amount of data in a single table.
- Supports identity authentication and multi-tenancy. Comprehensive access control and isolation mechanisms are provided to safeguard your data. VPC and access over HTTPS are supported. Tablestore provides multiple authentication and authorization mechanisms and RAM user management. This allows you to define access permissions on individual tables and operations.
- Table operations

You can list all tables in an instance, create a table, query and update the configurations of a table, and delete a table.
- Data operations

You can call the PutRow, GetRow, UpdateRow, and DeleteRow operations to perform operations on a single row or call the BatchWriteRow, BatchGetRow, and GetRange operations to perform operations on multiple rows. You can read and write data in a table by calling an operation to perform operations on a single row or multiple rows in the table.
- Data versions and TTL

You can use the max versions and TTL features to manage the lifecycle of your data to optimize storage efficiency and reduce costs.
- Auto-increment primary key column

You cannot set the partition key to an auto-increment primary key column. If you write data to a table that contains an auto-increment primary key column, you do not need to specify values for the auto-increment primary key column. Tablestore automatically generates values for the auto-increment primary key column. The values in the auto-increment primary key column are unique and increase monotonically within a partition that shares the same partition key value.
- Conditional update

If you use conditional update, data in the table can be updated only when the conditions are met. If the conditions are not met, the update fails.
- Filter

Filters sort results on the server side. Only results that meet the filter conditions are returned. Filters effectively reduce the volume of transferred data and shorten the response time.
- Tunnel Service

Tunnel Service is built on the Tablestore API to provide tunnels that are used to consume data in full, incremental, and differential modes. After you create a tunnel for a table, you can use the tunnel to consume historical and incremental data in the table.

Feature	Description
Tunnels for full and incremental data consumption	Tunnel Service supports incremental data consumption and allows you to concurrently consume full data and differential data.
Orderly incremental data consumption	Tunnel Service sequentially distributes incremental data to one or more logical partitions based on the write time. Data in different logical partitions can be concurrently consumed.

Consumption latency monitoring	Tunnel Service allows you to call the DescribeTunnel operation to view the recovery point objective (RPO) information about the consumed data on each client. Tunnel Service also allows you to monitor data that is consumed through tunnels in the Tablestore console.
Horizontal scaling of data consumption capabilities	Tunnel Service supports automatic load balancing among logical partitions to accelerate data consumption.

- Secondary index

The secondary index feature allows you to create one or more index tables for a data table. Then, you can query data based on the primary key columns of the index tables instead of the data table. This improves query performance.

Tablestore provides global secondary indexes and local secondary indexes to meet your requirements, such as strong query consistency.

Feature	Description
Single-column index and combined index	You can create an index on one or more columns in a data table.
Index synchronization	<p>Global secondary indexes and local secondary indexes synchronize data in different modes.</p> <ul style="list-style-type: none">◦ When you use the global secondary index feature, Tablestore automatically synchronizes the data from the index columns and primary key columns of a data table to the columns of an index table in asynchronous mode. The synchronization latency is within a few milliseconds.◦ When you use the local secondary index feature, Tablestore automatically synchronizes the data from the index columns and primary key columns of a data table to the columns of an index table in synchronous mode. After data is written to the data table, you can query the data from the index table.
Covering index	<p>Index tables can contain attribute columns. You can predefine several attribute columns as predefined columns when you create a data table. Then, you can create an index table based on predefined columns and all primary key columns of the data table. You can specify predefined columns as attribute columns in the index table. You can also specify no predefined columns as attribute columns in the index table.</p> <p>If you specify predefined columns of a data table as attribute columns in an index table, you can query the index table to obtain the values of the predefined columns. You do not need to query the data table.</p>
Index that contains existing data of a data table	You can create an index table that contains the existing data of a data table.
Sparse index	You can specify a predefined column of a data table as an attribute column of an index table. If a row in the data table does not contain the predefined column but contains all index columns, an index is created on the row. However, an index cannot be created on a row if the row does not contain all index columns.

- Search index

The search index feature is implemented by using inverted indexes and column stores. This feature provides query methods to resolve issues in complex big data scenarios. The query methods include queries based on non-primary key columns, full-text search, prefix query, fuzzy query, Boolean query, nested query, and geo query. Aggregation can be implemented by using the max, min, count, sum, avg, distinct_count, group_by, percentiles, and histogram functions.

Feature	Description
Search index management	After you create search indexes, you can query the description of the search indexes, query the search index list, and delete a search index.
Lifecycle management	The TTL can be configured for search indexes. TTL is an attribute of search indexes that specifies the retention period of data in search indexes. When data in a search index is retained for a period of time that exceeds the TTL value, Tablestore automatically deletes the data to free up storage space and reduce costs. If the UpdateRow operation is disabled for a data table, you can use the TTL feature of the search index that is created for the data table.
Types of date data	The search index feature can identify date type data in various formats. You can index dates stored as strings or integers to use in the search index. Queries made on dates take less time than queries made on strings in search indexes.
ARRAY and NESTED field types	<p>Search indexes provide the following special field types in addition to the basic field types such as LONG, DOUBLE, BOOLEAN, KEYWORD, TEXT, and GEOPOINT:</p> <ul style="list-style-type: none">◦ ARRAY: ARRAY is a type that can be combined with basic field types such as LONG, DOUBLE, BOOLEAN, KEYWORD, TEXT, and GEOPOINT. For example, the combination of LONG with ARRAY is used to specify arrays of the LONG INTEGER type. LONG ARRAY fields can contain multiple long integers. If a query matches a component of an array, the corresponding row is returned.◦ NESTED: NESTED indicates nested documents. Nested documents are used when a row of data (document) contains multiple child rows (child documents). Multiple child rows are stored in a NESTED field. You must specify the schema of child rows in the NESTED field. The schema must include the fields of the child rows and the property of each field.

Sorting and paging	You can predefine a sorting method when you create a search index or specify a sorting method when you use the search index to query data. This way, the rows that meet the query conditions are returned based on the order that you predefined or specified. If a large number of rows are included in the response, you can locate data by configuring the limit and offset parameters or by using tokens.
Tokenization	After you specify a tokenization method for TEXT fields, Tablestore tokenizes field values into multiple tokens based on the tokenization method. You cannot specify tokenization methods for non-TEXT fields. The following tokenization methods are supported: single-word tokenization, delimiter tokenization, minimum semantic unit-based tokenization, maximum semantic unit-based tokenization, and fuzzy tokenization.
Collapse (distinct)	You can use the collapse (distinct) feature to collapse the result set based on the specified column when the results of a query contain large amounts of data of a specific type. Data of the specified type is displayed only once in the query results to ensure the diversity of the result types.
Match all query	You can use the match all query feature to match all rows in a table to query the total number of rows in the table or return several random rows.
Match query	You can use the match query feature to query data in a table based on approximate matches. Tablestore tokenizes the values in TEXT columns and the keywords that you use to perform match queries based on the analyzer that you specify. This way, Tablestore can perform match queries based on the tokens. We recommend that you use the match phrase query feature for columns for which fuzzy tokenization is used to ensure high performance in fuzzy queries.
Match phrase query	The match phrase query feature is similar to match query, except that the match phrase query feature evaluates the positions of tokens. A row meets the query condition only if the order and positions of the tokens in the row match the order and positions of the tokens that are contained in the keyword. If the tokenization method for the column that you want to query is fuzzy tokenization, a match phrase query is faster than a wildcard query.
Term query	You can use the term query feature to query data that exactly matches the specified value of a field. A term query is similar to queries based on string match conditions. If the type of a field is TEXT, Tablestore tokenizes the string and exactly matches tokens.

Terms query	This query is similar to a term query. A terms query supports multiple terms. A row of data is returned if at least one of the keywords matches the field value. Terms queries can be used in the same manner as the IN operator in SQL statements.
Prefix query	You can use the prefix query feature to query data that matches a specific prefix. If the type of a field is TEXT, Tablestore tokenizes the string and matches tokens by using the specified prefix.
Range query	You can use the range query feature to query data that falls within the specified range. If the type of a field is TEXT, Tablestore tokenizes the string and matches any of the tokens that fall within the specified range.
Wildcard query	When you perform a wildcard query, you can use the asterisk (*) and question mark (?) wildcard characters in the query to search for data. The asterisk (*) matches a string of any length at, before, or after a search term. The question mark (?) matches a single character in a specific position. The string to match can start with an asterisk (*) or a question mark (?).
Boolean query	You can use the Boolean query feature to query the rows based on one or more subqueries. Tablestore returns the rows that match the subqueries. Each subquery can be of any type, including Boolean query.
Nested query	You can use the nested query feature to query the data in the child rows of nested fields. Nested fields cannot be directly queried. To query a nested field, you must specify the path of the nested field and a subquery in a NestedQuery object. The subquery can be a query of any type.
Geo-distance query	You can use the geo-distance query feature to specify a circular geographical area consisting of a central point and a radius as a filtering condition in a query. Tablestore returns the rows in which the value of a field falls within the circular geographical area.
Geo-bounding box query	You can use the geo-bounding box query feature to query data that falls within a rectangular geographic area. You can specify the rectangular geographic area as a query condition. Tablestore returns the rows in which the value of a field falls within the rectangular geographic area.

Geo-polygon query	You can use the geo-polygon query feature to query data that falls within a polygon geographic area. You can specify the polygon geographic area as a query condition. Tablestore returns the rows in which the value of a field falls within the polygon geographic area.
Exists query	An Exists query is also called a NULL query or NULL-value query. This query is used in sparse data to determine whether a column of a row exists. For example, you can query the rows in which the value of the address column is not empty.
Aggregation	You can perform aggregation operations to obtain the minimum value, maximum value, sum, average value, count and distinct count of rows, percentile statistics, and rows in each group. You can also perform aggregation operations to group results by field value, range, geographical location, filter, or histogram, and perform nested queries. You can perform multiple aggregation operations for complex queries.
Parallel scan	<p>The search index feature allows you to call the Search operation to query data, sort data in a specific order, and aggregate data.</p> <p>In some cases, a faster query speed may be more important than the order of query results. For example, the query speed matters when you connect Tablestore to a cluster computing environment such as Spark or Presto, or when you query a specific group of objects. To improve query speeds, Tablestore provides the ParallelScan operation for the search index feature. Compared with the Search operation, the ParallelScan operation supports all query features but does not provide analytics capabilities such as sorting and aggregation. This way, query speeds are improved by more than five times. You can call the ParallelScan operation to export hundreds of millions of data rows within a minute. The capability to export data can be horizontally scaled without upper limits.</p>
Virtual columns	<p>The virtual column feature allows you to map a column in a data table to a virtual column in a search index when you create the search index. The type of the virtual column can be different from that of the column in the data table. This allows you to create a column without the need to modify the table schema and data. The new column can be used to accelerate queries or can be configured with different analyzers.</p> <ul style="list-style-type: none">◦ You can configure different analyzers for a TEXT field. A single STRING column can be mapped to multiple TEXT columns of a search index. Different TEXT columns use different tokens to meet various business requirements.◦ You can accelerate queries by only mapping the required columns of a data table to the columns in a search index without the need to cleanse data or rebuild a table schema. The column types can differ between the data table and the search index. For example, you can map the numeric type to the KEYWORD type to improve the performance of a term query, and map the STRING type to the numeric type to improve the performance of a range query.

Dynamic schema modification	Data tables of Tablestore are schema-free. However, search indexes have rigid schemas. When you create a search index, you must specify the columns that you want to add to the search index. Then, you can query these columns when you use the search index to query data. You can dynamically modify the schema of a search index. For example, you can add, update, or delete index columns for the search index, and modify the routing keys of the search index.
Fuzzy query	<p>You can select a method to perform a fuzzy query based on your business requirements.</p> <ul style="list-style-type: none"> ◦ If you use *word* for a wildcard query, you can use fuzzy tokenization to perform a fuzzy query. For example, if you use "123" to query mobile numbers that contain 123 at any position, you can use fuzzy tokenization to perform a fuzzy query. ◦ For other complex queries, you can perform wildcard queries.

Simple Log Service

Simple Log Service is an all-in-one service that is developed by Alibaba Group based on extensive big data analytics scenarios. You can use Simple Log Service to collect, process, consume, query, and analyze log data without the need to invest in in-house data collection or processing systems. Simple Log Service helps improve O&M efficiency and allows you to process large amounts of data.

Feature	Description
Real-time log collection and consumption (LogHub)	LogHub allows you to collect logs without data loss by using various methods. These methods include clients, websites, protocols, SDKs, and API operations (for mobile applications and games). You can also consume logs by using SDKs, Storm Spout, and Spark Client. LogHub supports real-time log collection and consumption in multiple formats. You can use LogHub to streamline the collection and consumption of logs across multiple devices and sources.
Real-time log query and analysis (Search/Analytics)	You can use Simple Log Service to index, query, and analyze collected log data in real time. Simple Log Service can generate dynamic reports based on the query and analysis results. It can also generate visualized reports of log data in multiple scenarios.
Alert management	You can create an alert rule for the query and analysis results. After you create an alert rule, Simple Log Service periodically checks the related query and analysis results. If the query and analysis results meet the trigger condition that you specified in the alert rule, Simple Log Service sends an alert notification. This way, you can monitor the service status in real time.

Scheduled SQL	Simple Log Service provides the Scheduled SQL feature. You can use the feature to analyze data at a scheduled time and aggregate data for storage. You can also use the feature to project and filter data. You can view the background information, features, terms, usage scenarios, and usage notes of Scheduled SQL in the related topic.
Data processing	Simple Log Service provides the data transformation feature. This feature is fully hosted and provides high availability and scalability. You can use the data transformation feature to standardize, enrich, transfer, mask, and filter data.
Time series storage	Simple Log Service provides the time series storage feature. You can use the feature to collect, query, analyze, and visualize data.

-
- Scenarios

9.1.2. Benefits

Cloud Defined Storage (CDS) provides flexible deployment, high performance, and high reliability. When different cloud services are deployed in CDS, CDS inherits the benefits of the cloud services.

Overall advantages

Flexible deployment

CDS provides storage based on cloud services and allows you to deploy cloud services together or separately. In addition, CDS supports dynamic and flexible resource scaling.

High performance

- CDS can integrate resources such as CPU resources and hard disk resources on all storage nodes, and distribute data storage and process tasks in real time in a dynamic and balanced manner. This helps implement concurrent data processing, prevents issues caused by single points of failure (SPOFs), and improves the processing capabilities of clusters.
- The performance of CDS storage clusters can be easily improved to meet the growing storage requirements of applications.

High reliability

CDS storage clusters use the erasure coding mechanism to store data as data blocks on different servers in different racks. If a data block error occurs, you can easily restore the data block.

High availability

- CDS storage clusters use fully redundant architectures to prevent issues caused by SPOFs. In addition, CDS provides automatic failure detection and data migration features to shield applications from server-related and network-related hardware faults. This ensures the high availability of applications.
- The erasure coding mechanism used by storage clusters ensures proper data redundancy and improves space utilization.

High scalability

CDS improves the service capability of a storage cluster by using various methods, such as expanding the cluster, adding and upgrading server hardware in the cluster, and adding storage nodes to the cluster. CDS storage clusters support smooth online upgrades and hot upgrades, and allow you to dynamically add or remove storage nodes. In addition, the storage clusters support automatic data migration and do not require shutdown maintenance.

Access security

CDS provides multiple permission management mechanisms, and authenticates each request from applications to prevent unauthorized access. This ensures data security.

Easy management

- CDS frees you from complex O&M tasks, such as data partition management, software and hardware upgrades, configuration updates, and cluster scale-outs.
- CDS allows you to store audit logs to Simple Log Service and download logs from Simple Log Service. This facilitates the long-term storage and management of audit logs.
- The centralized O&M platform CDS Ops supports daily O&M operations on storage clusters and cloud services that are deployed in the storage clusters, such as Object Storage Service (OSS), Simple Log Service, and Elastic Block Storage (EBS).

Advantages of OSS

Advantages of OSS over self-managed storage

Item	OSS	Self-managed server storage
Reliability	<ul style="list-style-type: none">• Provides automatic backup for redundancy.• Tolerates faults at the hard disk, node, rack, and cluster levels. Read and write operations are not interrupted in the event of failures of up to two nodes. This ensures business continuity.	<ul style="list-style-type: none">• Is prone to errors due to low hardware reliability. If a disk has a bad sector, data may be irreversibly lost.• Requires manual restoration of data, which can be a complex, time-consuming, and labor-intensive process.
Security	<ul style="list-style-type: none">• Provides multi-level security protection for enterprises.• Provides resource isolation mechanisms for multiple tenants and supports zone-disaster recovery.• Provides various authentication and authorization mechanisms. It also provides features such as allowlists, hotlink protection, Resource Access Management (RAM), and Security Token Service (STS) for temporary access.	<ul style="list-style-type: none">• Requires additional scrubbing devices and blackhole policy-related services.• Requires a separate security mechanism.
Data processing	Provides Image Processing (IMG).	Requires separate purchase and deployment of data processing capabilities.

Other advantages of OSS

- Ease of use

- OSS provides standard RESTful API operations, some of which are compatible with Amazon S3 API operations, a wide range of SDKs, client tools, and the OSS console. You can use any one of these options to upload, download, query, and manage large amounts of data used in your apps and websites in the same way you would with regular file systems.
- OSS supports streaming writes and reads. It is suitable for business scenarios that require simultaneous write and read of large files such as videos.
- OSS supports lifecycle management. You can configure lifecycle rules to delete expired objects in batches.
- OSS provides plenty of storage space that is also scalable. You can add nodes to increase your storage space. A single bucket can contain trillions of objects.
- Powerful and flexible security mechanisms
OSS provides STS and URL-based authentication and authorization mechanisms, allowlists, hotlink protection, and RAM.
- Rich image processing features
OSS supports the conversion between formats such as JPG, PNG, BMP, GIF, WebP, and TIFF. OSS also supports various of operations on image objects, such as thumbnails, cropping, watermarking, and resizing.

Advantages of EBS

EBS provides devices based on a distributed storage architecture.

EBS provides ECS instances with block-level storage that features low latency and high performance, durability, and reliability. EBS devices use a triplicate distributed storage mechanism to ensure data durability for ECS instances. You can create, release, and resize EBS devices at any time.

You can resize EBS devices without interrupting your business, including system disks and data disks. When you resize an EBS device, you do not need to stop the Elastic Compute Service (ECS) instance to which the EBS device is attached or detach the EBS device from the ECS instance.

Advantages of Tablestore

Scalability

- Tablestore does not impose any limits on the amount of data that can be stored in tables. As data increases, Tablestore adjusts data partitions to provide more storage space for tables and improve the capability of handling sudden spikes of access requests.
- Tablestore supports CPUs, disks, memory, and network interface controllers (NICs) of different specifications in a single-component cluster without affecting cluster running performance. This ensures maximum compatibility with existing devices.

High performance

High-performance Tablestore instances provide single-digit millisecond latency when you access single rows of data. The read/write performance is not affected by the size of data in a table.

Data reliability

- Tablestore provides high data reliability. It stores multiple copies of data and restores data when any of the copies become damaged.
- Tablestore supports automatic fault tolerance for server disk failures in a cluster and supports hot swapping of disks. In the event of a disk failure, services can be restored within a minute.
- Tablestore supports full and incremental backup and data restoration from storage.
- Tablestore supports the backup between data clusters in different data centers. You can

view and manage the backup process.

- Tablestore supports the backup and restoration of the metadata, files, and tables of key components.

High availability

Tablestore uses automatic failure detection and data migration to shield applications from host- and network-related hardware faults, providing high availability for your applications.

Easy management

- Tablestore automatically performs complex O&M tasks, such as the management of data partitions, software and hardware upgrades, configuration updates, and cluster scale-out.
- You can store audit logs in Log Service and download logs from Log Service. This facilitates the long-term storage and management of audit logs.

Access security

- Tablestore provides multiple permission management mechanisms. It verifies and authenticates the identity of each application request to prevent unauthorized data access, which improves data security.
- Tablestore supports the management of data access permissions, including logon permissions, table creation permissions, read and write permissions, and whitelist-related permissions.
- Tablestore allows you to use the Apsara Uni-manager Management Console to manage administrative permissions, including administrator classification. You can use the console to manage user permissions in a centralized manner. You can manage the access control features of all components in the system. You can also block regular users from querying access control details and simplify access control for administrators. This improves the usability of access control.

Strong consistency

Tablestore ensures high data consistency for data writes. After a write operation succeeds, three replicas are written to a disk. Applications can read the latest data immediately.

Flexible data models

Tablestore tables do not require a rigid schema. Each row can contain a different number of columns. Tablestore supports multiple data types, including Integer, Boolean, Double, String, and Binary.

Powerful data indexing capabilities

Tablestore provides powerful data query features, such as primary key-based queries, secondary index-based queries, and search index-based queries.

- Secondary index: defines a custom data structure that can be used to improve query efficiency based on your business requirements.
- Search index: supports query methods such as boolean queries, fuzzy queries, geo queries, and full-text searches based on inverted indexes and column-oriented storage.

Monitoring service integration

You can log on to the Tablestore console to obtain monitoring information in real time, including the number of requests per second and the average response latency.

Multi-tenant management

- Isolation: allows tasks of multiple tenants to be submitted to different queues and run separately. Resources are isolated among tenants.
- Permission: allows you to manage tenants in a centralized manner, dynamically configure and manage tenant resources, isolate resources, view statistics for resource usage, and manage tenants at multiple levels in the console.
- Scheduling: supports multi-tenant scheduling of multiple clusters and multiple resource

pools.

Advantages of Simple Log Service

Fully managed service

- Log Service is easy to access and easy to use.
- LogHub provides all features of Kafka, provides monitoring and alerting data, and supports auto scaling (by petabytes per day).
- LogSearch/Analytics provides the saved search feature, and allows you to view log data on dashboards and configure alerts.
- Log Service provides more than 30 access methods to connect to open source software, such as Storm and Spark Streaming.

Comprehensive ecosystem

- LogHub supports more than 30 types of data sources, including embedded devices, web pages, servers, and programs. LogHub can connect to different consumption systems, such as Storm and Spark Streaming.
- LogSearch/Analytics supports complete query and analysis syntax and is compatible with SQL-92 syntax. Log Service can be connected to Grafana by using Java Database Connectivity (JDBC).

Real-time response

- LogHub: Data can be immediately consumed after the data is written to Log Service. Logtail is used as an agent to collect and send data to Log Service in real time.
- LogSearch/Analytics: Data can be queried 3 seconds after the data is written to Log Service. If you specify multiple conditions to query data from hundred billions of data records, the result can be returned in seconds.

High throughput

Log Service provides high throughput after Log Service software is optimized. If Log Service is deployed on a single server, the throughput performance can reach the following levels:

- Data write: Raw logs can be written to Simple Log Service at a speed that exceeds 400 MB/s, and indexed logs can be written to Simple Log Service at a speed that exceeds 150 MB/s.
- Real-time consumption: Data can be consumed in real time at a speed that exceeds 400 MB/s.
- QPS: The queries per second (QPS) can exceed 10,000.

 **Note** The throughput performance varies based on the hardware of a server.

Scalability

Simple Log Service provides flexible scaling capabilities to process petabytes of data.

9.1.3. Scenarios

Cloud Defined Storage (CDS) can provide one-stop storage solutions for enterprises in different industries in a cost-efficient, secure, and reliable manner. Enterprises can deploy different services in CDS to meet requirements in different scenarios.

OSS

High storage capacity for image, audio, and video applications

OSS can be used to store large amounts of data, such as images, audio and video data, and logs. Various devices, websites, and mobile applications can directly write data to and read data from OSS. You can write data to OSS by uploading files or using streams.

Offline data storage

OSS is a cost-effective storage service that offers high data availability. You can use OSS to store enterprise data that needs to be archived offline for a long period of time.

Cross-region disaster recovery of data

You can use cross-region replication or cross-cloud replication to asynchronously replicate your data between two clusters or clouds in near real time. This way, you can build a storage architecture with three data centers across two zones and store your data in different regions for backup and disaster recovery. This ensures the continuity of your business even in case of severe disaster events.

EBS

Elastic Block Storage (EBS) is a persistent random block storage service that provides low latency and high reliability and is designed for Elastic Compute Service (ECS). EBS provides devices based on a distributed storage architecture. EBS devices are classified into the following types based on whether they can be attached to multiple ECS instances: disks and shared disks.

Disks

Disks are classified by performance into the following types: premium performance disks, standard performance disks, ultra disks, and standard SSDs.

- Standard performance disks and premium performance disks are ideal for online transaction processing (OLTP) databases (such as MySQL, PostgreSQL, Oracle, and SQL Server), NoSQL databases (such as MongoDB, HBase, and Cassandra), and Elasticsearch, Logstash, and Kibana (ELK) distributed logs. Premium performance disks deliver up to 25,000 random IOPS for ECS instances.
- Ultra disks are ideal for medium I/O load scenarios and deliver up to 3,000 random IOPS for ECS instances. Ultra disks can be used as system disks of ECS instances.
- Standard SSDs are ideal for I/O-intensive applications and deliver stable and high random IOPS performance. Standard SSDs can be used in small and medium-sized development and test environments that require high data reliability.

Tablestore

Scenario 1: Big data storage and analytics

Tablestore provides cost-effective, highly concurrent, and low-latency storage for large amounts of data, and online access to the data. It provides full and incremental data tunnels and supports direct SQL-based read and write operations for various big data analytics platforms such as MaxCompute. An efficient incremental streaming read operation is provided for easy processing of real-time data streams.

- Tablestore supports various big data computing platforms, streaming processing services, and real-time computing services.
- Tablestore provides high performance and capacity instances to meet the requirements of different business.

Scenario 2: Social media feeds on the Internet

You can use Tablestore to store large amounts of instant messaging (IM) messages and social media feed information such as comments, posts, and likes. The elastic resources available for Tablestore can meet application requirements including handling significant traffic fluctuations, high concurrency, and low latency at relatively low costs.

- Built-in auto-increment primary key columns reduce the number of external system

dependencies.

- Average read and write performance of high performance instances are not affected by data volumes.
- Highly reliable storage for large amounts of messages, and multi-terminal message synchronization are supported.

Scenario 3: Storage and real-time queries of large amounts of transaction records and user models

Tablestore instances are elastic, low latency, and highly concurrent, which provides optimal running conditions for risk control systems. This helps you control transaction risks. Furthermore, the flexible data structure allows your business model to rapidly evolve to meet market demands.

- A table can store full historical transaction records.
- Data is stored in three copies to ensure high consistency and data security.
- The schema-free data model allows you to add attribute columns based on your requirements. This allows rapid business development.

Scenario 4: Efficient and flexible storage of large amounts of IoV data

The schema-free data model eliminates the need for table sharding and simplifies access to the data collected from different vehicle-mounted devices. Tablestore can be seamlessly connected to multiple big data analytics platforms and real-time computing services to implement real-time online queries and business report analysis.

- Data is stored in a table without sharding, which simplifies business logic.
- The query performance for vehicle conditions and recommended routes is stable and predictable.
- The schema-free model allows you to store data collected from different vehicle-mounted devices.

Scenario 5: Storage of large amounts of IoT data for efficient queries and analysis

Tablestore can be used to store time series data from IoT devices and monitoring systems. It supports direct SQL-based read operations for big data analytics and provides API operations to read incremental data streams, which allow you to implement offline data analysis and real-time streaming processing.

- Tablestore can meet the data write and storage requirements of ultra-large-scale IoT devices and monitoring systems.
- Tablestore can connect to a variety of offline or stream data analysis platforms. This allows you to use a single piece of data for multiple analysis and computing purposes.
- Tablestore supports the lifecycle management of data.

Scenario 6: Databases for large-scale e-commerce transaction orders and user-specific recommendations

Tablestore can help manage large amounts of historical transaction data and improve access performance. Tablestore can be used together with MaxCompute to implement precision marketing. The elastic resources available for Tablestore allow Tablestore to handle requests during peak hours when all users go online.

- Resources can be scaled based on data volumes and access concurrency, which allows the service to handle scenarios that feature high access fluctuations during various periods.
- Various big data analytics platforms are supported for direct analysis of user behaviors.
- Single-digit millisecond latency for queries on large amounts of transaction orders is supported.

Simple Log Service

Simple Log Service is applicable to the following scenarios: data collection, real-time computing, data warehousing and offline analysis, product operation and analysis, O&M, and management.

Data collection and consumption

You can use the LogHub module of Simple Log Service to collect large amounts of log data in real time. The log data can be metrics, events, binary logs, text logs, and clickstream data.

- Easy to use: More than 30 real-time data collection methods are provided for you to quickly set up your platform and reduce O&M workload.
- Automatically scalable: Log Service scales based on traffic and business requirements, helping you handle traffic spikes and respond to growing business demands.

ETL and stream processing

You can connect LogHub to multiple real-time computing engines and services. LogHub can monitor the processing progress and generate alerts based on the monitoring results. You can also use SDKs or call API operations to consume logs.

- Easy operations: LogHub provides SDKs in multiple programming languages and programming frameworks. It can interconnect with various stream computing engines.
- Comprehensive features: LogHub supports alert mechanisms and provides large amounts of monitoring data.
- Elastic scaling: PB-grade elasticity and zero latency.

Real-time query and analysis

The LogAnalytics module allows you to index LogHub data in real time and query data by using keywords, fuzzy match, context, or SQL aggregate functions. You can also query data within the specified range.

- Strong real-timeliness: Data can be queried immediately after it is written.
- High efficiency at low cost: LogSearch/Analytics is able to index PBs of data each day. Costs are 85% lower compared with self-built systems.
- Strong analysis capability: LogSearch/Analytics supports multiple query methods and SQL for aggregation analysis. It also provides visualization and alerting capabilities.

9.2. Simple Log Service

Simple Log Service is a cloud-native observability and analytics platform that provides large-scale, low-cost, and real-time services to process multiple types of data such as logs, traces, and metrics. Simple Log Service allows you to collect, transform, query, analyze, visualize, ship, and consume data. You can also configure alert rules in the console. Simple Log Service helps you improve digital capabilities in R&D, O&M, and data security.

Simple Log Service provides the following features:

- Data collection: You can collect log data, trace data, or metric data from more than 50 types of data sources. These data sources include Alibaba Cloud services, servers, applications, IoT devices, mobile terminals, and open source software. You can also collect data that is transferred over standard protocols.
- Data query and analysis: Simple Log Service allows you to query and analyze petabytes of data in real time. This feature supports more than 10 operators, more than 10 machine learning functions, and more than 100 SQL functions. Simple Log Service also provides the Scheduled SQL feature and Dedicated SQL feature.
- Data transformation: Simple Log Service provides more than 200 built-in functions, more than 400 regular expressions, and flexible custom functions to filter, split, convert, enrich, and replicate data. This feature meets your business requirements for multiple scenarios, such as data distribution, data standardization, and data integration.
- Visualization: Simple Log Service allows you to visualize query and analysis results. Simple

Log Service provides more than 10 types of charts, such as tables, line charts, bar charts, and maps. You can customize dashboards based on charts. You can also perform embedded analysis and drill-down analysis.

- Alerting: Simple Log Service can automatically run query statements at regular intervals after you create an alert rule. If the query results meet the trigger conditions of the alert rule, Simple Log Service sends an alert notification to the specified recipients in real time.
- Data consumption and shipping: Simple Log Service allows you to consume data in real time by using Storm, Flume, or Flink. You can also ship data to Alibaba Cloud services such as Object Storage Service (OSS) in real time.

9.2.1. Features

Data collection

Simple Log Service allows you to collect the following types of data from more than 50 data sources:

- Logs, traces, and metrics from servers and applications
- Logs from IoT devices
- Logs from Alibaba Cloud services
- Data from mobile terminals
- Data from open source software such as Logstash, Flume, Beats, FluentD, and Telegraph
- Data transferred over standard protocols such as HTTP, HTTPS, Syslog, Kafka, and Prometheus

Data query and analysis

Simple Log Service provides the following features that allow you to query and analyze data in real time:

- Simple Log Service supports exact search, fuzzy search, full-text search, and field search.
- Simple Log Service supports features such as contextual query, LogReduce, LiveTail, and reindexing.
- Simple Log Service supports the SQL-92 syntax.
- Simple Log Service provides the Dedicated SQL feature.

Data transformation

Simple Log Service provides the data transformation feature to help you standardize, enrich, transfer, mask, and filter data.

- Data standardization: Simple Log Service can extract fields from logs in different formats and convert the log formats to obtain structured data for stream processing and computing in data warehouses.
- Data enrichment: Simple Log Service can join the fields of logs and dimension tables to link logs with dimension information, which facilitates data analysis. For example, Simple Log Service can join the fields of order logs and a user information table.
- Data transfer: Simple Log Service can transfer logs from regions outside China to one region by using the global acceleration feature. This way, global logs can be managed in a centralized manner.
- Data masking: Simple Log Service can mask sensitive information that is contained in data. The sensitive information includes passwords, mobile phone numbers, and addresses.
- Data filtering: Simple Log Service can filter logs to obtain key service logs. This helps further analysis.

Data consumption and shipping

Simple Log Service provides the data consumption feature and data shipping feature. You can use an SDK or call an API operation to consume data in real time. You can also use the Simple Log Service console to ship data to Alibaba Cloud services such as Object Storage Service (OSS) in real time.

- You can consume data by using third-party software, such as Splunk, QRadar, Logstash, and Flume.
- You can consume data by using different programming languages, such as Java, Python, and Go.
- You can consume data by using Alibaba Cloud services, such as Function Compute and Realtime Compute for Apache Flink.
- You can consume data by using different stream processing platforms, such as Apache Flink, Apache Spark, and Apache Storm.
- You can ship data to OSS.

Visualization

Simple Log Service allows you to use the following methods to visualize query and analysis results:

- Built-in charts on dashboards: Simple Log Service provides various statistical charts, such as tables, line charts, and column charts. You can select chart types to visualize query and analysis results on a dashboard and save the results to the dashboard.
- Third-party visualization tools: Simple Log Service is compatible with third-party visualization tools, such as Grafana and DataV.

Alerting

Simple Log Service allows you to configure alert rules for query and analysis results. After you create an alert rule, Simple Log Service periodically checks the related query and analysis results. If the query and analysis results meet the trigger condition that you specified in the alert rule, Simple Log Service sends an alert notification. This way, you can monitor the service status in real time.

9.2.2. Benefits

Unified

Log Service allows you to collect data from multiple types of data sources.

Intelligent

Log Service provides AIOps capabilities to intelligently detect exceptions and analyze root causes.

Efficient

Log Service can collect, query, and analyze hundreds of billions of logs in real time.

End-to-end

Log Service allows you to collect, transform, query, analyze, and visualize data from end to end. You can also configure alerts for the data.

Elastic

Log Service provides automatic scaling capabilities to process petabytes of data.

Cost-effective

Log Service provides visualized and automated installation, deployment, and scaling. The high-availability architecture and service self-check feature significantly reduce O&M costs.

9.2.3. Scenarios

Data collection and consumption

You can use the LogHub module of Simple Log Service to collect large amounts of log data in real time. The log data can be metrics, events, binary logs, text logs, and clickstream data.

Benefits:

- **Ease of use:** Simple Log Service supports more than 50 data sources to help you build platforms. Simple Log Service also provides powerful configuration and management capabilities to reduce your O&M workloads.
- **Elastic scalability:** Simple Log Service can handle traffic spikes and business growth.

Real-time query and analysis

The LogAnalytics module can index the data in LogHub in real time and provides multiple query methods such as keyword-based search, fuzzy search, contextual query, data query within a specific time range, and SQL-based aggregation. Simple Log Service also provides the Scheduled SQL feature and Dedicated SQL feature.

- **Timeliness:** You can perform real-time query after data is written to LogHub.
- **High efficiency and low costs:** You can index petabytes of data each day. Costs are 85% lower compared with self-managed systems.
- **Powerful analysis:** Simple Log Service supports multiple query methods and SQL-based aggregate functions. Simple Log Service also provides visualized reports and allows you to configure alerts.

Data transformation

Simple Log Service provides the data transformation feature to help you standardize, enrich, transfer, mask, and filter data.

- **Data standardization:** Simple Log Service can extract fields from logs in different formats and convert the log formats to obtain structured data for stream processing and computing in data warehouses.
- **Data enrichment:** Simple Log Service can join the fields of logs and dimension tables to link logs with dimension information, which facilitates data analysis. For example, Simple Log Service can join the fields of order logs and a user information table.
- **Data transfer:** Simple Log Service can transfer logs from regions outside China to one region by using the global acceleration feature. This way, global logs can be managed in a centralized manner.
- **Data masking:** Simple Log Service can mask sensitive information that is contained in data. The sensitive information includes passwords, mobile phone numbers, and addresses.
- **Data filtering:** Simple Log Service can filter logs to obtain key service logs. This helps further analysis.

ETL and stream processing

You can connect LogHub to multiple real-time computing engines and services. LogHub can monitor the processing progress and generate alerts based on the monitoring results. You can also use SDKs or call API operations to consume logs.

- **Ease of use:** LogHub provides SDKs in multiple programming languages and frameworks. You can connect LogHub to various stream computing engines.

- Comprehensive features: LogHub can monitor large amounts of data and generate alerts based on the monitoring results.
- Elastic scalability: Simple Log Service supports auto scaling to process petabytes of data without latency.

Data shipping (LogShipper)

The LogShipper module of Simple Log Service can ship log data to Object Storage Service (OSS). During the shipping process, you can compress the data, define custom partition formats, and specify row store or column store.

- Large data capacity: An unlimited amount of data can be shipped.
- Multiple formats: Various storage formats such as row store, column store, and text files are supported.
- Flexible configurations: Different configurations are supported, which allows you to define custom partition formats.

10. Network services

10.1. What is SLB?

This topic provides an overview of Server Load Balancer (SLB). SLB distributes inbound network traffic across multiple Elastic Compute Service (ECS) instances that function as backend servers based on forwarding rules. You can use SLB to improve the responsiveness and availability of your applications.

Overview

After you attach ECS instances that are deployed in the same region to an SLB instance, SLB uses virtual IP addresses (VIPs) to virtualize these ECS instances into backend servers in a high-performance server pool that ensures high availability. Client requests are distributed to the ECS instances based on forwarding rules.

SLB checks the health status of the ECS instances and automatically removes unhealthy ones from the server pool to eliminate single points of failure (SPOFs). This enhances the resilience of your applications.

Components

SLB consists of three components:

- SLB instances

An SLB instance is a running SLB service entity that receives traffic and distributes traffic to backend servers. To get started with SLB, you must create an SLB instance, configure at least one listener for the SLB instance, and attach at least two ECS instances to the SLB instance.

- Listeners

A listener checks client requests and forwards them to backend servers. It also performs health checks on backend servers.

- Backend servers

ECS instances are used as backend servers to receive distributed requests. You can add ECS instances one by one to the server pool, or use vServer groups or primary/secondary server groups to add and manage multiple ECS instances at a time.

Benefits

- High availability

SLB is designed with full redundancy that prevents SPOFs and supports zone-disaster recovery.

SLB can be scaled based on application loads and can provide continuous services during traffic fluctuations.

- High scalability

You can increase or decrease the number of backend servers to adjust the load balancing capability of your applications.

- High cost efficiency

SLB can save 60% of load balancing costs compared with traditional hardware solutions.

- Outstanding security

You can integrate SLB with Apsara Stack Security to defend your applications against DDoS attacks of up to 5 Gbit/s.


- High concurrency

An SLB cluster supports hundreds of millions of concurrent connections. A single SLB instance supports tens of millions of concurrent connections.

10.1.1. Features

This topic describes the features of Apsara Stack Server Load Balancer (SLB). SLB provides load balancing at Layer 4 and Layer 7. It supports features such as health checks, session persistence, and domain name-based forwarding rules to ensure the high availability of backend services.

In the following table, Y indicates that the feature is supported, and N indicates that the feature is not supported.

Feature	Layer 4 SLB	Layer 7 SLB
<p>Scheduling algorithms</p> <p>SLB supports the round-robin (RR), weighted round-robin (WRR), and consistent hashing scheduling algorithms.</p>	Y	Y <div> Note Layer 7 SLB does not support the consistent hashing scheduling algorithm.</div>
<p>Health checks</p> <p>SLB checks the health status of backend servers. When an unhealthy backend server is detected, SLB stops distributing inbound traffic to the backend server. Network traffic is distributed to other backend servers that work as expected.</p>	Y	Y
<p>Session persistence</p> <p>SLB supports session persistence. After session persistence is enabled, SLB can distribute requests from a client during a session to the same backend server.</p>	Y	Y
<p>Access control</p> <p>SLB uses whitelists to control access to your applications.</p>	Y	Y
<p>High availability</p> <p>SLB distributes inbound traffic to backend servers that are deployed in different zones. In most regions, Apsara Stack allows you to deploy SLB instances in primary/secondary zone mode across multiple zones. If the primary zone fails, a failover is triggered to redirect requests to servers in the secondary zone.</p>	Y	Y
<p>Security</p> <p>You can integrate SLB with Apsara Stack Security to defend your applications against DDoS attacks of up to 5 Gbit/s.</p>	Y	Y

Network types Apsara Stack provides Internet-facing and internal-facing SLB instances. To process network traffic within a virtual private cloud (VPC), you can create an internal-facing SLB instance. To process network traffic from the Internet, you can create an Internet-facing SLB instance.	Y	Y
IPv6 Internet-facing SLB instances can forward requests from IPv6 clients.	Y	Y
Certificate management SLB manages certificates for HTTPS in a centralized way. You do not need to upload certificates to backend servers. Requests are decrypted on SLB instances before the requests are sent to backend servers. This reduces the CPU utilization on backend servers.	N	Y
WebSocket Secure and WebSocket WebSocket is an HTML5 protocol that provides full-duplex communication channels between clients and servers. You can use WebSocket to save server resources and bandwidth, and enable real-time communication.	N	Y
HTTP/2 HTTP/2 is the second major version of the HTTP protocol and is backward compatible with HTTP/1.x. In addition, HTTP/2 improves performance by optimizing the flow of content.	N	Y

10.1.2. High availability

This topic describes the high availability of Server Load Balancer (SLB). SLB provides a high-availability architecture based on system design and product configurations. To meet your business requirements, you can use SLB together with services such as Apsara Stack DNS to implement geo-disaster recovery.

High availability of SLB

SLB instances are deployed in clusters to synchronize sessions and protect backend servers from single points of failure (SPOFs). This improves redundancy and ensures service stability. Layer 4 SLB uses the Linux Virtual Server (LVS) and Keepalived software to balance loads, whereas Layer 7 SLB uses Tengine. Tengine, a web server project launched by Taobao, is based on NGINX and adds advanced features that are designed for high-traffic websites.


Requests from the Internet are forwarded to LVS clusters based on equal-cost multi-path (ECMP) routing. In an LVS cluster, each LVS machine uses multicast packets to synchronize sessions with other LVS machines. This way, sessions are synchronized among all machines in the LVS cluster. LVS clusters also perform health checks on Tengine clusters. To ensure the availability of Layer 7 SLB, unhealthy devices are removed from Tengine clusters.

Best practices:

You can use session synchronization to prevent persistent connections from being affected by server failures within a cluster. However, for short-lived connections, server failures in the cluster may affect user access. This also occurs when a connection does not trigger session synchronization because the three-way handshake is not complete. To prevent session interruptions due to server failures within the cluster, you can add a retry mechanism to the service logic. This reduces the impact of server failures on user access.

High-availability solution with one SLB instance

Apsara Stack allows you to deploy SLB instances across multiple zones in different regions. You can deploy an SLB instance in primary/secondary zone mode. This ensures the high availability of the SLB instance. If the primary zone fails or becomes unavailable, a failover is triggered to redirect requests to servers in the secondary zone in about 30 seconds. After the primary zone recovers, traffic is automatically switched back to servers in the primary zone.

 **Note** Zone-disaster recovery is implemented between the primary and secondary zones. SLB implements failovers only when the entire SLB cluster within the primary zone is unavailable due to factors such as power outage and optical cable failures. A failover is not triggered when a single instance in the primary zone fails.

Best practices:

1. We recommend that you create SLB instances in regions that support primary/secondary zone deployment for zone-disaster recovery.
2. You can determine the primary and secondary zones for an SLB instance based on the distribution of Elastic Compute Service (ECS) instances. Select the zone where most ECS instances are deployed as the primary zone to minimize access latency.


However, we recommend that you do not deploy all ECS instances in the primary zone. You must deploy a small number of ECS instances in the secondary zone. This way, requests can be redirected to backend servers in the secondary zone when the primary zone becomes unavailable.

High-availability solution with multiple SLB instances

If you require extremely high availability, the high-availability solution with one SLB instance may be insufficient for your needs. If an SLB instance becomes unavailable due to network attacks or invalid configurations, failovers between the primary and secondary zones are not triggered. To avoid such issues, you can create multiple SLB instances. Then, you can use Apsara Stack DNS to schedule requests, or use a global load balancing solution to achieve cross-region backup and disaster recovery.

Best practices:

You can deploy SLB instances and ECS instances in multiple zones within a region or across multiple regions, and schedule requests by using Apsara Stack DNS.



High-availability solution with backend ECS instances

SLB performs health checks to verify the availability of backend ECS instances. The health check feature improves the availability of frontend services by minimizing the impacts of ECS instance exceptions on the services.

After you enable the health check feature, when an unhealthy ECS instance is detected, SLB distributes new requests to other healthy ECS instances. After the ECS instance recovers and is confirmed healthy, SLB automatically sends requests to the ECS instance. For more information about the health check feature, see [Health check overview](#) in User Guide.

Best practices:

To perform health checks, make sure that the health check feature is enabled and properly configured. For more information, see [Configure health checks](#) in User Guide.

10.1.3. Scenarios

Server Load Balancer (SLB) can be used to improve the availability and reliability of applications with high access traffic.

Balance the loads of your applications

You can configure listening rules to distribute heavy traffic among Elastic Compute Service (ECS) instances that are attached as backend servers to SLB instances. You can also use the session persistence feature to forward all requests from the same client to the same backend ECS instance to enhance access efficiency.

Scale your applications

You can add or remove backend ECS instances to scale the service capability of your applications based on your business requirements. SLB is applicable to both web servers and application servers.

Eliminate single points of failure

You can attach multiple ECS instances to an SLB instance. When ECS instances malfunction, SLB automatically isolates these ECS instances and distributes inbound requests to other healthy ECS instances. This ensures that your applications continue to run as expected.

Implement zone-disaster recovery (multi-zone disaster recovery)

To provide more stable and reliable load balancing services, Apsara Stack allows you to deploy SLB instances across multiple zones in different regions for disaster recovery. You can deploy an SLB instance in primary/secondary zone mode. If the primary zone fails or becomes unavailable, a failover is triggered to redirect requests to servers in the secondary zone in about 30 seconds. After the primary zone recovers, traffic is automatically switched back to servers in the primary zone. We recommend that you plan the deployment of backend servers based on your business requirements. We recommend that you add at least one ECS instance in each zone to achieve the highest load balancing efficiency.

ECS instances in different zones can be attached to an SLB instance, as shown in the following figure. In most cases, the SLB instance distributes inbound traffic to ECS instances in the primary zone (Zone A) and those in the secondary zone (Zone B). If Zone A fails, the SLB instance distributes inbound traffic only to ECS instances in Zone B. This deployment mode avoids service interruptions due to zone-level failures and reduces latency.

For example, you deploy all ECS instances in the primary zone (Zone A), and no ECS instances are deployed in the secondary zone (Zone B), as shown in the following figure. If the primary zone fails, your services are interrupted because no ECS instances are available in the secondary zone. This deployment mode achieves low latency but mitigates high availability.

Cross-region disaster recovery

You can deploy SLB instances in different regions and attach ECS instances of different zones within the same region to an SLB instance. You can use Apsara Stack DNS to resolve domain names to the endpoints of SLB instances in different regions for global load balancing purposes. When a region becomes unavailable, you can temporarily stop DNS resolution in the region without affecting user access.


10.2. Virtual private cloud

A virtual private cloud (VPC) is a private network in the cloud. A VPC provides network services for resources in the cloud. Different VPCs are logically isolated from each other. You have full control over your VPC. For example, you can specify CIDR blocks, and configure route tables and gateways for your VPC. In your VPCs, you can use Alibaba Cloud resources, such as Elastic Compute Service (ECS) instances, ApsaraDB RDS instances, and Server Load Balancer (SLB) instances.

10.2.1. What is a VPC?


A virtual private cloud (VPC) is a private network in the cloud. You can configure the CIDR block, route tables, and gateways of your VPC. You can use Alibaba Cloud services in a VPC, such as Elastic Compute Service (ECS), Server Load Balancer (SLB), and ApsaraDB RDS.

You can connect your VPC to other VPCs or on-premises networks to create a custom network environment. This way, you can migrate applications to the cloud and extend data centers.



Components

Each VPC consists of one vRouter, at least one private CIDR block, and at least one vSwitch.



- Private CIDR blocks

When you create a VPC and a vSwitch, you must specify the private IP address range for the VPC in CIDR notation.

You can use the standard private CIDR blocks listed in the following table and their subsets as CIDR blocks for your VPCs. For more information, see [Plan networks](#). For more information, see the **Plan and design a VPC** topic in **User Guide**.

CIDR block	Number of available private IP addresses (excluding system reserved IP addresses)
192.168.0.0/16	65,532
172.16.0.0/16	65,532

- vRouters

A vRouter is the hub of a VPC. As a core component, it connects the vSwitches in a VPC and serves as a gateway between a VPC and other networks. After you create a VPC, the system automatically creates a vRouter. Each vRouter is associated with a route table.

For more information, see [Route table overview](#).

For more information, see the **Route table overview** topic in **User Guide**.

- vSwitches

A vSwitch is a basic network component that connects different cloud resources in a VPC. After you create a VPC, you can create vSwitches to create one or more subnets for the VPC. vSwitches in the same VPC can communicate with each other. You can deploy your applications in vSwitches that belong to different zones to improve service availability.

For more information, see [vSwitches](#).

For more information, see the **Create a vSwitch** topic in **User Guide**.

10.2.2. Benefits

This topic describes the benefits of virtual private clouds (VPCs). VPCs are secure, reliable, flexible, easy to use, and scalable.

Security and reliability

Each VPC is identified by a unique tunnel ID, which corresponds to a virtual network. Different VPCs are isolated by tunnel IDs:

- Similar to a traditional network, you can create vSwitches and vRouters to divide a VPC into multiple subnets. Elastic Compute Service (ECS) instances in the same subnet use the same vSwitch to communicate with each other, while ECS instances in different subnets use vRouters to communicate with each other.
- VPCs are completely isolated from each other. Cloud resources in different VPCs can communicate with each other by using elastic IP addresses (EIPs) or NAT IP addresses.
- The IP packets of an ECS instance are encapsulated by using the tunneling technology. Therefore, information at the data link layer (the MAC address) of the ECS instance is not transferred to the physical network. This way, ECS instances in different VPCs are isolated at Layer 2.
- ECS instances in a VPC use security groups and firewalls to control inbound and outbound traffic at Layer 3.

Flexible management

You can use security group rules and access control lists (ACLs) to manage inbound and outbound traffic to cloud resources in a VPC in a flexible manner.

Ease of use

You can easily create and manage VPCs in the VPC console. When you create a VPC, the system automatically creates a vRouter and a route table for the VPC.

High scalability


You can create different subnets in a VPC to deploy different services. Additionally, you can connect a VPC to a data center or another VPC to extend the network architecture.

10.2.3. Use scenarios

Virtual private clouds (VPCs) are virtual networks that are isolated from each other. VPCs support flexible configurations to meet the requirements of different scenarios.


Deploy applications in a safe manner

You can deploy applications in a VPC to provide services to external networks. To control access to the applications over the Internet, you can create security group rules and configure whitelists. You can also isolate application servers from databases to implement access control. For example, you can deploy web servers in a subnet that can access the Internet, and deploy databases in another subnet that cannot access the Internet.



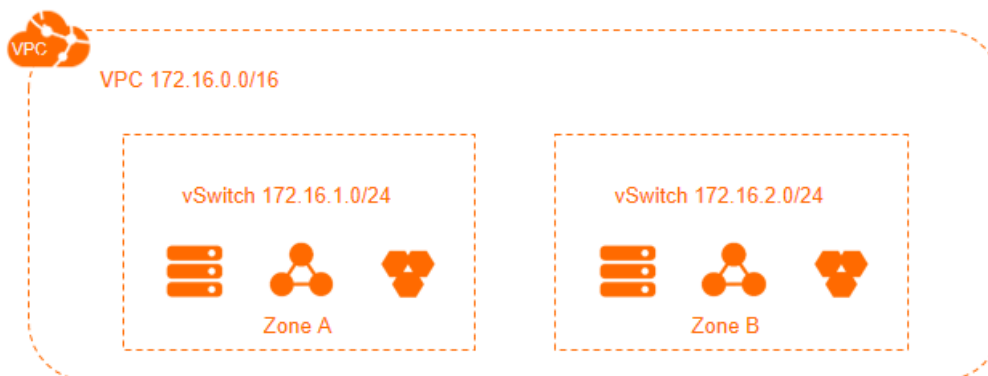
Deploy applications that require access to the Internet

You can deploy applications that require access to the Internet in a subnet of a VPC and use an Internet NAT gateway to route network traffic. You can configure SNAT entries to allow instances in the subnet to access the Internet without the need to expose the private IP addresses. In addition, you can change the elastic IP addresses (EIPs) specified in the SNAT entries to prevent attacks from the Internet.




Implement cross-zone disaster recovery

You can create one or more vSwitches to create one or more subnets for the VPC. vSwitches within the same VPC can communicate with each other. To implement cross-zone disaster recovery, you can deploy resources across vSwitches in different zones.




Isolate business systems

VPCs are logically isolated from each other. You can use multiple VPCs to isolate business systems in different environments such as production and test environments. To allow business systems deployed in two VPCs to communicate with each other, you can create a peering connection between the VPCs.



Build a hybrid cloud

To expand your on-premises network, you can establish a dedicated connection between a VPC and your data center. This allows you to seamlessly migrate the applications in your data center to the cloud. You do not need to change the access method for the applications.



10.3. EIP overview

An elastic IP address (EIP) is a public IP address that you can purchase and use as an independent resource. You can associate an EIP with an Elastic Compute Service (ECS) instance, an internal-facing Server Load Balancer (SLB) instance, or a secondary elastic network interface (ENI) deployed in a virtual private cloud (VPC). You can also associate an EIP with a NAT gateway or a high-availability virtual IP address (HAVIP).

An EIP is a NAT IP address provisioned in the Internet-facing gateway of Alibaba Cloud and is mapped to the associated cloud resource by using NAT. After an EIP is associated with a cloud resource, the cloud resource can use the EIP to communicate with the Internet.

Differences between an EIP and the static public IP address of an ECS instance

The following table describes the differences between an EIP and the static public IP address of an ECS instance.

Item	EIP	Static public IP address
------	-----	--------------------------

Supported network	VPC	VPC
Used as an independent resource	Supported	Not supported
Associated with and disassociated from an ECS instance at any time	Supported	Not supported
Displayed in the ENI information of the associated ECS instance	No	No

Benefits

EIPs have the following benefits:

- Purchase and use as independent resources

You can purchase and use an EIP as an independent resource. EIPs are not bundled with other computing or storage resources.

- Associate with resources at any time

You can associate an EIP with a cloud resource as needed. You can also disassociate and release an EIP at any time.

- Modify bandwidth limits on demand

You can modify the bandwidth limit of an EIP at any time to meet your business requirements. The modification immediately takes effect.

10.3.1. Features

An elastic IP address (EIP) is a public IP address that you can purchase and use as an independent resource. After an EIP is associated with a cloud resource, the cloud resource can use the EIP to communicate with the Internet.

Associate an EIP with a cloud resource

You can associate an EIP with the following resources: an Elastic Compute Service (ECS) instance in a virtual private cloud (VPC), a Server Load Balancer (SLB) instance in a VPC, a secondary elastic network interface (ENI), and a NAT gateway. After an EIP is associated with one of the preceding cloud resources, the cloud resource can use the EIP to communicate with the Internet.

Upgrade the maximum bandwidth of an EIP

You can upgrade the maximum bandwidth of an EIP. After you upgrade the maximum bandwidth of an EIP, the new maximum bandwidth immediately takes effect.

Disassociate an EIP from a cloud resource

You can disassociate an EIP from a cloud resource if the cloud resource no longer needs to communicate with the Internet.

Release an EIP

You can release an EIP that you no longer need.

10.3.2. Benefits

An elastic IP address (EIP) is a public IP address that you can purchase and use as an independent resource. EIPs support flexible association and disassociation, and allow you to modify the maximum bandwidth.

Independent purchase and use

You can purchase and use an EIP as an independent resource. EIPs are not bundled with other computing or storage resources.

Flexible association and disassociation

You can associate an EIP with a cloud resource as needed. You can also disassociate and release an EIP as needed.

Adjustable maximum bandwidth

You can modify the maximum bandwidth of an EIP as needed. The new maximum bandwidth immediately takes effect.

10.3.3. Scenarios

Elastic IP addresses (EIPs) are ideal for scenarios that require Internet connection and disaster recovery.

Internet connection

EIPs can be associated with different cloud resources to meet different requirements:

- After you associate an EIP with an Elastic Compute Service (ECS) instance in a virtual private cloud (VPC), the ECS instance can communicate with the Internet.
- After you associate an EIP with a Server Load Balancer (SLB) instance in a VPC, the SLB instance can forward requests from the Internet.
- After you associate EIPs with a NAT gateway, you can configure DNAT and SNAT entries by using the EIPs. Then, ECS instances can use DNAT and SNAT entries to communicate with the Internet.
- If you associate EIPs with a secondary elastic network interfaces (ENI) and associate the ENI with an ECS instance, the ECS instance can use multiple EIPs. This improves the availability, flexibility, and scalability of your service.

Disaster recovery

You can implement disaster recovery by using EIPs. If the primary server is down, you can disassociate the EIP from the primary server and associate the EIP with the secondary server. This ensures service continuity.

10.4. Express Connect

Express Connect allows you to establish private connections to enable fast, stable, and secure communication between different network environments. You can use Express Connect to ensure network stability and prevent data breaches.

10.4.1. What is Express Connect?

Apsara Stack Express Connect allows you to establish private, flexible, stable, and secure connections between virtual private clouds (VPCs). Network traffic does not traverse the Internet, which prevents data breaches and ensures network stability.

Express Connect allows you to connect VPCs within the same region and account. You can also connect VPCs across different regions and accounts.

Benefits

Express Connect provides the following benefits:

- High-speed connections

Powered by the network virtualization technology of Apsara Stack, Express Connect allows networks to communicate with each other through direct, private, and high-speed connections. Network traffic does not traverse the Internet. The impact of distance on network performance is minimized to ensure low-latency and high-bandwidth communication.

- Stability and reliability

Express Connect provides services based on the high-quality infrastructure of Apsara Stack. This guarantees stable and reliable communication between networks.

- Security

Express Connect allows you to virtualize your networks and enables communication based on the infrastructure provided by Apsara Stack. Networks owned by different accounts are independent of each other. Data packets are exchanged through private connections to prevent breaches.

- On-demand purchase

The maximum bandwidth of Express Connect circuits varies. You can set the maximum bandwidth based on your business requirements.

Differences between Express Connect and Internet connections

A VPC is a logically isolated network. Network traffic between VPCs or between VPCs and data centers are transferred through separate connections.

Compared with Internet connections, Express Connect provides higher performance and security. The following table describes the differences.

Item	Connect VPCs through Internet connections	Connect VPCs through Express Connect
Network performance and availability	When network traffic is transferred over a long distance, a low network latency and a low packet loss rate cannot be guaranteed.	Connections are built based on the infrastructure provided by Apsara Stack, which offers higher quality and availability.
Costs	You are charged a high Internet bandwidth fee or data transfer fee.	The cost for cross-region communication is minimized. You are not charged for connecting VPCs that belong to the same region.
Security	Network traffic is transferred over the Internet. Therefore, data breaches may occur.	Networks are virtualized on Apsara Stack and isolated from each other to ensure high security.

10.4.2. Benefits

A virtual private cloud (VPC) is a logically isolated network. Network traffic between VPCs is transferred through separate connections.

Compared with Internet connections, Express Connect offers higher performance and security

for communication between VPCs. The following table describes the differences.

Item	Connect VPCs through Internet connections	Connect VPCs through Express Connect
Network performance and availability	When network traffic is transferred over a long distance, a low network latency and a low packet loss rate cannot be guaranteed.	Connections are built based on the infrastructure provided by Apsara Stack, which offers higher quality and availability.
Costs	You are charged a high Internet bandwidth fee or data transfer fee.	The cost for cross-region communication is minimized. You are not charged for connecting VPCs that belong to the same region.
Security	Network traffic is transferred over the Internet. Therefore, data breaches may occur.	Networks are virtualized on Apsara Stack and isolated from each other to ensure high security.

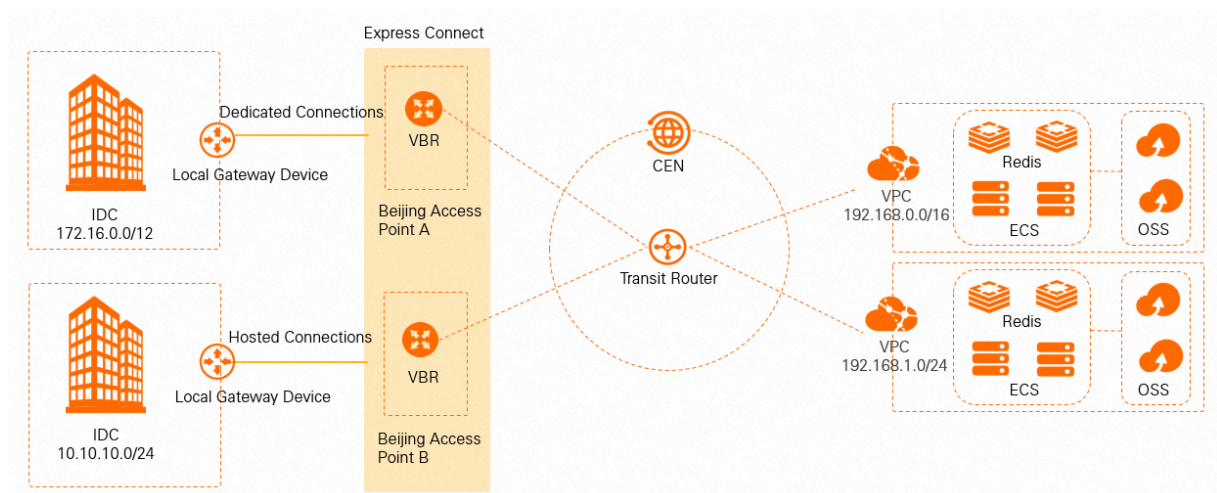
10.4.3. Scenarios

Express Connect enables reliable, secure, and high-speed communication between data centers and VPCs. You can use Express Connect in the following scenarios to facilitate communication in different network architectures.

Scenario 1: Implement disaster recovery for large and medium-sized enterprises

You can establish multiple physical connections between different access points and a VPC to provide high availability for your services. This prevents service outages caused by severed fiber-optic cables, device faults, or access point malfunctions.

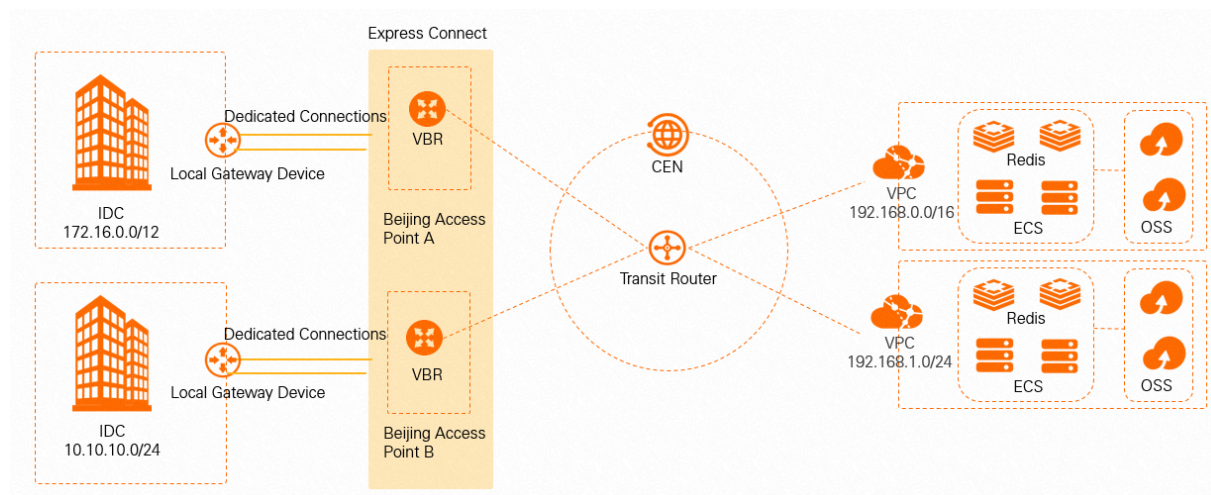
You can use a physical connection interface or use a shared pre-deployed leased line of a partner in this scenario.



Scenario 2: Build highly-scalable and high-availability architectures for large enterprises

When you deploy services on the cloud to provide better support for fast-growing workloads together with on-premises solutions, you can use the equal-cost multi-path routing (ECMP) strategy of Express Connect to offer substantial increases in bandwidth that handles terabytes of data. This prevents service outages caused by device faults, leased line failures, or access point malfunctions.

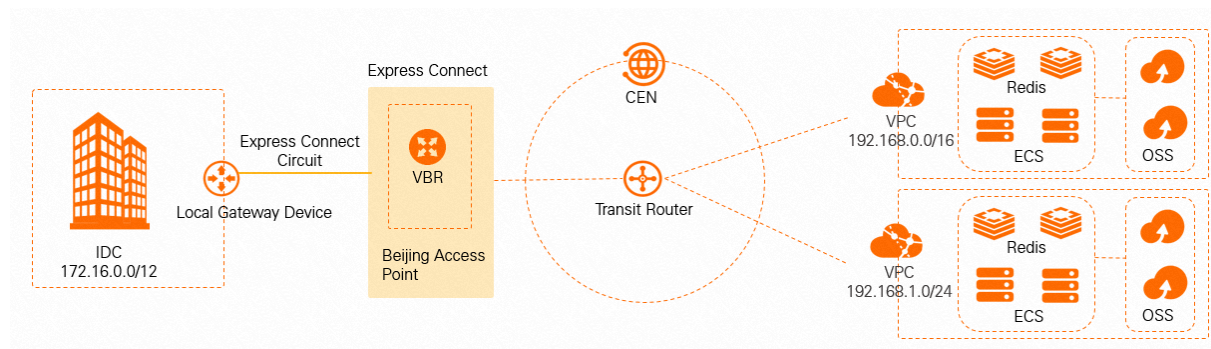
You can only use a physical connection interface in this scenario.



Scenario 3: Establish a basic network architecture for non-critical services

For services that do not require high scalability and availability, such as running tests in a testing environment on the cloud, we recommend that you use Express Connect to directly connect your data center to Alibaba Cloud. This ensures security and reliability of communication between your data center and Alibaba Cloud.

You can use a physical connection interface or use a shared pre-deployed leased line of ad partner in this scenario.




10.5. NAT Gateway

NAT gateways are enterprise-class gateways that support the SNAT and DNAT features. NAT gateways provide a forwarding capacity of 10 Gbit/s and cross-zone disaster recovery capabilities.

10.5.1. What is NAT Gateway?

NAT gateways are enterprise-class gateways that provide the SNAT and DNAT features. Each NAT gateway provides a throughput capacity of up to 10 Gbit/s. NAT gateways also support cross-zone disaster recovery.



Overview

A NAT gateway works as expected only after an elastic IP address (EIP) is associated with the NAT gateway. After you create a NAT gateway, you can associate an EIP with the NAT gateway.

NAT gateways provide the SNAT and DNAT features. The following table describes the features.

Feature	Description
SNAT	SNAT allows Elastic Compute Service (ECS) instances that are deployed in a virtual private cloud (VPC) to access the Internet when no public IP addresses are assigned to the ECS instances.
DNAT	DNAT maps the EIPs that are associated with a NAT gateway to ECS instances. This way, the ECS instances can provide Internet-facing services.

10.5.2. Benefits

NAT Gateway features easy configuration, high performance, high availability, and on-demand purchase.

Easy configuration

A NAT gateway is an enterprise-grade Internet gateway that provides the SNAT and DNAT features. NAT gateways are reliable, flexible, and easy-to-use. NAT gateways save you the trouble of building an Internet gateway by yourself.

High performance

Alibaba Cloud NAT gateways are distributed gateways that use the software-defined networking (SDN) technology. Each NAT gateway provides a throughput capacity of up to 10 Gbit/s, and can serve a large number of Internet applications.

High availability

You can deploy a NAT gateway across zones to implement high availability. Failures in one zone do not affect service continuity.

On-demand purchase

You can modify the specifications of NAT gateways, adjust the number of EIPs, and modify bandwidth of EIPs based on your business requirements.

10.5.3. Use scenarios

NAT gateways allow Elastic Compute Service (ECS) instances in virtual private clouds (VPCs) to access the Internet and receive requests from the Internet.


Configure SNAT to enable ECS instances to access the Internet

You can create a NAT gateway for a VPC, associate an elastic IP address (EIP) with the NAT gateway, and then create an SNAT entry on the NAT gateway. This way, the ECS instances in the VPC can access the Internet by sharing the EIP. This saves public IP resources.

Configure DNAT to provide Internet-facing services

You can create a NAT gateway for a VPC, associate EIPs with the NAT gateway, and then create a DNAT entry on the NAT gateway. This way, ECS instances in a VPC can receive requests from the Internet through port mapping or IP mapping.

- Note** Descriptions of port mapping and IP mapping:
- Port mapping: A NAT gateway forwards requests destined for an EIP to the specified ECS instance based on the specified ports and protocol.
 - IP mapping: A NAT gateway forwards all requests destined for an EIP to the specified ECS instance.

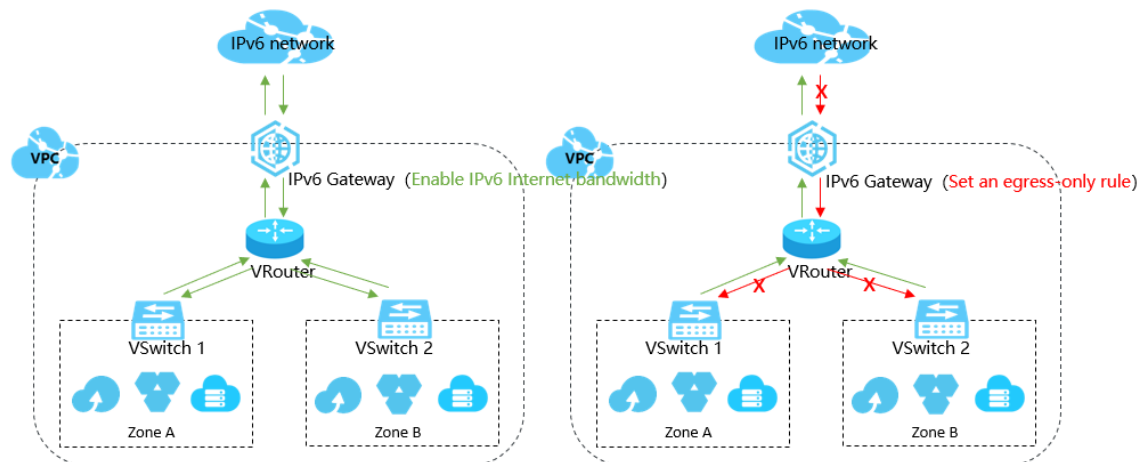


10.6. IPv6 Gateway

IPv6 gateways route IPv6 traffic to and from virtual private clouds (VPCs). IPv6 gateways support high availability across zones and 10-gigabit throughput to process a large number of requests. IPv6 gateways help you build a secure and reliable IPv6 network.

10.6.1. What is an IPv6 Gateway?

This topic provides an overview of the IPv6 Gateways of Virtual Private Cloud (VPC). An IPv6 Gateway functions as an IPv6 traffic gateway for a VPC. You can configure the IPv6 Internet bandwidth and egress-only rules to manage the inbound and outbound IPv6 traffic.



Functions

The functions of an IPv6 gateway are as follows:

- **IPv6 internal network communication**

By default, an IPv6 address in a VPC is allocated with an Internet bandwidth of 0 Mbit/s and only supports communication over the internal network. Specifically, the cloud instances in a VPC can only access other IPv6 addresses in the same VPC through the IPv6 address. The resources cannot access the Internet with these IPv6 addresses or be accessed by IPv6 clients over the Internet.

- **IPv6 public network communication**

You can purchase an Internet bandwidth for the IPv6 address for which you have applied. In this way, the resources in the VPC can access the Internet through the IPv6 address and be accessed by IPv6 clients over the Internet.

You can set the Internet bandwidth to 0 Mbit/s at any time to deny the IPv6 address Internet access. After this configuration, the IPv6 address can only communicate over the internal network.

- **IPv6 public network communication with an egress-only rule**

You can set an egress-only rule for an IPv6 Gateway. In this way, the IPv6 address can access the Internet, but IPv6 clients are denied access to your cloud resources in the VPC over the Internet.

You can delete the egress-only rule at any time. After the rule is deleted, your resources in the VPC can access the Internet through the IPv6 address for which you have purchased Internet bandwidth, and IPv6 clients can access the resources in the VPC over the Internet.

The network access capability of IPv6 addresses is dependent on the settings of the network type, Internet bandwidth, and egress-only rule, as shown in the following table.

IPv6 network type	Enable IPv6 Internet bandwidth?	Set an egress-only rule?	IPv6 network access capability
Internal network	No	No	Internal network communication
Public network	Yes	No	Internal network communication Public network communication
		Yes	Internal network communication Public network communication when access is initiated by VPCs

Benefits

IPv6 Gateway provides the following benefits:

- **High availability**

IPv6 Gateways provide cross-zone high availability and stable IPv6 Internet gateway services.

- **High performance**

A single IPv6 Gateway provides a 10-gigabit level throughput.

- **Flexible management of public network communication**

You can manage the Internet communication capability of an IPv6 Gateway by adjusting its Internet bandwidth and setting an egress-only rule.

10.6.2. Terms

This topic describes the terms used in IPv6 Gateway.

Term	Description
IPv6 address	The IPv6 address allocated by the system to an instance in a VPC. An IPv6 address is made of 128 binary bits that are divided into eight 16-bit groups separated by colons (:). Each group is represented as a 4-digit hexadecimal number. The following is an example of an IPv6 address: 2001:xxx:0102::0304
IPv6 gateway	The Internet gateway for IPv6 traffic flowing in and out of a VPC. You can use an IPv6 gateway to control and manage the bandwidth used by IPv6 traffic. IPv6 gateways allow you to create egress-only rules to funnel egress traffic.
IPv6 Internet bandwidth	The Internet bandwidth of an IPv6 address that limits the bandwidth of Internet connectivity for the IPv6 address. You must purchase and add IPv6 Internet bandwidth to an IPv6 address before the IPv6 address can be used to communicate over the Internet.

Egress-only rule	A rule by which an IPv6 gateway implements egress control for IPv6 traffic. After you configure an egress-only rule for an IPv6 address, the IPv6 gateway allows outbound only communication to the Internet over IPv6 using the IPv6 address, and prevents the Internet from initiating IPv6 connections with the instance associated with the IPv6 address.
IPv6 CIDR block for VPC	A /61 IPv6 CIDR block automatically allocated to a VPC after IPv6 is enabled for the VPC.
IPv6 CIDR block for vSwitch	An IPv6 CIDR block allocated to a vSwitch. The default subnet mask for the IPv6 CIDR block of a vSwitch is /64. When you enable IPv6 for a vSwitch, you can specify the last eight bits of the IPv6 CIDR block of the vSwitch.


10.6.3. Common scenarios

IPv6 gateways help you build a secure and reliable IPv6 environment.

Scenario 1: Enable IPv6 for a virtual private cloud (VPC) and build an isolated IPv6 environment

If you enable IPv6 for an existing VPC, the VPC supports both IPv4 and IPv6. You can assign IPv6 addresses to the Elastic Compute Service (ECS) instances on which services are deployed. This way, the ECS instances can use IPv4 addresses and IPv6 addresses. By default, the IPv6 address of an ECS instance can be used only for communication within the VPC.

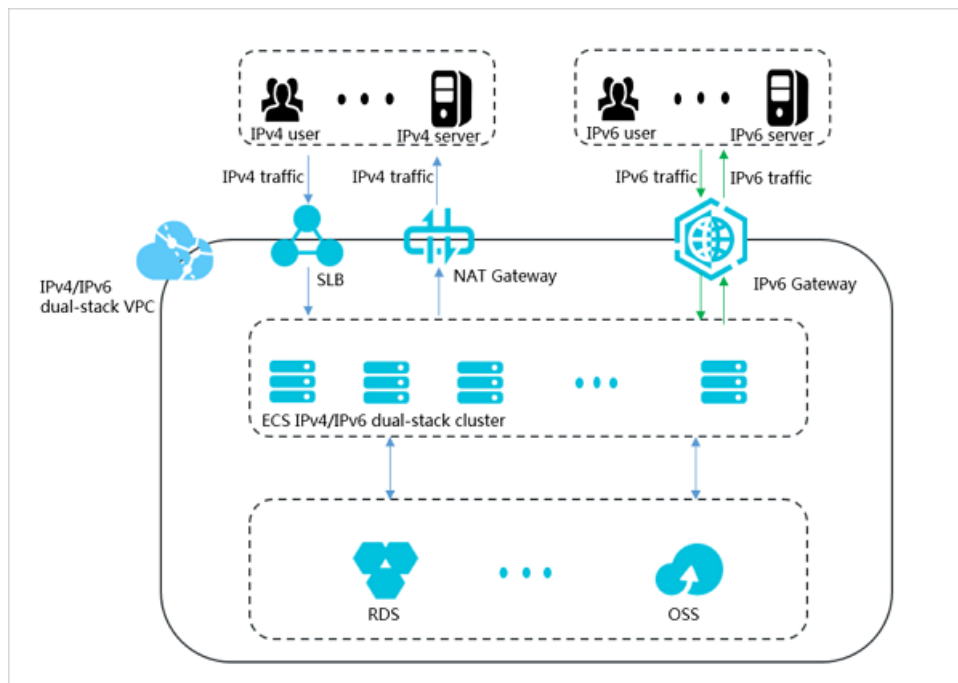
ECS instances for which IPv4 and IPv6 are enabled can use IPv4 addresses or IPv6 addresses to communicate with other resources in the VPC. The ECS instances cannot use IPv6 addresses to access the Internet or provide services to IPv6 clients over the Internet.



Scenario 2: Enable ECS instances in a VPC to communicate with the Internet by using IPv6 addresses

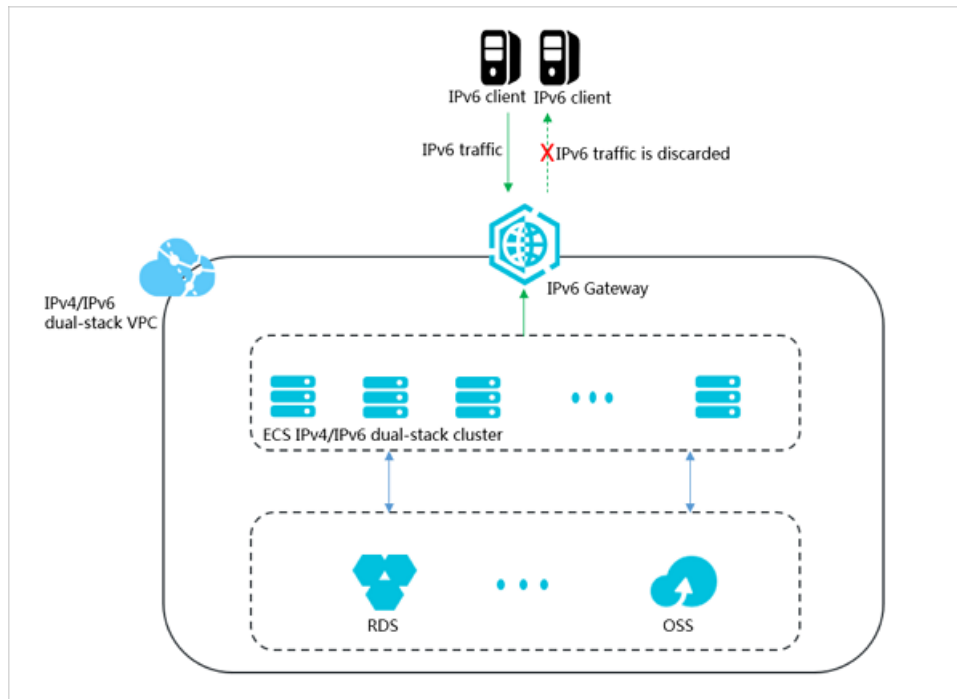
After you enable IPv6 Internet bandwidth for an IPv6 address, the IPv6 address can be used for communication over the Internet. IPv6 traffic between ECS instances in a VPC and IPv6 networks traverses the IPv6 gateway. The IPv6 gateway processes inbound and outbound IPv6 traffic.

IPv4 traffic between ECS instances and IPv4 networks traverses the Server Load Balancer (SLB) instance and the NAT gateway. The SLB instance and the NAT gateway process inbound and outbound IPv4 traffic.



Scenario 3: Configure egress-only rules to manage IPv6 traffic

If you want an ECS instance to access IPv6 clients and deny access from IPv6 clients, you can configure an egress-only rule for the ECS instance. This way, the ECS instance can access IPv6 networks, but does not receive requests from IPv6 clients.



10.7. Cloud Gateway

Cloud Gateway (CGW) is designed to provide scenario-specific network services to help you efficiently build hybrid clouds and migrate workloads to the clouds. CGW is an ideal service for migrating your infrastructure from a public or private cloud to a hybrid cloud.

10.7.1. Service description

Cloud Gateway (CGW) provides scenario-specific network access services to help enterprises efficiently build hybrid clouds and migrate workloads to clouds.

ONSP

Capabilities

- Allows you to create and delete service clusters.
- Allows you to manage resources.
- Allows you to route network traffic from VPCs to service clusters.

Supported services

- Load balancing: allows you to configure load balancing services of ecosystem partners, such as F5 load balancers, Sangfor Technologies load balancers, and Radware load balancers, for VPCs.
- Virtual firewalls: allows you to protect traffic between VPCs and traffic between a data center and a VPC and enable address translation. Hillstone firewalls are supported.
- Network performance monitoring (NPM): allows you to collect and analyze network traffic by using third-party services, such as DeepFlow provided by Yunshan Networks and Network Performance Management (NPM) provided by Netis.

10.7.2. Benefits

Cloud Gateway (CGW) provides you with comprehensive, efficient, secure, easy-to-use, and visual network services.

ONSP

ONSP allows you to connect to, deploy, and manage network elements from ecological partners.

- Automatic deployment: Virtual network elements from ecological partners can be automatically deployed by using standard procedures and models.
- High reliability: Virtual network elements can be deployed in active/standby clusters, which support automatic switchover for network traffic.
- Centralized management: Virtual network elements can be deployed in an exclusive VPC. You can centrally manage the virtual network elements at a low cost.

10.8. Apsara Stack DNS

Apsara Stack DNS is a service that complies with the Domain Name System (DNS) protocol and runs on Apsara Stack to resolve domain names over your internal networks, such as virtual private clouds (VPCs), physical networks, and networks of self-managed data centers. You can configure rules and policies to map domain names to IP addresses. Apsara Stack DNS then distributes the DNS requests from clients to cloud resources, self-managed business applications, business systems on your internal networks, or the service resources of Internet service providers (ISPs).

10.8.1. Service description

Apsara Stack DNS provides the Domain Name System (DNS) resolution and global server load balancing (GSLB) services for virtual private clouds (VPCs), networks of self-managed data centers, and classic networks.

Internal DNS resolution management

You can manage the global internal domain names, global forwarding configurations, and global recursive settings created on Apsara Stack. The changes to these configurations take effect for all VPCs and classic networks.

Apsara Stack DNS provides undifferentiated global DNS resolution services for all servers in a VPC. DNS servers use anycast virtual IP addresses in specific regions. This delivers seamless failovers for the disaster recovery scenarios of data centers in these regions.

Global internal domain names

Apsara Stack DNS allows you to manage the data of global internal domain names. You can register, search for, describe, and delete global internal domain names. You can also add, delete, and modify the hostnames for these domain names. Apsara Stack DNS supports the following types of DNS records: A, AAAA, CNAME, MX, PTR, TXT, SRV, NAPTR, CAA, and NS.

- You can add multiple A, AAAA, and PTR records for a hostname of a domain name. DNS servers return all record values by default through random round-robin to achieve load balancing.
- You can add multiple A, AAAA, and CNAME records for a hostname of a domain name. DNS servers return a record value based on the weights configured for the record values of these records to achieve load balancing.

Global forwarding configurations

Apsara Stack DNS can forward the DNS requests of specific domain names to other DNS servers for resolution. Apsara Stack DNS also supports global default forwarding, which forwards the DNS requests of the domain names that do not have forwarding configurations to other DNS servers for resolution.

Apsara Stack DNS can forward requests with or without recursion.

- In the mode of forwarding without recursion, only the specified DNS server is used to resolve domain names. If the resolution fails or times out, a message is returned to the DNS client to indicate that the current request fails.
- In the mode of forwarding with recursion, the specified DNS server is preferentially used to resolve domain names. If the resolution fails, a local DNS server is used.

Global recursive resolution

Apsara Stack DNS provides recursive resolution of Internet domain names so that enterprises can access Internet services. You can enable or disable the global recursive resolution feature and modify the configurations of this feature.

PrivateZone

The PrivateZone feature allows you to create tenant-specific domain names in VPCs. You can bind the domain names to VPCs and unbind the domain names from VPCs as required to isolate tenants. Changes to these configurations take effect only in the VPCs to which the domain names are bound.

The PrivateZone feature provides personalized DNS resolution services for the servers in the VPCs to which the domain names are bound. DNS servers use anycast virtual IP addresses in specific regions. This delivers seamless failovers for the disaster recovery scenarios of data centers in these regions.

Note

This feature is only supported in Apsara Stack DNS Standard Edition.

Tenant internal domain name

Apsara Stack DNS allows you to manage the data of tenant internal domain names. You can register, search for, describe, and delete tenant internal domain names. You can also add, delete, and modify the hostnames for these domain names. Apsara Stack DNS supports the following types of DNS records: A, AAAA, CNAME, MX, PTR, TXT, SRV, NAPTR, CAA, and NS.

- You can add multiple A, AAAA, and PTR records for a hostname of a domain name. DNS servers return all record values by default through random round-robin to achieve load balancing.
- You can add multiple A, AAAA, and CNAME records for a hostname of a domain name. DNS servers return a record value based on the weights configured for the record values of these records to achieve load balancing.

Tenant forwarding configurations

Apsara Stack DNS can forward the DNS requests of specific domain names to other DNS servers for resolution. Apsara Stack DNS also supports global default forwarding, which forwards the DNS requests of the domain names that do not have forwarding configurations to other DNS servers for resolution.

Apsara Stack DNS can forward requests with or without recursion.

- In the mode of forwarding without recursion, only the specified DNS server is used to resolve domain names. If the resolution fails or times out, a message is returned to the DNS client to indicate that the current request fails.

- In the mode of forwarding with recursion, the specified DNS server is preferentially used to resolve domain names. If the resolution fails, a local DNS server is used.

Internal Global Traffic Manager

Internal Global Traffic Manager (GTM) provides multi-cloud disaster recovery for your domain names. You can connect your domain names to an internal GTM instance to manage traffic loads between Apsara Stack systems.

Internal GTM supports global server load balancing (GSLB) for enterprises. This feature intelligently allocates IP addresses for DNS queries from request sources based on configured scheduling policies. It also supports multi-cloud, hybrid deployment, and configuration data synchronization among clouds.

Note

This feature is only supported in internal GTM Standard Edition of Apsara Stack DNS.

Scheduling instance management

- Apsara Stack DNS allows you to manage scheduling instances. One scheduling instance corresponds to one application instance.
- Apsara Stack DNS allows you to manage address pools. One address pool corresponds to one service cluster of one application instance.
- Apsara Stack DNS allows you to manage scheduling domains and set the scheduling domains to which scheduling instances belong. In this way, you can manage and code global scheduling instances based on your naming conventions.

Resolution line management

You can customize lines and their priorities to allow clients to access the nearest nodes and implement intelligent traffic scheduling based on geographical locations and application groups. This accelerates access to applications.

Data synchronization management

- Apsara Stack DNS allows you to manage global data synchronization. You can create data synchronization tasks, manage data synchronization configurations, and view data synchronization information about multiple internal GTM services. The information includes local system information, information about the cluster nodes on which a data synchronization relationship has been established, and primary and secondary relationships.
- Apsara Stack DNS allows you to manage the messages for changes to data synchronization tasks. This helps you confirm the messages for requesting primary nodes to actively add secondary nodes.

10.8.2. Benefits

You can perform the following operations by using Apsara Stack DNS in virtual private clouds (VPCs), networks of self-managed data centers, and physical networks:

- Access other Elastic Compute Service (ECS) instances deployed in the same VPC.
- Access other cloud service instances on Apsara Stack.
- Access the custom business systems of your enterprise.
- Access services over the Internet.

- Use the global server load balancing (GSLB) feature to implement multi-active solutions and disaster recovery, such as local active-active disaster recovery, active zone-redundancy, remote active-active disaster recovery, active geo-redundancy, and geo-disaster recovery.
- Connect to Apsara Stack DNS with your own DNS servers over a leased line.

Enterprise domain name management

Apsara Stack DNS provides management and resolution services for your domain names. It supports the following features:

- Performs forward and reverse DNS lookups for domain names of cloud service instances, such as ECS instances.
- Performs forward and reverse DNS lookups for your internal domain names.
- Adds, modifies, and removes common DNS records, such as A, AAAA, CNAME, NS, MX, TXT, SRV, and PTR records.
- Adds multiple A, AAAA, and PTR records for a hostname of a domain name, and returns all record values by default through random round-robin to achieve load balancing.

Flexible integration with data centers

Apsara Stack DNS can forward enterprise domain names and provide the following services for you to flexibly build your network and cascade Apsara Stack DNS servers with the DNS servers that you create:

- Global default forwarding
- Query forwarding for specific domain names

Internet access from enterprise servers

Apsara Stack DNS supports recursive resolution for Internet domain names, which allows your servers to access the Internet.

Tenant isolation

Apsara Stack DNS allows you to manage private zones in VPCs, resolve internal domain names, and isolate DNS records and resolution based on tenants.

- You can add, delete, modify, and query private authoritative zones. You can also bind and unbind private authoritative zones to and from VPCs.
- You can add, delete, modify, and query private forwarding zones. You can also bind and unbind private forwarding zones to and from VPCs.

Note

This feature is available only in Apsara Stack DNS Standard Edition.

Global Traffic Manager

Apsara Stack DNS provides the following Global Traffic Manager (GTM) features for internal networks:

- Adds multiple A, AAAA, and CNAME records for a hostname of a domain name, and returns a record value based on the weights configured for the record values of these records to achieve load balancing.
- Supports scheduling line management. You can customize lines and their priorities to allow clients to access the nearest nodes and implement intelligent traffic scheduling based on geographical locations and application groups. This accelerates access to applications.

- Synchronizes DNS configuration data among multiple clusters for which GTM is activated. This feature is supported in multi-cloud scenarios.
- Supports address pool management. You can manage your applications in a centralized manner by application service cluster.
- Supports custom global scheduling domains. You can manage and code global scheduling instances in a centralized manner based on your naming conventions.

Centralized management console

- You can access Apsara Stack DNS and any other cloud services on the Apsara Uni-manager Management Console with one account.
- Apsara Stack DNS supports web operations for data and service management, which facilitates your use of the DNS service.

10.8.3. Scenarios

Apsara Stack DNS is a key network service that controls data traffic for Apsara Stack. This service resolves domain names, balances server loads, and connects Apsara Stack with data centers and Alibaba Cloud public cloud. Apsara Stack DNS provides a complete suite of solutions to deploy a cloud environment, achieve high availability and disaster recovery for data centers, and balance server loads to secure your IT services.

Basic DNS resolution

Access cloud resource instances in a VPC

You can use Apsara Stack DNS to access ApsaraDB RDS, Server Load Balancer (SLB), and Object Storage Service (OSS) instances from the Elastic Compute Service (ECS) instances or Docker instances in a virtual private cloud (VPC).

Access ECS instances or Docker instances by using hostnames in a VPC

You can create hostnames for the ECS instances or Docker instances in a VPC based on the rules that you define, and access and manage these instances by using the hostnames.

Access the domain names of internal services in a VPC

You can develop PaaS or SaaS services on Apsara Stack and access the PaaS or SaaS services by using domain names in a VPC.

Schedule the requests to internal services on Apsara Stack by using the round-robin algorithm

If a SaaS service that you develop on Apsara Stack is deployed in multiple data centers or regions, you can use Apsara Stack DNS to access the SaaS service and evenly distribute requests to different nodes in a VPC.

Access Internet services from a VPC or physical network

You can use Apsara Stack DNS to access Internet services from a VPC or physical network.

Establish connections between Apsara Stack and other networks

You can use Apsara Stack DNS to connect your internal network, a public cloud network, or other external networks to Apsara Stack to implement DNS resolution.

Tenant isolation

Isolate tenant resources on Apsara Stack

You can use Apsara Stack DNS to isolate tenant resources on Apsara Stack. The internal DNS records and the default forwarding configurations of each tenant are isolated from those of other tenants. Tenants can configure their own private DNS records to complete service addressing and scheduling.

Establish connections among global resources on Apsara Stack

If you want to allow all tenants to share global resources and configurations on Apsara Stack, system administrators can configure global DNS records and settings to complete service addressing and scheduling.

Perform VPC-based intelligent scheduling on Apsara Stack

You can use Apsara Stack DNS to provide different resolution results for different VPCs of a tenant based on the same hostname.

Global scheduling

Schedule the traffic loads of internal network services on Apsara Stack based on weights

If a PaaS or SaaS service that you created on Apsara Stack is deployed in multiple data centers or regions, you can use Apsara Stack DNS to access the service from within or outside the cloud and distribute requests to different nodes in backend service clusters based on the weights of these nodes.

Schedule the traffic loads of internal network services on Apsara Stack based on the geographical locations of request sources or the application groups to which request sources belong

If a PaaS or SaaS service that you created on Apsara Stack is deployed in multiple data centers or regions, you can use Apsara Stack DNS to access the service from within or outside the cloud and distribute requests to different nodes based on the geographical locations or application groups of the request sources.

Enterprise disaster recovery

Schedule the traffic loads of internal network services on Apsara Stack to achieve disaster recovery

If a PaaS or SaaS service that you created on Apsara Stack is deployed in multiple data centers or regions, you can use Apsara Stack DNS to access the service from within or outside the cloud and switch the access traffic of internal services from backend service cluster A (primary data center) to backend service cluster B (secondary data center) in a disaster recovery scenario.

11. Database services

11.1. ApsaraDB RDS

ApsaraDB RDS is a stable, reliable, and scalable online database service. Based on high-performance storage, ApsaraDB RDS provides a set of solutions for disaster recovery, backup, restoration, monitoring, and migration.

11.1.1. What is ApsaraDB RDS?

ApsaraDB RDS is a stable, reliable, and scalable online database service. Based on the distributed file system and high-performance storage, ApsaraDB RDS provides a set of solutions for disaster recovery, backup, restoration, monitoring, and migration.

ApsaraDB RDS supports four database engines, which are MySQL, SQL Server, and PostgreSQL. You can create database instances based on these database engines to meet your business requirements.

RDS MySQL

Originally based on a branch of MySQL, ApsaraDB RDS for MySQL provides excellent performance. It is a tried and tested solution that handled the high-volume concurrent traffic during Double 11. ApsaraDB RDS for MySQL provides basic features, such as whitelist configuration, backup and restoration, Transparent Data Encryption (TDE), data migration, and management for instances, accounts, and databases. ApsaraDB RDS for MySQL also provides the following advanced features:

- **Read-only instance:** In scenarios where ApsaraDB RDS for MySQL handles a small number of write requests but a large number of read requests, you can create read-only instances to scale up the reading capability and increase the application throughput.
- **Read/write splitting:** The read/write splitting feature provides a read/write splitting endpoint. This endpoint enables an automatic link for the primary instance and all of its read-only instances. An application can connect to the read/write splitting endpoint to read and write data. Write requests are distributed to the primary instance and read requests are distributed to read-only instances based on their weights. To scale up the reading capability of the system, you can add more read-only instances.

RDS SQL Server

ApsaraDB RDS for SQL Server provides strong support for a variety of enterprise applications under the high-availability architecture, has the capability of restoring data to any point in time, and covers Microsoft licensing fee.

ApsaraDB RDS for SQL Server provides basic features such as whitelist configuration, backup and restoration, TDE, data migration, and management for instances, accounts, and databases.

RDS PostgreSQL

ApsaraDB RDS for PostgreSQL is an advanced open source database service that is fully compatible with SQL and supports a diverse range of data formats such as JSON, IP, and geometric data. In addition to support for features such as transactions, subqueries, multi-version concurrency control (MVCC), and data integrity check, ApsaraDB RDS for PostgreSQL integrates a series of features including high availability, backup, and restoration to ease operations and maintenance loads.

11.1.2. Benefits

11.1.3. Scenarios

11.2. PolarDB

PolarDB is a relational database service developed by Alibaba Cloud. It uses the cloud-native architecture, decoupled compute and storage resources, and integrated software and hardware. This database service is characterized by high elasticity, high performance, high availability, high reliability, and cost-effectiveness. PolarDB supports SQL and NoSQL access modes, and are compatible with MySQL and PostgreSQL protocols. This simplifies its integration with existing databases to provide efficient data processing capabilities.

PolarDB uses the common Apsara Stack infrastructure, which allows you to enjoy the core capabilities of PolarDB at low costs.

- PolarDB uses the database engine that has been optimized by Alibaba Cloud for years and that provides much higher performance than open-source engines.
- The high-performance computing and storage infrastructure of PolarDB can significantly reduce user costs.
- PolarDB decouples computing and storage. Each cluster consists of one primary node and multiple read-only nodes. Configuration changes and adding nodes can be implemented in minutes.
- Multiple compute nodes share the same storage. You only need to add compute resources when you add read-only nodes. This greatly reduces scale-out costs.

11.2.1. Features

PolarDB is a powerful cloud-native relational database that features high availability, high performance, and high scalability. It is suitable for enterprise-grade applications of all sizes.

Decoupled storage and computing

PolarDB decouples computing and storage. PolarProxy and compute nodes are deployed on separate ECS instances. Enhanced SSDs are used as the shared storage. This greatly reduces PolarDB costs.

Note

If you select one node when you create a PolarDB for MySQL cluster, the cluster uses the single-node architecture.

- You can add read-only nodes to change the single-node architecture to the multi-node architecture. Then, you can enable PolarProxy.
- You can also change the multi-node architecture to the single-node architecture by removing read-only nodes.

Turbocharged MySQL

- Fully compatible with open source MySQL and ApsaraDB RDS for MySQL: You can migrate data from these databases to phph without changing the code or configurations of your applications.
- Powerful features: one-to-many and many-to-many association for primary and read-only nodes, shared storage, high-speed processing for large tables (6 billion rows), smooth failover, flashback queries, and fast DDL queries (response times within seconds).

11.2.2. Benefits

You can use PolarDB in the same way as using traditional databases. Compared with traditional databases, PolarDB has the following benefits:

- **Cost-effectiveness**

- Compute nodes share storage resources. You pay only for compute nodes when you add read-only nodes, which greatly reduces scale-out costs.

- **High elasticity**

- Nodes can be added or deleted in minutes.
- Storage space is automatically scaled without business interruptions as your data volume changes.

- **High performance**

The database engine that PolarDB uses has been optimized upon the open source engines. PolarDB also leverages capabilities like physical replication, high-speed network protocol, and shared distributed storage to deliver six times the performance of open source MySQL databases.

- **High availability, reliability, and security**

- Shared distributed storage eliminates the data inconsistency issue that may occur when data is asynchronously synchronized from the primary node to secondary nodes. This ensures zero data loss if a single point of failure occurs in a cluster.
- A hot standby storage cluster is deployed in the secondary zone or in a different data center in the same zone of the region where a PolarDB cluster is deployed. The hot standby storage cluster uses independent storage for hot standby of the PolarDB cluster. When the PolarDB cluster or the primary zone is unavailable, the hot standby storage cluster quickly becomes the primary node to perform read and write operations on the PolarDB cluster.

- **Data security**

- PolarDB adopts various security measures such as IP whitelists, VPCs, and multiple data replicas to protect your data in terms of access, storage, and management.

- **Lock-free backup**

- Snapshots that are implemented based on the distributed storage can back up a database with terabytes of data in a few minutes. During the entire backup process, no locks are required, which ensures high efficiency and minimized impacts on your business.

11.2.3. Scenarios

PolarDB can be used in various scenarios to meet different business requirements and provide the database service featured by high performance, high availability, and high scalability.

- PolarDB single-node architecture is a highly recommended cost-effective service for individual users who need to test and learn about database services. This edition is suitable for startups in the early development stage of their applications.
- PolarDB multi-node architecture is suitable for large and medium-sized enterprises whose production databases need to process a large number of read requests during peak hours or need to perform intelligent data analysis. These enterprise users include governments, online retailers, automobile enterprises, education enterprises, and Enterprise Resource Planning (ERP) service providers.

11.3. ApsaraDB for MongoDB

ApsaraDB for MongoDB is a high-performance document database service that is fully compatible with the MongoDB protocol. Based on the Apsara distributed system and a highly reliable storage engine, ApsaraDB for MongoDB provides features such as high availability, elastic scaling, read/write splitting, high data security, backup and restoration, intelligent O&M, and online database management.

11.3.1. What is ApsaraDB for MongoDB?

ApsaraDB for MongoDB is a MongoDB-compatible document-oriented database service that is developed based on the Apsara system and a high-reliability storage engine. ApsaraDB for MongoDB uses a multi-node architecture to ensure high availability and supports elastic scaling, disaster recovery, backup and restoration, and performance optimization.

Data structure

MongoDB is a document-oriented NoSQL database. MongoDB stores data in JSON-like documents that consist of field-value pairs. Example:

```
{
  name: "John",
  sex: "male",
  age: 30
}
```

Storage structure

The storage structure of MongoDB is different from that of conventional relational databases. Data in MongoDB is organized at the following levels:

- Document

Documents are the basic unit of data in MongoDB. A document consists of BSON key-value pairs and is equivalent to a row in a relational database.
- Collection

A collection can contain multiple documents. Collections are equivalent to tables in a relational database.
- Database

A database can contain multiple collections. You can create multiple databases in MongoDB. Databases are equivalent to relational databases.

Instance architectures


ApsaraDB for MongoDB provides the following instance architectures:

- Replica set instances

An ApsaraDB for MongoDB replica set instance consists of a primary node, one or more secondary nodes, and a hidden node.

 - Primary node: processes all read and write operations.
 - Secondary node: synchronizes data from the primary node. If the primary node fails, a secondary node becomes the new primary node to ensure high availability.

- Hidden node: ensures high availability of the instance. If a secondary node fails, the hidden node becomes the secondary node.


 **Note** The hidden node is used only to ensure high availability. It is invisible to users.

- Sharded cluster instances
An ApsaraDB for MongoDB sharded cluster instance consists of mongos, shard, and Configserver nodes.
 - Mongos node: routes queries and write operations to the corresponding shard node.
 - Shard node: stores database data.
 - Configserver node: stores data of shard nodes.


Deployment suggestions

When you deploy an ApsaraDB for MongoDB instance, you can consider the following aspects:

- Regions and zones
You can select a region and zone based on your location, the availability of Alibaba Cloud services, your application availability requirements, and whether internal network communication is required.

 **Note** A region is an Alibaba Cloud data center. A zone is a physical area within a region that has its own independent power supply and network.

For example, if your application is deployed on an Elastic Compute Service (ECS) instance and requires an ApsaraDB for MongoDB instance to serve as its database, you must select the same region and zone as the ECS instance when you create your ApsaraDB for MongoDB instance.

-  **Note**
- An ECS instance and an ApsaraDB for MongoDB instance within the same zone can be connected by using an internal network with minimal network latency.
 - The region and zone determine the physical location of an ApsaraDB for MongoDB instance. You cannot change the region of an ApsaraDB for MongoDB instance after the instance is created.

- Network planning

A virtual private cloud (VPC) helps you build an isolated network environment on Apsara Stack. You can customize route tables, CIDR blocks, and gateways within a VPC. We recommend that you select VPC for improved security.

- Security policies

ApsaraDB for MongoDB provides comprehensive security measures to improve data security. You can ensure database security by means of zone-disaster recovery, audit logs, network isolation, whitelists, password authentication, SSL encryption, and Transparent Data Encryption (TDE).

11.3.2. Benefits

ApsaraDB for MongoDB is a MongoDB-compatible document-oriented database service that is developed based on the Apsara system and a high-reliability storage engine. ApsaraDB for MongoDB ensures high availability, supports elastic scaling, and delivers security.

High availability

- High-availability architecture, zone-disaster recovery, and automatic backup to ensure

business availability

- ApsaraDB for MongoDB provides the three-node replica set and sharded cluster architectures. Multiple data nodes are deployed on various physical servers. When a node in an instance fails, other nodes automatically synchronize data to ensure the high availability of the instance.
- ApsaraDB for MongoDB allows you to create dual-zone instances. When a zone for a dual-zone instance becomes unavailable due to unexpected events, the data in the instance can be synchronized in the other zone to ensure the continued availability of the instance.
- ApsaraDB for MongoDB provides the automatic backup feature. The system automatically backs up data and uploads the data to Object Storage Service (OSS) during the specified time period. This improves disaster recovery capabilities and reduces consumed disk capacity. Backup files can be used to restore instance data to their source instance and prevent irreversible effects on business data caused by accidental changes and other errors.
- Primary/secondary failover to ensure service availability
ApsaraDB for MongoDB provides the primary/secondary failover feature. When a node of an instance fails, the system triggers a primary/secondary failover to ensure availability of the instance.

Elastic scaling

- Flexible configuration changes to meet business requirements
ApsaraDB for MongoDB allows you to change instance configurations. Multiple instance types are available for configuration change. You can change instance configurations as your business needs change.
- Multiple chip architectures for hybrid deployment
ApsaraDB for MongoDB provides a variety of chip architectures such as x86 and ARM. You can select a chip architecture based on your needs for scenarios such as scale-out events, disaster recovery, and hybrid deployment.

Security

- Pre-protection
ApsaraDB for MongoDB provides the DDoS mitigation feature. ApsaraDB for MongoDB monitors inbound traffic in real time, filters source IP addresses to scrub large amounts of malicious traffic, and triggers blackhole filtering if traffic scrubbing becomes ineffective.
- In-event protection
 - Configuration of IP address whitelists for enhanced database access security
IP addresses or CIDR blocks used to access databases can be added to a whitelist of the instance to ensure security and stability of databases.
 - SSL encryption for enhanced link security
SSL encryption can encrypt network connections at the transport layer to improve data security and ensure data integrity. After SSL encryption is enabled, you can install SSL certificates issued by certification authorities (CAs) on your application to improve link security.
 - TDE for improved data security
Transparent Data Encryption (TDE) is used to encrypt data before the data is written from data files to a disk and decrypt data before the data is read from a disk and written to the memory. TDE does not increase the sizes of data files. You can use TDE without the need to modify the configuration data of your application. You can enable TDE for an instance to encrypt instance data and improve data security.

- Post-auditing

ApsaraDB for MongoDB automatically stores audit logs in Log Service and allows you to download these logs from Log Service. This facilitates the long-term storage and management of audit logs.

Intelligent O&M

- Comprehensive monitoring to help O&M personnel understand the running status of instances

ApsaraDB for MongoDB provides a variety of performance monitoring metrics such as CPU utilization, memory usage, and disk usage for you to view the running status of your instances in real time.

- Performance optimization to ensure stability, security, and efficiency of databases

ApsaraDB for MongoDB allows you to view and customize instance performance trends and view performance, storage, and slow query logs of instances in real time. This helps you eliminate service failures caused by manual operations and ensure the stability, security, and efficiency of databases.

Network isolation

ApsaraDB for MongoDB uses Virtual Private Cloud (VPC) to implement advanced network access control. VPC and IP address whitelists greatly improve the security of ApsaraDB for MongoDB instances.

A VPC can help you build an isolated network environment by using underlying network protocols. You can resolve resource conflicts by customizing route tables, IP addresses, and gateways in VPCs.

By default, ApsaraDB for MongoDB instances deployed in a VPC can be accessed only by the Elastic Compute Service (ECS) instances in the same VPC. If necessary, you can apply for a public IP address to allow access requests from the Internet (not recommended). Before you apply for a public IP address, you must configure a whitelist. For example, you can allow access requests from elastic IP addresses (EIPs) of ECS instances and the Internet egress of your data center.

Online management of databases

Data Management (DMS) is an integrated and visualized database solution that offers data management, structure management, user authorization, security auditing, data trend analysis, data tracking, business intelligence (BI) charts, performance optimization, and server management. You can use DMS to log on to the ApsaraDB for MongoDB console and obtain a list of ApsaraDB for MongoDB instances for remote access and online management.

11.3.3. Scenarios

ApsaraDB for MongoDB is a MongoDB-compatible document-oriented database service that can store large amounts of data and provides features such as security auditing and backup and restoration. It is widely used in scenarios such as IoT and gaming.

Read/write splitting

ApsaraDB for MongoDB uses a three-node replica set architecture to ensure high availability. Three data nodes are deployed on different physical servers and automatically synchronize data. The primary and secondary nodes are configured with different endpoints. MongoDB drivers allocate read/write requests to these nodes.

Flexible business scenarios

ApsaraDB for MongoDB has no schema and is suitable for startups that do not want to go through the hassles of changing table schemas. You can store structured data that has a fixed schema in ApsaraDB RDS, data that has flexible schemas in ApsaraDB for MongoDB, and hot data in ApsaraDB for Redis or ApsaraDB for Memcache. This ensures that business data can be stored and retrieved in an efficient manner to reduce data storage costs.

Mobile apps

ApsaraDB for MongoDB supports two-dimensional spatial indexes. Therefore, ApsaraDB for MongoDB can provide support for location-based apps. ApsaraDB for MongoDB uses a dynamic storage method, which is suitable for storing heterogeneous data from multiple systems and meets the needs of mobile apps.

IoT scenarios

- ApsaraDB for MongoDB provides high performance and allows data to be asynchronously written.
- Three types of components are available in ApsaraDB for MongoDB sharded cluster instances: mongos, shard, and Configserver nodes. You can configure the number and specifications of mongos and shard nodes when you create sharded cluster instances to provide different levels of performance. This makes ApsaraDB for MongoDB suitable for IoT scenarios that involve highly concurrent write operations.
- ApsaraDB for MongoDB provides the secondary index feature for dynamic queries. It can use the MapReduce aggregation framework of MongoDB to conduct multidimensional data analysis.

Applications in various fields

- Gaming

ApsaraDB for MongoDB can be used as a database service for game servers to store user information. The in-game equipment and credits of players are directly stored in the form of an embedded document to facilitate queries and updates.

- Logistics

ApsaraDB for MongoDB can store order information. The order status is constantly updated during the shipping process and is stored in the form of an embedded array in ApsaraDB for MongoDB. You can read all the changes in an order in a straightforward manner by performing a single query.

- Social networking

- ApsaraDB for MongoDB can store the information of users and their published WeChat moments. Geographical location indexes can be used to search nearby people and places.
- Additionally, ApsaraDB for MongoDB is suitable for storing chat history because ApsaraDB for MongoDB provides rich query abilities and is fast in both writing and reading.

- Live video streaming

ApsaraDB for MongoDB can store user information and gift information.

11.4. KVStore for Redis

KVStore for Redis is a database service that is compatible with the open source Redis protocol. KVStore for Redis provides cluster instances and allows you to deploy your instances in hot standby mode. With these features, KVStore for Redis can scale to meet high-throughput and low-latency requirements.

11.4.1. Product details

KVStore for Redis supports multiple versions of the Redis storage engine, including Redis 4.0 and 5.0. KVStore for Redis provides multiple architectures and offers features such as persistent data storage, high availability, elastic scalability, and intelligent O&M.

Flexible architectures

Standard architecture

An instance that uses the standard architecture contains a master node and a replica node. Data is synchronized between the two nodes in real time. The master node serves your workloads, and the replica node stays in hot standby mode to ensure high availability. If the master node fails, the system switches the workloads to the replica node to ensure that your business continues to run smoothly.

The standard architecture is applicable to the following scenarios:

- Support for more native Redis features.
- Persistent data storage in KVStore for Redis instances.
- Stable query rate for a single KVStore for Redis instance.
- Use of simple Redis commands, where only a few sorting and computing commands are required.

Cluster architecture

A cluster instance contains proxy nodes, data shards, and config servers. You can scale out a cluster instance by increasing the number of data shards. A cluster master-replica instance contains multiple data shards. Each data shard works in a high availability architecture in which a master node and a replica node are deployed on different devices. If the master node is faulty, the cluster master-replica instance fails over to the replica node to ensure high service availability.

The cluster architecture is applicable to the following scenarios:

- Large data volume.
- High queries per second (QPS).
- High throughput and high performance.

Read/write splitting architecture

A read/write splitting instance contains proxy nodes, master and replica nodes, and read replicas. Read replicas support chained replication. This allows you to scale out read replicas to increase the read capacity.

The read/write splitting architecture is applicable to the following scenarios:

- High QPS scenarios, such as data hotspots.
- Support for more native Redis features. Read/write splitting instances are free of the limits of cluster instances.

Data security

Data backup and restoration

- Data backup: KVStore for Redis uses Redis Database (RDB) backup snapshots to persist data. One backup is automatically created each day from the replica node based on the default backup policy. You can modify the backup policy or manually create a temporary backup.
- Data restoration: You can restore data from a specified backup set to the current instance or a new instance. When you restore data to a new instance, the data in the new instance is the same as that in the backup set. This feature is applicable to scenarios such as data restoration, quick business deployment, and data verification.

- Download of backup files: Backup files of KVStore for Redis are retained for seven days. If you want to retain backup files for more than seven days, you can download the backup files to your on-premises device. For example, you may want to retain backup files for more than seven days due to regulatory or security requirements.

Multi-layer network security protection

- Virtual private clouds (VPCs) are supported. VPCs are logically isolated from each other to provide higher security and performance.
- The anti-DDoS feature is supported to monitor and guard against Distributed-Denial-of-Service (DDoS) attacks.
- Each whitelist can block access from up to 1,000 IP addresses. You must configure a whitelist before you connect to an instance.
- Password authentication is provided out of the box to ensure secure and reliable access. For an instance, you can create up to 20 accounts. You can grant appropriate permissions to accounts based on business requirements. This allows you to flexibly manage instances and prevent accidental operations.
- SSL encryption is supported. You can install an SSL certificate on your application server after SSL encryption is enabled. This helps you encrypt network connections at the transport layer to improve data security and ensure data integrity.

In-depth kernel optimizations

Alibaba Cloud performed in-depth engine optimizations on Redis source code to prevent out of memory (OOM) issues and fix security vulnerabilities.

High availability

Master-replica deployment

Data shards use the master-replica architecture. The master and replica nodes implement real-time data synchronization by using both RDB and append-only file (AOF) persistence mechanisms. The master node serves your workloads, and the replica node stays in hot standby mode to ensure high availability. If the master node fails, the system switches the workloads to the replica node to ensure that your business continues to run smoothly.

Redundancy design and automatic detection

- A redundancy design is implemented for each system component to eliminate single points of failure and ensure service continuity.
- The system automatically detects hardware failures. In the case of failures, the system performs a failover and restores services within seconds.

Elastic scalability

KVStore for Redis provides flexibility in adjusting instance configurations to meet performance and capacity requirements as business needs evolve. If the performance of an instance becomes insufficient or excessive after the instance is created, you can change the architecture or memory specifications of the instance.

Intelligent O&M

Performance monitoring

KVStore for Redis provides abundant performance monitoring metrics such as CPU utilization and connections. You can query the monitoring data during a specified period of time within the past month. This helps you check the health status and troubleshoot issues of KVStore for Redis instances.

Visualized management

The KVStore for Redis console is a web-based visual management console that provides rich O&M and management features such as data backup and parameter settings. You can manage instances in a convenient and visualized manner.

Database kernel version management

KVStore for Redis continuously optimizes the kernel and fixes security vulnerabilities to improve service stability. You can update an instance to the latest minor version with a few clicks in the KVStore for Redis console.

11.4.2. Benefits

KVStore for Redis provides features such as ultra-high performance, elastic scale-out, resource isolation, high data security, high availability, and ease of use.

Ultra-high performance

- Supports cluster instances with a memory capacity of 128 GB or larger. These instances can meet large capacity and high performance requirements.
- Supports master-replica instances with a maximum memory capacity of 32 GB. The instances can meet common capacity and performance requirements.
- Supports CPUs, disks, memory, and network interface controllers (NICs) of different specifications in a cluster without affecting the operational performance of the cluster. This ensures compatibility with your existing devices.

Elastic scale-out

- Easy scaling: You can scale the instance storage capacity with only a few clicks in the KVStore for Redis console.
- Online scaling: You can scale the instance storage capacity without service interruption.

Resource isolation

- Supports instance isolation. This ensures the stability of individual services.
- Supports multi-tenant isolation. Each instance can use exclusive resources, such as CPUs, memory, I/O resources, and disk capacity.
- Supports multi-tenant parallel execution on a cluster by using multiple instances. Tasks from tenants are submitted to queues in different instances for execution. KVStore for Redis isolates resources among tenants by using instances.

High data security

- Data persistence: KVStore for Redis supports high-speed data read/write capabilities and provides data persistence by using the hybrid storage of memory and disks. KVStore for Redis allows you to load data from a persistent database into a cache database.
- Master/replica backup: KVStore for Redis maintains two backup copies of all data on master and replica nodes to prevent data loss.
- Access control: KVStore for Redis supports password authentication to ensure secure and reliable access to databases.
- Data transmission encryption: KVStore for Redis supports encryption based on SSL and Transport Layer Security (TLS) to secure data transmission.

High availability

- Master-replica architecture: Each instance runs in the master-replica architecture to eliminate single points of failure and ensure high availability.

- Automatic failure detection and recovery: The system automatically detects hardware failures and performs a failover within seconds. This minimizes the impacts of sudden hardware failures.
- Automatic fault tolerance: KVStore for Redis supports automatic fault tolerance for server disk failures in a cluster and hot swapping of disks. If a disk fails, services can be recovered within 2 minutes.

Ease of Use

- Supports Redis commands. You can use a Redis client to connect to a KVStore for Redis instance and manage data.
- Supports multiple commands in each query.

Permission management

- Supports data access permission management. Permissions include logon permissions, table creation permissions, read and write permissions, and whitelist management permissions.
- Allows you to log on to the KVStore for Redis console to manage permissions, including setting administrative permissions.
- Provides a unified permission management feature. This feature allows you to manage various permissions for each component of the system in the KVStore for Redis console. This isolates common users from internal permission management details, simplifies the permission management for administrators, and improves user experience.
- Allows you to manage multiple tenants in a centralized manner in the KVStore for Redis console. For example, you can dynamically configure and manage tenant resources, isolate resources, view statistics on resource usage, and manage tenants of multiple levels.

Scheduling

Supports multi-tenant scheduling across multiple clusters and resource pools.

11.4.3. Scenarios

KVStore for Redis is ideal for use in industries such as gaming, livestreaming, and e-commerce.

Gaming applications

KVStore for Redis can serve as an important architecture component in the gaming industry.

Use KVStore for Redis as storage

Gaming applications can be deployed in a simple architecture, in which the main program runs on an Elastic Compute Service (ECS) instance and the business data is persistently stored in a KVStore for Redis instance. KVStore for Redis can be used for persistent storage. It uses the master-replica architecture to implement redundancy.

Use KVStore for Redis as cache to accelerate application access

You can use a KVStore for Redis instance as a cache to accelerate connections for applications and use an ApsaraDB RDS instance as a backend storage database.

The high availability of KVStore for Redis is essential. If your KVStore for Redis service becomes unavailable, backend databases may be overwhelmed by requests sent from applications. KVStore for Redis adopts the master-replica architecture to ensure high availability. In this architecture, the master node provides services for your business. If this node fails, the system automatically switches workloads to the replica node. The complete failover process is transparent.

Livestreaming applications

Livestreaming heavily relies on KVStore for Redis to store user data and chat records.

Use hot standby mode to ensure high availability

KVStore for Redis can be deployed in the master-replica architecture in which the master node serves your workloads and the replica node stays in hot standby mode. This significantly improves service availability.

Use cluster instances to ensure high performance

KVStore for Redis provides cluster instances to eliminate the performance bottleneck of the native Redis single-threaded mechanism. Cluster instances can effectively handle traffic spikes to ensure high performance.

Use high scalability to support traffic spikes

KVStore for Redis allows you to deal with traffic spikes during peak hours by scaling out an instance with a few clicks. The scale-out is completely transparent to users.

E-commerce applications

In the e-commerce industry, KVStore for Redis is widely used in modules such as commodity presentation and recommendation.

Build shopping systems that support flash sales

An online shopping system may be overwhelmed by user traffic during large flash sales. Most databases cannot handle the heavy load.

To resolve this issue, you can use KVStore for Redis for persistent storage.

Build inventory management systems that support stocktaking

A KVStore for Redis instance can be used to count the inventory, and an ApsaraDB RDS instance can be used to store item quantity statistics. This way, the KVStore for Redis instance reads quantity statistics, and the ApsaraDB RDS instance stores quantity statistics. The KVStore for Redis instance is deployed on a physical server. The system provides a high-level data storage capability based on high-performance SSDs.

11.5. What is DTS?

Data Transmission Service (DTS) is a data service that is provided by Alibaba Cloud. DTS supports data transmission between various types of data sources, such as relational databases and big data systems.

Features

DTS has the following advantages over traditional data migration and synchronization tools: high compatibility, high performance, security, reliability, and ease of use. DTS allows you to simplify data transmission and focus on business development.

Feature	Description
---------	-------------

Data migration	You can use DTS to migrate data between homogeneous and heterogeneous data sources. This feature applies to the following scenarios: data migration to Alibaba Cloud, data migration between instances within Alibaba Cloud, and database splitting and scale-out.
Data synchronization	You can use DTS to synchronize data between data sources. This feature applies to the following scenarios: disaster recovery, data backup, load balancing, cloud BI systems, and real-time data warehousing.
Change tracking	You can use DTS to track data changes from user-created MySQL databases, ApsaraDB RDS for MySQL instances, and PolarDB-X instances (formerly known as DRDS) in real time. This feature applies to the following scenarios: cache updates, business decoupling, asynchronous data processing, synchronization of heterogeneous data, and synchronization of extract, transform, and load (ETL) operations.

11.5.1. Overview

This topic provides an overview of the features of Data Transmission Service (DTS), such as data synchronization, data migration, and change tracking.

Data synchronization

You can use DTS to synchronize data between data sources in real time.

The following table describes the capabilities supported by the data synchronization feature.

Capability	Description
Add or remove the objects to be synchronized	You can add or remove the objects to be synchronized at any time when a data synchronization task is running.
View and analyze the synchronization performance	DTS provides trend charts that allow you to view and analyze the performance of your data synchronization tasks. The synchronization performance is measured based on bandwidth, records per second (RPS), and synchronization latency.
Monitor data synchronization tasks	DTS allows you to monitor the status of data synchronization tasks. If the threshold for synchronization latency is reached, the system sends an alert notification to you. You can set the alert threshold based on the sensitivity of your business to synchronization latency.

Data migration

You can use DTS to migrate data between homogeneous or heterogeneous data sources. DTS provides the following extract, transform, and load (ETL) capabilities: object name mapping for columns, tables, and databases, and data filtering. These capabilities allow you to migrate data between various types of data sources with ease.

DTS uses online migration. You need to only configure the source instance, the destination instance, and the objects to be migrated. DTS automatically completes the entire data migration process. You can select all of the supported migration types to minimize the impact of online data migration on your services. However, you must ensure that DTS servers can connect to both the source and destination instances.

The following table describes the ETL capabilities supported by the data migration feature.

Capability	Description
------------	-------------

Enable object name mapping for columns, tables, and databases	You can migrate data between two columns, tables, or databases that have different names.
Filter the data to be migrated	You can set SQL conditions to filter the data to be migrated in a specific table. For example, you can specify a time range to migrate only the latest data.

Change tracking

DTS allows you to track incremental data in databases in real time. You can consume the tracked data based on your business requirements.

The following table describes the capabilities supported by the change tracking feature.

Capability	Description
Add or remove the objects for change tracking	You can add or remove the objects for change tracking at any time when a change tracking task is running.
View the tracked data	You can view the incremental data that is tracked by a change tracking task in the DTS console.
Modify the consumption checkpoint	You can modify the consumption checkpoint at any time.
Monitor change tracking tasks	DTS allows you to monitor the status of change tracking tasks. If the threshold for consumption latency is reached, the system sends an alert notification to you. You can set the alert threshold based on the sensitivity of your business to consumption latency.

ETL

DTS provides the ETL feature. When you configure a data migration or synchronization task, you can rename databases, tables, and columns in the destination instance and set conditions to filter specific data.

The following table describes the capabilities supported by the ETL feature.

Capability	Description
High computing effectiveness	Integrated with the powerful capabilities of DTS in collecting streaming data from databases, the ETL feature ensures data accuracy and provides high computing effectiveness in the industry.
Flexible task monitoring and management	You can monitor and manage ETL tasks in the DTS console. For example, you can start a task, stop a task, and view task details.

Data consistency

DTS uses the data reading module, data loading module, data verification feature, and resumable transmission feature to ensure data consistency between the source and destination databases.

11.5.2. Benefits

DTS supports transmitting data between data sources such as relational databases and OLAP databases. DTS provides you with multiple data transmission methods such as data migration, real-time data subscription, and real-time data synchronization. Compared with other third-party data migration and synchronization tools, DTS provides multiple transmission channels with high performance, security, and reliability. DTS also makes it easy to create and manage transmission channels.

Diverse transmission methods

DTS supports multiple data transmission features, including data migration, data subscription, and data synchronization. In data subscription and data synchronization, data is transmitted in real time.

Data migration enables you to migrate data between databases without interrupting application operations. The application service downtime during data migration is reduced to minutes.

High performance

DTS uses servers with high specifications to ensure high data transmission performance for each synchronization or migration channel.

At the underlying layer, multiple measures are taken to improve DTS performance.

Compared with traditional data synchronization tools, the real-time synchronization feature of DTS enables you to concurrently transmit transactions. It also allows you to synchronize table data you want to update at a time. This greatly improves synchronization performance.

High security and reliability

DTS is implemented using clusters. If a node in a cluster is down or faulty, the control center quickly moves all tasks from this node to another healthy node in the cluster.

DTS provides a 24 x 7 mechanism for validating data accuracy in some transmission channels to quickly locate and correct incorrect data. This helps ensure reliable data transmission.

Secure transmission protocols and tokens are used for authentication across DTS modules to ensure reliable data transmission.

Easy-to-use

The DTS console is a visual management interface that provides a wizard-like process to assist you in creating data transmission channels.

You can also view data transmission information in the DTS console, including the transmission status, progress, and performance, to better manage the transmission channels.

DTS supports resumable transmission, and regularly monitors channel status to avoid interruptions resulting from network or system exceptions. When DTS detects a channel exception, it automatically repairs or restarts the channel. In cases where manual operations are needed, you can directly repair the channel and restart it in the DTS console.

11.5.3. Scenarios

DTS supports multiple features including data migration, real-time data subscription, and real-time data synchronization to meet the following scenarios.

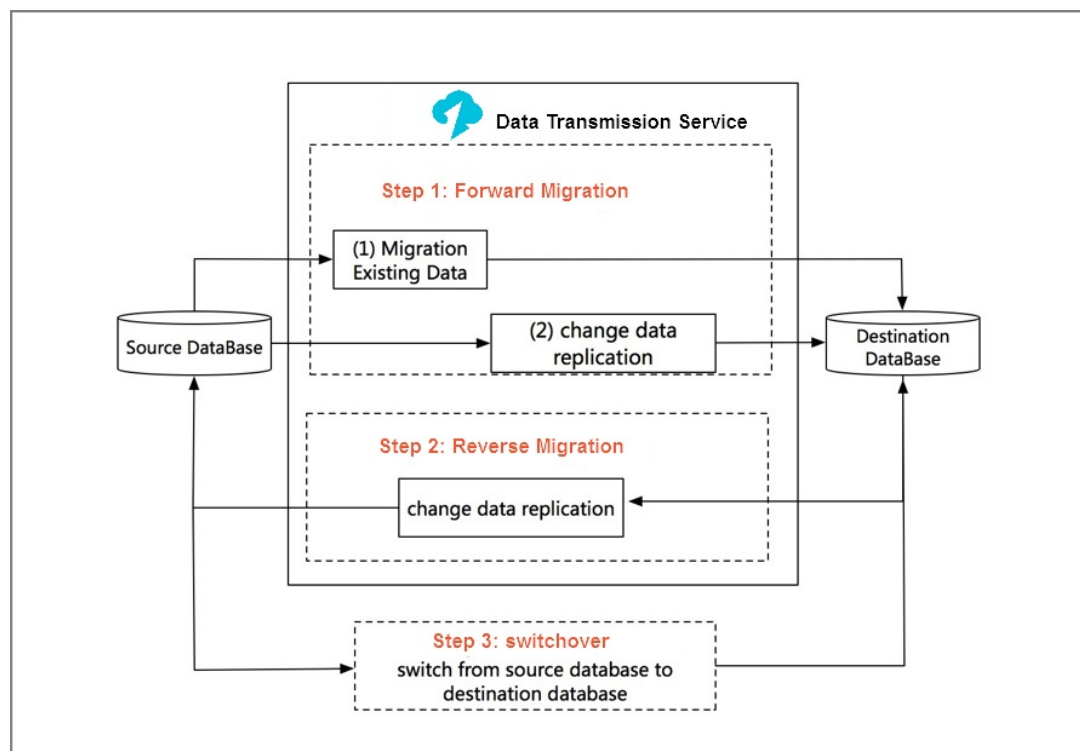
Migration with service downtime reduced to minutes

Many users seek for a way to migrate systems without affecting their services. However, data changes if services are not suspended during the migration. To ensure data consistency, many third-party migration tools require that the service be suspended during data migration. It may take hours or even days throughout the migration and result in a significant loss in service availability.

To reduce the barrier of database migration, DTS provides an interruption-free migration solution that minimizes the service downtime to minutes.

Interruption-free migration shows how interruption-free migration works.

Figure 1. Interruption-free migration



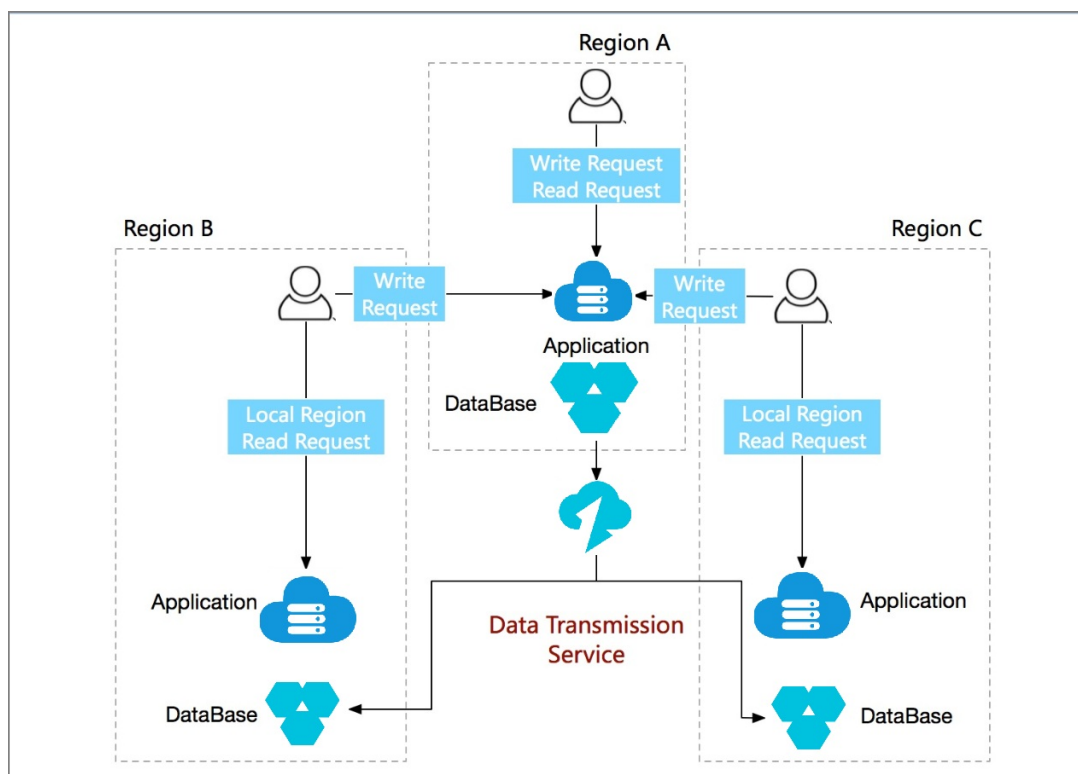
The interruption-free data migration process involves schema migration, full data migration, and incremental data migration. In the incremental data migration phase, data is synchronized between the source and destination instances in real time. You can validate the service in the destination database. After the validation is complete, the service is migrated to the destination database. The entire system is then eventually migrated.

Throughout the migration process, the service experiences interruptions only when it is switched from the source instance to the destination instance.

Accelerated access to global services to empower cross-border businesses

If services with widely distributed users, such as global services, are deployed only in one region, users in other regions have to access them remotely, resulting in high access latency and poor user experience. To accelerate the access to global services and improve access experience, you can adjust the architecture, as shown in **Reduced cross-region access latency**.

Figure 2. Reduced cross-region access latency



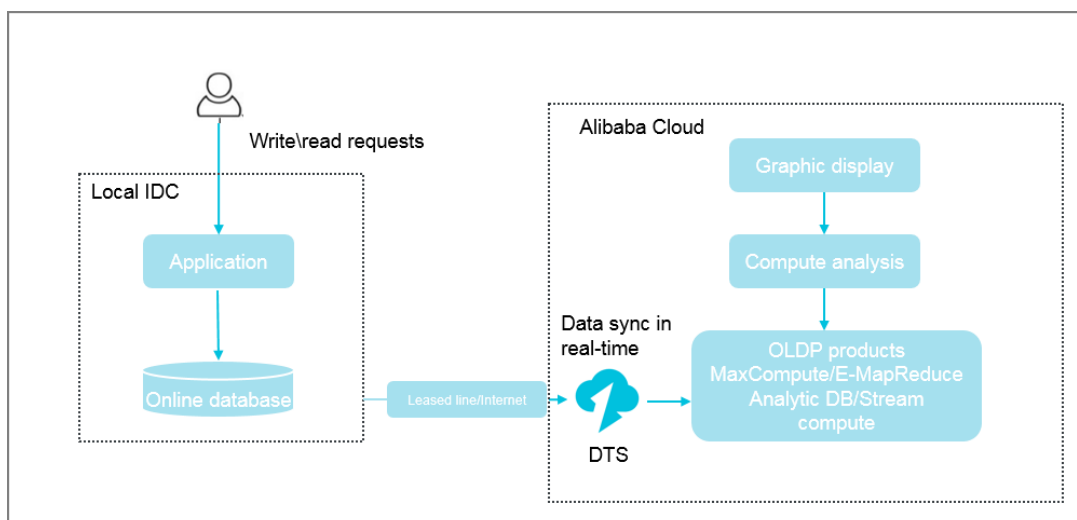
This architecture consists of one center and multiple units. Write requests of users in all regions are routed back to the center. DTS synchronizes data in the center to all units. Read requests of users in different regions can be routed to nearby units to avoid remote access and reduce access latency. In this way, access to global services is accelerated.

Custom cloud BI system built with more efficiency

User-created business intelligence (BI) systems cannot meet the increasing demand for real-time performance and are difficult to manipulate. With the Apsara Stack BI architecture, you can quickly build a BI system without affecting the current architecture. For this reason, more and more users choose to build BI systems that meet their own business requirements on Apsara Stack.

DTS can help you synchronize data stored in local databases to an Apsara Stack BI system (such as MaxCompute or StreamCompute) in real time. You can then perform subsequent data analysis with various compute engines while viewing the computing results in real time with a visualization tool. You can also synchronize those results back to the local IDC with a migration tool. [Cloud BI architecture](#) shows the implementation architecture.

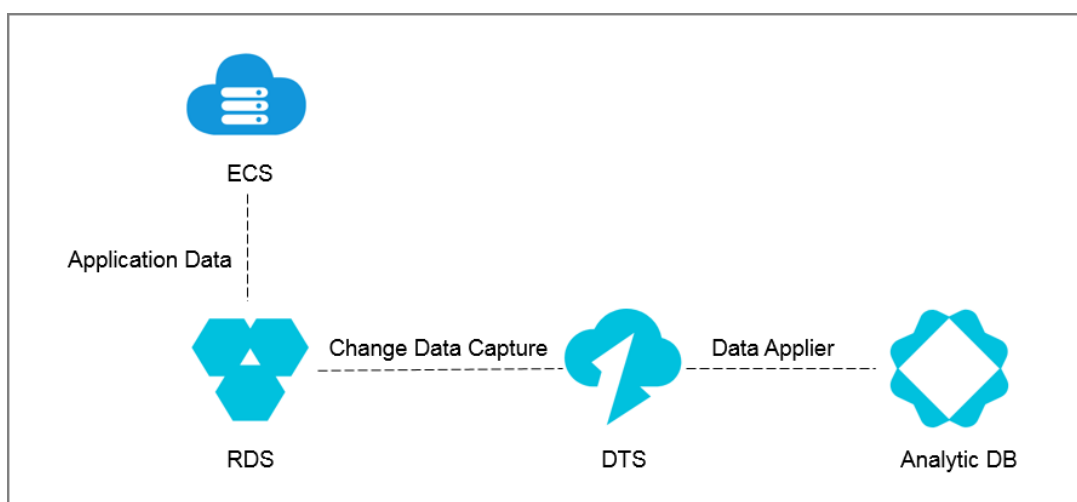
Figure 3. Cloud BI architecture



Real-time data analysis to rapidly respond to market conditions

Data analysis is essential in improving enterprise insights and user experience. Real-time data analysis enables enterprises to adjust marketing strategies more quickly and flexibly so that they can adapt to the rapidly changing marketing conditions and demands for higher user experience. To implement real-time data analysis without affecting online services, service data needs to be synchronized to the analysis system in real time. For this reason, acquiring service data in real time becomes essential. In DTS, the data subscription feature can help you acquire real-time incremental data without affecting online services and synchronize the data to the analysis system using the SDK for real-time data analysis, as shown in [Real-time data analysis](#).

Figure 4. Real-time data analysis

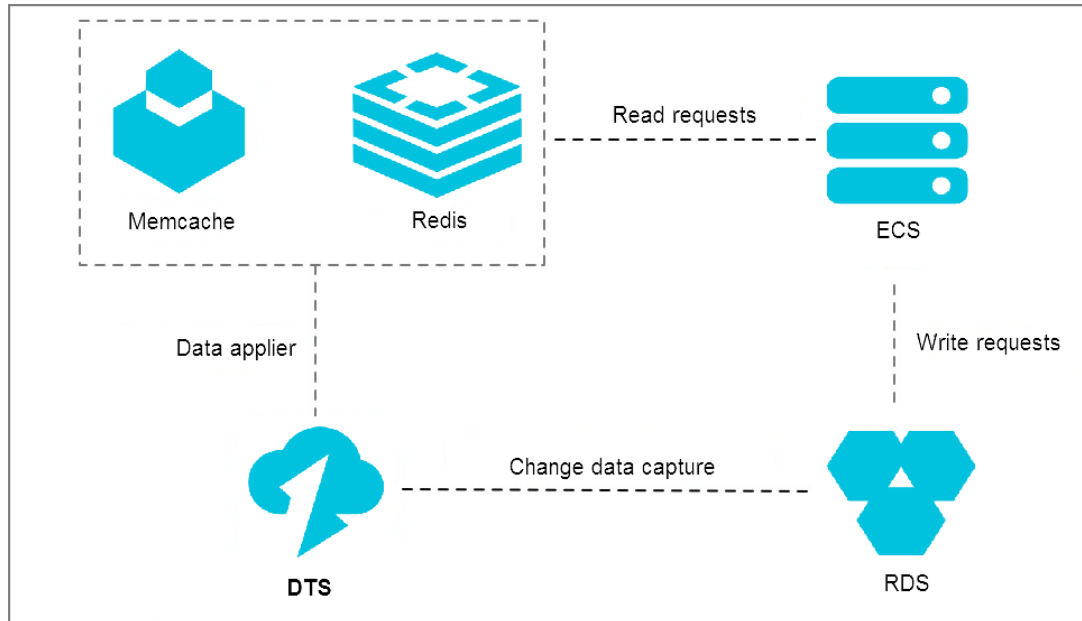


Lightweight cache update policies to make core services more simple and reliable

To accelerate service access and improve concurrent read performance, many enterprises introduce the caching layer to the service architecture. In this architecture, all the read requests are routed to the caching layer, and the memory reading mechanism greatly improves read performance. Cached data cannot persist. If caching ends abnormally, data in the cache memory is lost. To ensure data integrity, the updated service data is kept in a persistent storage medium, such as a database.

In this condition, the service data is inconsistent between the cache and the persistent databases. The data subscription feature can help asynchronously subscribe to the incremental data in those databases and update the cached data to implement lightweight cache update policies. [Cache update policies](#) shows the architecture of these policies.

Figure 5. Cache update policies



Cache update policies offer the following benefits:

- Quick update with low latency

Cache invalidation is an asynchronous process, and the service returns data directly after the database update is complete. For this reason, you do not need to consider the cache invalidation process, and the entire update path is short with low latency.

- Simple and reliable applications

The complex doublewrite logic is not required for the application. You only need to start the asynchronous thread to monitor the incremental data and update the cached data.

- Application updates without extra performance consumption

Because data subscription acquires incremental data by parsing incremental logs in the database, the acquisition process does not damage the performance of services and databases.

Asynchronous service decoupling to make core services simpler and more reliable

Data subscription optimizes intensive coupling to asynchronous coupling by using real-time message notifications. This makes the core service logic simpler and more reliable. This application has been widely implemented in Alibaba. Tens of thousands of downstream services in the Taobao ordering system acquire real-time data updates through data subscription to trigger the business logic every day.

The following uses a simple example to describe the benefits of implementing data subscription in this scenario.

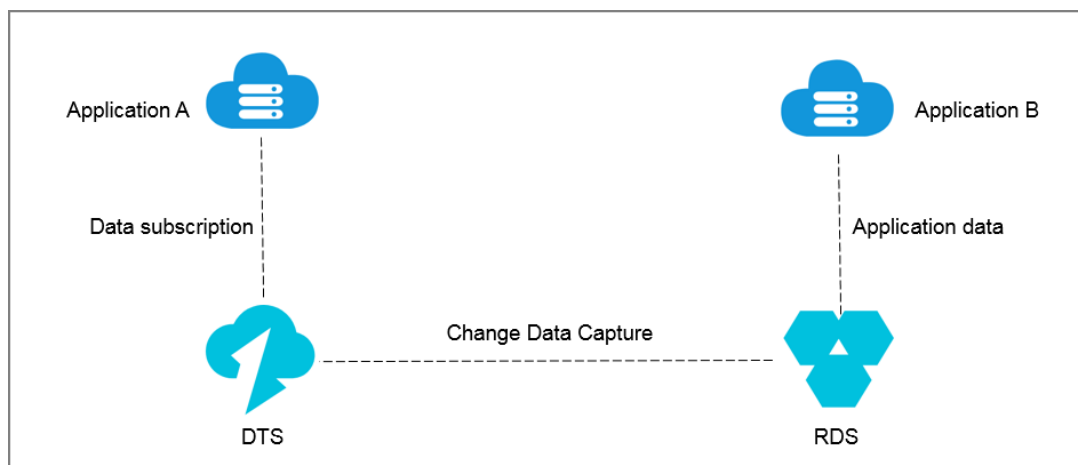
The e-commerce industry involves multiple services including the order management system, inventory management, and the shipping of goods. An ordering process with all of those services included is as follows: After a user places an order, downstream services including seller inventory notification and goods shipping are modified. When all logic modifications are complete, the order result is returned to the user. However, this ordering logic has the

following issues:

- The lengthy ordering process results in poor user experience.
- The system is unstable and any downstream fault directly affects the availability of the ordering system.

To improve user experience of core applications, you can decouple the core applications and the dependent downstream services so that they can work asynchronously. In this way, the core applications become more stable and reliable. [Asynchronous service decoupling](#) shows how to adjust the logic.

Figure 6. Asynchronous service decoupling



The ordering system returns the order result directly after order placement. With DTS, the underlying layer acquires the updated data from the ordering system in real time. Then, the downstream service subscribes to the modified data using the SDK and triggers the service logic such as inventory and shipping. In this way, the ordering system becomes simpler and more reliable.

Horizontal scaling to improve read performance and quickly adapt to business growth

A single RDS instance may not be able to support a large number of read requests, which may affect the main service process. To elastically improve the read performance and reduce database workload, you can create read-only instances using the real-time synchronization feature of DTS. These read-only instances take on large amounts of the database reading workload and expand the throughput of applications.


11.6. What is DMS?

Data Management (DMS) is a fully managed service that is provided by Apsara Stack. You can use this service to manage data, table schemas, R&D processes, R&D specifications, users, permissions, and access security.

Supported databases

- Relational databases:
 - MySQL: ApsaraDB RDS for MySQL, PolarDB-X, MySQL databases from other cloud service providers, and self-managed MySQL databases
 - SQL Server: ApsaraDB RDS for SQL Server, SQL Server databases from other cloud service providers, and self-managed SQL Server databases
 - PostgreSQL: ApsaraDB RDS for PostgreSQL, PostgreSQL databases from other cloud service providers, and self-managed PostgreSQL databases

- Self-managed Dameng (DM) databases
- Self-managed Oracle databases
- ApsaraDB for OceanBase and self-managed OceanBase databases
- NoSQL databases:
 - Redis: ApsaraDB for Redis, Redis databases from other cloud service providers, and self-managed Redis databases
 - MongoDB: ApsaraDB for MongoDB, MongoDB databases from other cloud service providers, and self-managed MongoDB databases
 - Graph Database (GDB)
- Online analytical processing (OLAP) databases:
 - AnalyticDB for MySQL
 - AnalyticDB for PostgreSQL

 **Note** Self-managed databases are databases that are installed on Apsara Stack Elastic Compute Service (ECS) instances, instances from other cloud service providers, or servers in data centers.

Features

- DMS provides support for the entire database development process. You can design table schemas in an on-premises environment based on the design specifications. Before you publish SQL statements to an online environment, DMS can review the add, remove, modify, and query operations in the statements. Then, you can publish schemas to the specified environment as needed.
- DMS provides fine-grained access control at the database, table, or field level. You can perform all operations on databases in the DMS console. The operations can be traced and audited.
- DMS allows you to configure operation specifications and approval processes for multiple modules based on your business requirements. These modules include the schema design, data change, data export, and permission application.
- DMS integrates database development with database interaction. You can manage databases without the need to switch between database endpoints at a high frequency by using database accounts and passwords.
- DMS provides the task orchestration feature that allows you to orchestrate and schedule SQL tasks for databases. You can use this feature to perform a variety of operations with ease. For example, you can use this feature to dump historical data or generate periodical reports.

11.6.1. Features

Data Management (DMS) provides seven features to cater to various scenarios. These features are: data asset management, SQL console, database development, data transmission service (DTS), security and compliance, solutions, and operations and maintenance (O&M) management.

Data asset management

- Instance management: You can add, edit, and view the information of database instances that are added to DMS. You can also perform management operations on these instances, including permission management, enable or disable instances, and delete instances.
- Data classification: You can classify and manage instances, databases, and tables in a database instance. This helps administrators, developers, and O&M engineers manage data in tables.

- Sensitive data management: You can manage all medium and high sensitivity data centrally from DMS. This improves control over sensitive data within enterprises.

SQL console

- Visual operation section: You can view all tables, fields, and indexes of the current database. You can also edit the schema of a table and import and export data from it.
- SQL execution section: You can write, format, and execute SQL statements. You can also modify or update result sets.
- Execution result section: You can view the execution results and execution history.
- Quick actions: You can perform quick actions such as obtaining a list of tables, metadata synchronization, export, table schema version management, operation auditing, risk auditing, and super SQL mode.

Database development

- Change schemas
 - Schema design: You can design table schemas for databases or tables. These table schemas conform to the predefined development specifications. You can customize development processes for different lines of business based on your business requirements to ensure the consistency of schemas among multiple environments, such as the development environment, test environment, and production environment.
 - Schema synchronization: You can compare the schemas of different databases and generate scripts to ensure consistency across databases. This can be used to compare and synchronize table schemas across multiple database environments.
 - Shadow table synchronization: This feature can automatically create a shadow table based on the schema of a source table. You can use this feature for end-to-end stress testing.
 - Empty database initialization: You can use this feature to synchronize the schemas of a database to an empty database. This allows you to quickly initialize databases across regions and branches.
 - Table consistency management: You can use this feature to compare a batch of tables with a base table, and generate and execute scripts to repair consistency issues between the tables. You can use this to ensure schema consistency between the two environments, or to compare table schemas between logical databases.
- Change data
 - General data operations: You can use SQL statements such as INSERT, UPDATE, DELETE, and TRUNCATE for data initialization, historical data clearing, error fixing, and feature test.
 - Historical data clearing: You can clear historical data periodically to prevent data accumulation from impacting the stability of the environment.
 - Data import: You can quickly import a large amount of data to databases. This helps you save labor and resources.
- Lockless changes
 - Lockless data change: You can use this feature to ensure execution performance and reduce the impact of the SQL statement on database performance or database storage. This feature is especially useful when you need to change a large amount of data.
 - Lockless schema change: You can use this feature to change schemas without the need to lock tables. This prevents table locking and synchronization latency between primary and secondary databases.
- Data export
 - SQL result set export: You can export the result set of an SQL query statement.

- Database export: You can export a database or individual table. You can also export only the schemas or table data of a database or individual table for data analysis.
- SQL review: This feature can review the submitted SQL statements and provide optimization suggestions. This feature checks whether your queries use indexes and conform to database development standards, helping you write efficient and secure queries and reducing the risk of SQL injection attacks.
- Database clone: You can use this feature to replicate and synchronize database data and initialize databases in multiple environments.
- Test data generation: You can use this feature to generate a large amount of test data, such as random values, region names, and virtual IP addresses. This helps you prepare test data with ease.
- DevOps: You can use this feature to customize development processes and manage the quality of these processes. This can help you effectively implement development processes, reduce accidental operations, protect data, and thus improve development efficiency.

Data Transmission Service (DTS)

- Data transmission and migration
 - Data migration: You can use this feature to migrate data between homogeneous and heterogeneous data sources. This is suitable for data migration to Alibaba Cloud, data migration between instances in Alibaba Cloud, and database partitioning and scaling.
 - Change tracking: You can use this feature to track data changes from databases in real time. Then, you can consume the tracked data. You can use this feature in cache updates, business decoupling, synchronization of heterogeneous data, and synchronization of extract, transform, and load (ETL) operations.
 - Data synchronization: You can use this feature to synchronize data between data sources in real time. This feature can be used for active geo-redundancy, geo-disaster recovery, zone-disaster recovery, cross-border data synchronization, query load balancing, cloud BI systems, and real-time data warehousing.
 - Advanced Database & Application Migration (ADAM): A solution for smooth database migration. You can use ADAM to evaluate the feasibility and costs of database migration and obtain recommendations on databases. ADAM helps reduce the risks, technical difficulties, and turnaround time of migrating databases and applications.
- Data integration, development, and application
 - Batch processing: This feature provides low-code data development tools to process complex big data in batches. This feature can be used for fine-grained enterprise operations, data-driven marketing, intelligent recommendation, and other big data scenarios.
 - Stream processing: This feature provides real-time data processing tools that extract, transform, process, and load streaming data. This feature provides more options for real-time data processing and computing, which helps enterprises accelerate digital transformation.
 - Task orchestration: This feature supports complex task orchestration and scheduling to improve the efficiency of data development.
 - Data service: You can use this feature to expose data managed on DMS to other entities. This feature helps manage data security by controlling the scope of the exposed data to the required minimum.
 - Data visualization: This feature can visualize data in the database in a variety of ways to help generate insights for decision-making.

Security and compliance

- Permissions: You can manage the query, change, export, and logon permissions of instances, databases, tables, data columns, and rows.

- Security rules: DMS manages database rule sets in a fine-grained manner to create database operation specifications and development processes.
- Approval process: You can select and configure approval processes based on different user behaviors.
- Access IP whitelist: This feature allows you to effectively control the use of DMS so that only authorized employees in a trusted environment can access DMS.
- Operation audit: You can query the information about operations that are performed in DMS, including SQL statements that are used in the SQL console, tickets, logon information, and operation logs. This feature helps you troubleshoot database issues and provides data for operation audit.
- Sensitive data management: This feature helps you detect and identify sensitive data assets, and helps you protect these data assets from abuse.

Solutions

T+1 full data snapshot: This feature creates snapshots for individual tables every hour or day on a T+1 basis. This facilitates statistical analysis on a daily or monthly basis. For example, this feature can collect statistics and calculate the total amount of orders on the previous day to provide you up-to-date information about business operations.

O&M management

- User management: You can add and delete DMS users, manage user permissions, and grant permissions on instances, databases, tables, rows, and sensitive columns.
- Task management: You can create SQL tasks and manage existing tasks.
- Configuration management: You can change the system configurations of DMS to suit the requirements of your business at any time. For example, you can enable automatic synchronization of instance resources.
- Database grouping: You can add multiple databases with the same environment or engine type to one group. You can use this feature to apply a data change or schema design to all of the databases in a database group with ease.

11.6.2. Benefits

Data Management (DMS) provides multiple benefits, including various data sources, secure and controllable processes, and fine-grained permission management. These benefits improve data security and simplify data management.

Various data sources

- Relational databases:
 - MySQL: ApsaraDB RDS for MySQL, PolarDB-X, MySQL databases from other cloud service providers, and self-managed MySQL databases
 - SQL Server: ApsaraDB RDS for SQL Server, SQL Server databases from other cloud service providers, and self-managed SQL Server databases
 - PostgreSQL: ApsaraDB RDS for PostgreSQL, PostgreSQL databases from other cloud service providers, and self-managed PostgreSQL databases
 - Self-managed Dameng (DM) databases
 - Self-managed Oracle database
 - ApsaraDB for OceanBase and self-managed OceanBase databases

- NoSQL databases:

Redis: ApsaraDB for Redis, Redis databases from other cloud service providers, and self-managed Redis databases

MongoDB: ApsaraDB for MongoDB, MongoDB databases from other cloud service providers, and self-managed MongoDB databases

Graph Database (GDB)

- Online analytical processing (OLAP) databases: AnalyticDB for MySQL and AnalyticDB for PostgreSQL

Unified operations and comprehensive audits

- After you add a database instance to DMS as an administrator, you can perform the required operations in the DMS console. The operations include querying databases, changing schemas, and changing data.
- You can query and audit all historical operations based on multiple dimensions, including the operator, database, table, and time.

Fine-grained access control

Regular users do not need to use database accounts and passwords. These users only need to request the query, export, or change permission on the destination database, table, or field in the DMS console based on their business requirements. After a permission expires, DMS revokes the permission.

Custom approval processes

You can create custom approval processes for the modules of each database instance based on your business requirements. This allows you to meet requirements from several aspects, such as efficiency and security. Examples:

- To impose loose controls on a test environment, you can reduce stages or set no approval process.
- To impose strict controls on a production environment, you can specify an approval process that includes the required operations for the production environment. The production environment takes effect until the specified engineers approve all these operations in sequence.

Custom design specifications for schemas

You can create custom design specifications for MySQL table schemas. These design specifications include the field type, index type, number of indexes, field length, table size, and release process.

Simple procedure to schedule and orchestrate periodical tasks

DMS allows you to quickly create the required orchestration and recurring schedules for the SQL task nodes of various databases. You can use this feature to perform operations, such as transferring historical data and generating periodic reports, on databases to mine value from data.

11.6.3. Scenarios

Data Management (DMS) is ideal for implementing data security and improving development efficiency. It is typically used in scenarios such as cloud migration, disaster recovery, T +1 data warehouses, real-time data warehouses, and cross-instance queries.

Data security and development efficiency

You can use DMS to apply custom security rules, sensitive data identification rules, and masking rules. This helps improve the data security and reliability of your applications. You can also make use of scheduled tasks and data management windows to reduce the risk of issues that may arise due to human error.

- Pain points:
 - High reliance on intra-person communication to perform O&M tasks and implement changes, leading to low efficiency within enterprises.
 - Manual labor is unable to keep up with the O&M responsibilities of hundreds of instances, which include the detection and prevention of errors. This leads to challenges in maintaining business reliability.
- Benefits:
 - You can use DMS to centrally manage a large number of instances, improving development efficiency.
 - You can use DMS to build a data security protection system from scratch to ensure data security.

T +1 data warehouse

DMS allows you to create snapshots for individual tables every hour or day on a T+1 basis. This way, you can view the statistics on data by hour, day, or month.

- Pain points:
 - Traditional T+1 data warehouses use complex operations to aggregate incremental data, which consumes a large amount of computing resources. This negatively affects the timeliness of the data.
 - Incremental data is fetched in real time or at a scheduled time. This has a significant impact on business systems.
 - Data slicing problems at a given time point.
- Benefits:
 - History tables record the life cycle of incremental information in real time. DMS generates hourly and daily snapshots five times more efficient than traditional solutions.
 - DTS does not make use of plug-ins or log parsing to fetch data. This reduces the resource overhead to a minimum, with little to no impact on the performance of business systems.
 - Data can be sliced at any point-in-time.

Cross-instance query

DMS allows you to perform real-time join queries across online or heterogeneous data sources that are deployed in different environments. You can perform queries on MySQL databases, SQL Server databases, PostgreSQL databases, and Redis databases. Even when such database instances are deployed in different regions and environments, you can write a dynamic SQL statement to perform join queries on these databases.

Benefits:

- Cross-database queries after vertical and horizontal partitioning: You can perform real-time join queries for common commodity databases and order databases across multiple regions and units.
- Join queries on heterogeneous databases: For example, after data is migrated from a SQL Server database to a MySQL database, you can perform join queries to verify the consistency between both databases.
- Join queries in hybrid clouds: You can perform join queries on databases across on-premise IDC and cloud environments.

11.7. DBS

Database Backup (DBS) is a cost-efficient backup platform that provides continuous data protection for databases. DBS offers strong protection for data stored in multiple environments, such as enterprise data centers, third-party clouds, and Alibaba Cloud public cloud. DBS offers a comprehensive solution for data backup and restoration. DBS supports incremental backup in real time and a recovery point objective (RPO) within seconds.

DBS allows you to back up data in real time. When data changes online, DBS obtains the data change and backs up the changed data to cloud storage in real time. This helps you achieve an RPO within seconds.

11.7.1. Features

Database Backup (DBS) is a cost-efficient and highly reliable cloud-native database backup platform. You can create backup schedules, back up data, download backup sets, and restore data in the DBS console based on your business requirements.

Storage pool

DBS allows you to use storage pools to store backup data. In addition to the built-in storage space of DBS, you can select an Object Storage Service (OSS) bucket or a self-managed NAS file system as the storage pool. After you add a storage pool, you can specify a backup policy for the storage pool. This way, data backups that are generated based on the backup policy are stored in the storage pool.

Backup gateway

DBS provides the backup gateway feature. You can download the installation package in the DBS console and install the backup gateway on the on-premises database server to back up the database to OSS.

Backup method

Logical backup and physical backup are common data backup methods.

- Logical backup: backs up database objects such as tables, indexes, and stored procedures. Logical backup tools include MySQL mysqldump and Oracle exp and imp.
- Physical backup: backs up database files on the operating system. Physical backup tools include MySQL XtraBackup and Oracle RMAN.

Data restoration

DBS allows you to restore data at any point in time at the second level and select the objects to restore at any level.

Emergency recovery

DBS provides the emergency recovery feature. This feature allows you to create sandbox instances based on backup sets. Sandbox instances allow backup data to be immediately available. Read and write operations that are performed in each sandbox instance do not affect data in other sandbox instances or source databases.

11.7.2. Benefits

Database Backup (DBS) provides database backup solutions for a variety of environments. You can connect databases to DBS by using various methods such as Express Connect and the Internet. DBS requires minimal configurations to perform full backup or incremental backup, and restore data. DBS allows you to back up ApsaraDB RDS databases and databases that are deployed in self-managed data centers, on Elastic Compute Service (ECS) instances, or on third-party cloud platforms.

Comparison between DBS and self-managed backup systems

Item	DBS	Self-managed backup system
Cost	<ul style="list-style-type: none"> • Cold data is separated from hot data for tiered storage. This is suitable for the long-term archiving of backup data. • The compressed and compact backup formats and incremental backup can significantly reduce storage costs. 	<ul style="list-style-type: none"> • A large amount of upfront asset investment is required. • Storage space is limited by hard disk capacity. The storage space must be manually increased. • Single-line or double-line access is slow, and bandwidth is limited. The bandwidth must be manually increased during peak hours. • The introduction of multi-level storage media leads to a sharp increase in O&M costs.
Security	<ul style="list-style-type: none"> • DBS uses SSL and AES-256 encryption to secure backup data during transmission and storage. • Resources are isolated between different users, and geo-disaster recovery is supported. • DBS provides a variety of authentication and authorization methods, such as whitelist configuration, hotlink protection, and user management based on Resource Access Management (RAM). • DBS allows you to verify the validity of backups at any time, and notifies you of the task status. • Custom authentication is supported. 	<ul style="list-style-type: none"> • Additional scrubbing devices and black hole policy-related services are required. • A separate security mechanism is required.
Ease of use	<ul style="list-style-type: none"> • The process of configuring a backup schedule and running a backup task takes only 5 minutes. • Fine-grained backup is supported. You can back up data of different granularities based on your business requirements, including an entire instance, a single database, multiple tables, and a single table. • DBS supports global rules for the lifecycle management of backup data. You can customize rules to automatically dump, clean up, duplicate, and distribute backup data. • DBS provides a web-based GUI for you to perform backup and restore operations with ease. 	<ul style="list-style-type: none"> • The backup process requires complex scripts and tools, which are difficult to learn. • A self-managed backup system is not flexible and provides only basic capabilities in most cases.

Performance	<ul style="list-style-type: none"> DBS captures in-memory logs in real time and achieves a recovery point objective (RPO) within seconds. DBS allows you to restore backup data to any point in time. DBS allows you to select a single table as the object to restore. This greatly reduces the recovery time objective (RTO). DBS supports streaming backup. Data is not flushed to disks. The entire backup window is unlocked. The backup speed can be adjusted based on the concurrency configuration. 	The shortcomings of multiple tools used for backup are performance bottlenecks.
Reliability	<ul style="list-style-type: none"> DBS uses Apsara Distributed File System to provide a distributed storage service with high reliability. During the backup process, data integrity is verified in real time. Tested by a large number of users, DBS can efficiently detect and fix vulnerabilities. 	<ul style="list-style-type: none"> The mixed use of multiple tools causes high risks. A self-managed backup system is prone to errors due to low hardware reliability. If a disk has a bad sector, data may be lost.
Scalability	<ul style="list-style-type: none"> DBS allows you to back up Alibaba Cloud databases and self-managed databases that are deployed on ECS instances, in self-managed data centers, or on third-party cloud platforms such as Amazon Web Services (AWS) and Tencent Cloud to Alibaba Cloud. In addition to restoring data to the source database, DBS also allows you to restore backup data to other environments. For example, you can restore an on-premises database to an Alibaba Cloud database by using DBS. 	Self-managed backup systems support only specific environments and are generally not scalable.

Low costs

DBS uses Object Storage Service (OSS) as built-in storage. Backup data is converted to a dedicated format, compressed, and then saved to the built-in storage. This reduces storage costs.

High security

Feature	Description
Encrypted transmission and storage	DBS uses SSL and AES-256 encryption to secure backup data during transmission and storage.
Geo-redundancy	This feature enhances the level of data protection.
Alerting	DBS sends notifications about key events such as backup errors, restoration errors, and restoration success.

Flexibility and ease of use

Feature	Description
Fine-grained backup	DBS allows you to back up data of different granularities based on your business requirements, including an entire instance, multiple databases, a single database, multiple tables, and a single table.
Single-table restoration	DBS allows you to select a single table as the object to restore. This reduces the RTO.
Lifecycle management	DBS globally controls lifecycle rules for backup data to automatically dump, clean up, duplicate, and distribute backup data.
GUI	DBS provides a web-based GUI for you to perform backup and restore operations with ease. You can purchase and configure a backup schedule and run a backup task on the GUI. The whole process takes only 5 minutes.

High performance

DBS reads and parses database logs in real time by using the real-time data streaming technology of Alibaba Cloud. Then, DBS stores data in the cloud to perform incremental backup. DBS keeps the latency within seconds during incremental backup. The latency varies based on network conditions.

DBS allows you to restore a database within seconds by using incremental backups. This ensures the security and integrity of your data.

Feature	Description
Real-time backup	DBS captures in-memory logs in real time and achieves an RPO within seconds.
Parallel backup	DBS can back up data in an unlocked manner, use multiple threads to back up data in parallel, and adaptively shard data during data pulling.
Restoration to any point in time	DBS provides a calendar and a timeline so that you can select a point in time to which data is restored.
Multiple specifications	DBS provides high scalability to seamlessly support the performance requirements of enterprises at different stages.

11.7.3. Scenarios

Database Backup (DBS) provides the full backup, incremental backup, and data restoration features to help you implement data backup and restoration in various scenarios.

Real-time backup

Databases are essential assets of an enterprise and must be backed up. The backup frequency varies from weekly backup and daily backup to hourly backup based on the importance of data. Even if you back up data once every hour, you may lose one hour of data if a database failure occurs. This is unacceptable for industries such as gaming and finance. To reduce the amount of data lost during failures, DBS provides real-time backup and a recovery point objective (RPO) within seconds.

Database- and table-level restoration

Full data backup is a common practice. If data is accidentally deleted and you use the traditional method to restore data, you need to restore all data of the instance, find the table that is deleted, and then drop data that is not stored in the table. The traditional restoration method is time-consuming and greatly prolongs the data restoration time. To improve data restoration efficiency, DBS provides the table-level data restoration feature to restore data of a specific table.

- **High efficiency:** When you perform a table-level data restoration, DBS reads backup data only from the specified table for restoration. This greatly improves data restoration efficiency.
- **High flexibility:** You can combine this feature with the incremental backup feature to restore data to a any point in time.

Long-term archiving

To meet the security and compliance requirements, enterprises must retain some data for an extended period of time. Traditionally, the data is stored in disks or tapes. However, this method has various disadvantages, such as poor backup reliability due to possible hardware failures. If you fail to restore critical data, the consequences are unacceptable. As the amount of data massively increases, you face more challenges in managing data.

To achieve long-term archiving of databases, DBS supports the lifecycle management for backup data and different cost-effective storage methods. You can select a storage method based on the access frequency of backup sets to automatically dump and clean up backup sets.

Geo-disaster recovery

To create a complete disaster recovery solution to protect your database, you must implement geo-redundancy in addition to local data backups. The traditional solution is to duplicate the backup sets to other local disks or devices, which cannot defend against natural disasters such as earthquakes and typhoons. To implement geo-disaster recovery, you must build your backup data centers in other regions. This requires a large amount of upfront investment.

DBS provides a pay-as-you-go geo-redundancy feature. You can use this feature to back up ApsaraDB RDS databases and databases that are deployed in self-managed data centers, on Elastic Compute Service (ECS) instances, or on third-party cloud platforms to Object Storage Service (OSS). This helps implement geo-disaster recovery for your databases.

In addition to cold backup centers, you can also build hot backup centers by using Data Transmission Service (DTS). When a failure occurs in the business center, the business traffic is switched to the on-premises data center or the backup center.

Note

- Cold backup center: requires a low investment cost. Data is stored in OSS and can be recovered within hours.
- Hot backup center: requires a high investment cost. Data is stored in databases and can be recovered within minutes.

12. Middleware services

12.1. What is Message Queue for Apache RocketMQ?

Message Queue for Apache RocketMQ is a distributed messaging middleware that is developed by Alibaba Cloud based on Apache RocketMQ. This Alibaba Cloud service features low latency, high concurrency, high availability, and high reliability.

Message Queue for Apache RocketMQ is a core service in the enterprise-level Internet architecture. Using the highly available distributed cluster technology, this service offers a series of messaging services in the cloud, including message subscription and delivery, message tracing and query, scheduled or delayed messages, and resource statistics. Message Queue for Apache RocketMQ provides asynchronous decoupling and load shifting for distributed application systems. It also supports various features for Internet applications, including massive message accumulation, high throughput, and reliable message consumption retry. It is one of the core Alibaba Cloud services that are used to support the Double 11 Shopping Festival.

Message Queue for Apache RocketMQ supports access by using TCP. It also supports the Java, C++, and .NET programming languages. This facilitates quick access to Message Queue for Apache RocketMQ for applications that are developed in different programming languages.

12.1.1. Features

Message Queue for Apache RocketMQ supports access by using TCP and multiple programming languages, and offers multi-dimensional management tools. In addition, it provides a series of features for different scenarios.

Protocol access

- Support for TCP: Message Queue for Apache RocketMQ provides professional, reliable, and stable access from SDKs by using TCP. The SDKs for Java, C, C++, and .NET are supported.

Management tools

- Web console: The Message Queue for Apache RocketMQ console supports topic management, group management, message query, message tracing, and resource statistics.
- API: The Message Queue for Apache RocketMQ API allows you to integrate management tools into your console.
- mqadmin command set: Apsara Stack provides a rich set of management commands for you to manage the Message Queue for Apache RocketMQ service.

Message types

- Regular messages: They are messages without special features in Message Queue for Apache RocketMQ. Such messages are different from those with special features.
- Scheduled or delayed messages: Message Queue for Apache RocketMQ allows producers to specify the period of time to wait before a scheduled or delayed message is delivered. The maximum period of time is 40 days.
- Transactional messages: Message Queue for Apache RocketMQ provides a distributed transaction processing feature that is similar to X/Open XA to ensure transaction consistency.

- Ordered messages: Consumers can consume messages in the order in which messages are delivered.

Feature highlights

- Large messages: Message Queue for Apache RocketMQ supports a message that has a maximum size of 4 MB, including message properties.
- Message query: Message Queue for Apache RocketMQ allows you to query messages by message ID, by message key, and by topic.
- Message tracing: This feature records the complete trace of a message from its delivery by a producer to a RocketMQ broker and then to a consumer. This facilitates troubleshooting.
- Clustering consumption and broadcasting consumption: In clustering consumption mode, a message needs to be processed by only one of the consumers in a group. In broadcasting consumption mode, Message Queue for Apache RocketMQ pushes each message to all registered consumers in a group to ensure that each message is consumed by each consumer at least once.
- Consumer offset reset: You can reset the consumption progress by time to analyze message traces or discard accumulated messages.
- Dead-letter queues: Messages that cannot be consumed are stored in a special dead-letter queue for subsequent processing.
- Resource statistics: You can use this feature to collect statistics about message production and consumption. It allows you to view the total number of messages that a topic receives from producers or the transactions per second (TPS) for message production in a specific period of time. It also allows you to view the total number of messages that a topic sends to a group ID or the TPS for message consumption in a specific period of time.

Apsara Stack deployment

- Customization: Technical solutions as well as onsite technical support and training are provided.
- Flexible deployment: Message Queue for Apache RocketMQ can be independently deployed in Apsara Stack or deployed in a hybrid cloud architecture.
- O&M management: Apsara Stack supports O&M management tools such as the mqadmin command set and API operations. This facilitates console integration and unified O&M.

12.1.2. Benefits

This topic describes the advantages of Message Queue for Apache RocketMQ over other messaging middleware.

Professionalism

- Message Queue for Apache RocketMQ is a professional messaging middleware in the industry. This service ensures data integrity and features a sophisticated technical system.
- The open source version named Apache RocketMQ has won several awards both in China and abroad.
- More than 1,000 core applications within Alibaba Group use Message Queue for Apache RocketMQ. This service forwards hundreds of billions of messages per day. It delivers stable and reliable performance in real-life business scenarios such as the Double 11 Shopping Festival.

High reliability

- Each message is stored in multiple replicas on disks. No message is found lost after rigorous power-off tests.
- Message Queue for Apache RocketMQ supports massive message accumulation. A single topic can reliably support more than 10 billion messages even at heavy traffic.

- By default, messages are stored for three days. Message Queue for Apache RocketMQ allows you to reset consumer offsets for messages that are sent at a point in time within three days.

High performance

- Scaling out is supported.
- Message Queue for Apache RocketMQ supports a message that has a maximum size of 4 MB, including message properties.

Support for TCP

- Message Queue for Apache RocketMQ provides professional, reliable, and stable access from SDKs by using TCP.

Independent deployment

- Message Queue for Apache RocketMQ can be independently deployed in Apsara Stack as well as on physical machines. Only a few machines are needed to create a complete Message Queue for Apache RocketMQ service.
- Apsara Stack provides the mqadmin command set and management API operations. O&M engineers can conveniently monitor the system health in real time.
- Message Queue for Apache RocketMQ can be deployed in a hybrid cloud architecture and allows service access by using Express Connect circuits.

12.1.3. Scenarios

Message Queue for Apache RocketMQ is applicable to distributed transactions, real-time computing, Internet of Things (IoT) applications, and large-scale cache synchronization. This topic describes the application of Message Queue for Apache RocketMQ in different scenarios.

Distributed transactions

In traditional transaction processing mode, interactions among systems are coupled into one transaction. This extends the response time and subsequently affects system availability. The introduction of distributed transactional messages creates a transaction processing process between transaction systems and Message Queue for Apache RocketMQ. The process ensures data consistency between distributed systems. Downstream business systems such as shopping carts and points are isolated from each other and concurrently process transactions.

Real-time computing

Data that is continuously generated by a source flows to a computing engine in real time by using Message Queue for Apache RocketMQ. This achieves real-time computing. You can use the following computing engines: Spark, Storm, E-MapReduce, Application Real-Time Monitoring Service (ARMS), and Beam Runner.

IoT applications

IoT devices connect to the cloud by using Message Queue for MQTT for bidirectional communication and data transmission. Device data is connected to a computing engine by using Message Queue for Apache RocketMQ. Analysis data or source data is efficiently written to a data store such as a time series database (TSDB), HiStore, or MaxCompute in real time.

Large-scale cache synchronization

In business promotion activities such as the Double 11 Shopping Festival, each branch has a wide range of products, and the price of each product changes in real time. A huge number of concurrent access requests for the product database results in long response time on the web page of each branch. In centralized caching mode, the bandwidth becomes a bottleneck and blocks the access requests for product prices.

Message Queue for Apache RocketMQ can reduce the page response time by means of large-scale cache synchronization for different branches. This satisfies customers' access requirements for product prices.

13. Big data services

13.1. Apsara Big Data Manager

Apsara Big Data Manager (ABM) is an operations and maintenance (O&M) platform tailored for big data products. ABM supports monitoring, management, and O&M in multiple dimensions, such as business, services, clusters, and hosts. Customer O&M teams or on-site personnel can use ABM to perform O&M operations on big data products and manage the products in an efficient and centralized manner.

13.1.1. Features

Apsara Big Data Manager (ABM) provides the following functional modules for big data products: Monitor, O&M, and Management. The Monitor module helps you monitor the status of big data products in real time. The O&M and Management modules allow you to perform O&M and management operations by using a graphical user interface (GUI) based on multiple dimensions, such as business, services, clusters, and hosts.

Monitor

- **Dashboard:** allows you to view the overview of the key metrics of MaxCompute, DataWorks, Realtime Compute, and DataHub.
- **Repository:** allows you to view the resource usage of MaxCompute, DataWorks, and DataHub.
- **Health:** allows you to view the monitoring items of each big data service and the alerts of different levels of urgency and check the monitoring items and alerts again after the alerts are handled.
- **Reports:** allows you to view the inspection results of each big data service.

O&M

- **Business:** provides O&M features that are dedicated for big data products.
 - MaxCompute: allows you to manage and optimize projects and jobs.
 - DataHub: displays information about projects and topics in DataHub clusters.
 - Realtime Compute: displays information about projects, jobs, and queues in Realtime Compute clusters.
 - Elasticsearch: allows you to configure clusters and system settings.
 - MaxCompute and DataWorks: allows you to configure and manage geo-disaster recovery in the configuration center. You can configure protection groups, rehearsal plans, and failure plans.
- **Services:** displays O&M information about big data products. The information includes all service roles in a cluster and the resource usage trend of each service role. This tab also provides specific information about the following big data products:
 - MaxCompute: displays information about Control, Fuxi, Pangu, and Tunnel Service.
 - DataWorks: displays information about data warehouses and data integration.
 - DataHub: displays information about Control, Fuxi, and Pangu.
 - Realtime Compute: displays information about Realtime Compute, Yarn, and Hadoop Distributed File System (HDFS).

- **Clusters:** displays O&M information about clusters. The information includes the overview and the health status of clusters.
- **Hosts:** displays O&M information about hosts. The information includes the overview of hosts.

Management

- **Jobs:** allows you to run jobs to perform O&M operations on big data products. Jobs are classified into cron jobs and common jobs. The system can automatically run cron jobs at a specified time on a schedule. You can also manually run cron jobs. Common jobs can only be manually run.
- **Health Checks:** includes Monitoring Configuration and Patrol Configuration.
 - The Monitoring Configuration tab provides various built-in check items for each big data service. The system automatically checks service faults and sends alerts based on the scheduling items and check items. If a check item does not meet the requirements, the system reports an alert. You can also configure custom scheduling intervals and modify the status of scheduling items and check items. The information displayed on this tab helps you identify and handle service faults at the earliest opportunity.
 - The Patrol Configuration tab allows you to manage inspection items and inspection scenarios. The system manages inspection items at a specified time on a schedule. You can also manually manage inspection items. This helps prevent potential risks of major issues or service faults.
- **Processes:** allows you to process manual tasks, define processes, and view process instances.
- **Operation Logs:** provides the records of all O&M operations that are performed in the ABM console and the details of each operation. You can identify the causes of service faults based on the operation records.

13.1.2. Benefits

Apsara Big Data Manager (ABM) allows you to perform O&M operations on big data products and manage the products in a centralized manner. ABM helps enhance the stability of big data products, reduce O&M costs, and improve the efficiency of O&M. ABM provides benefits in various aspects, such as cluster health monitoring, resource usage analysis, and a graphical user interface (GUI) for O&M and management operations.

Cluster health monitoring

Allows you to monitor and configure the devices, resources, and services of the clusters of big data products, and collects and displays performance metrics of the clusters in real time.

Resource usage analysis

Displays the status of devices, resources, and services of clusters in real time, provides the historical data of the devices, resources, and services, and supports data aggregation and analysis to help you evaluate the health status of a cluster. If the evaluation result indicates potential risks in the cluster, related engineers can be immediately notified.

GUI for O&M and management operations

Provides a GUI on which you can view performance metrics and perform common O&M operations.

13.1.3. Scenarios

If you are using Apsara Stack and have deployed one or more big data services, you must use Apsara Big Data Manager (ABM) to perform O&M operations on these big data services.

Apsara Stack Enterprise and big data services

If you are using Apsara Stack Enterprise and have deployed one or more big data services, such as MaxCompute or DataWorks, you must use ABM to perform O&M operations on these big data services.

13.2. What is MaxCompute?

MaxCompute is a highly efficient, highly available, and low-cost **EB-level** computing service for big data. It is independently developed by Alibaba Cloud. This service is used within Alibaba Group to process exabytes of data each day. MaxCompute is a distributed system designed for big data processing. As one of the core services in the Alibaba Cloud computing solution, MaxCompute is used to store and compute structured data.

MaxCompute is designed to support multiple tenants and provide features such as data security and horizontal scaling. It provides a centralized graphical user interface (GUI) and centralized APIs for various data processing tasks of different users based on an abstract job processing framework.

MaxCompute is used to store and compute large amounts of structured data at a time. It provides various data warehousing solutions as well as big data analytics and modeling services. MaxCompute aims to provide easy analysis and processing of large amounts of data. You can analyze big data without a deep knowledge of distributed computing.

MaxCompute has the following features:

- Uses a distributed architecture that can be horizontally scaled based on your business requirements.
- Provides automatic storage and fault tolerance mechanisms to ensure high data reliability.
- Allows all computing tasks to run in sandboxes to ensure high data security.
- Uses RESTful APIs to provide services.
- Supports both uploads and downloads of high-concurrency, high-throughput data.
- Supports two service models: the offline computing model and the machine learning model.
- Supports data processing methods based on programming models such as SQL, MapReduce, Graph, and MPI.
- Supports multiple tenants, which allows multiple users to collaborate on data analysis.
- Manages user permissions based on access control lists (ACLs) and policies, which allows you to flexibly configure access control policies to prevent unauthorized data access.
- Supports Spark on MaxCompute, the enhanced Spark application.
- Supports Elasticsearch on MaxCompute, the enhanced Elasticsearch application.
- Supports the access to and processing of unstructured data.
- Supports the deployment of multiple clusters in a single region.
- Supports multi-region deployment.
- Stores audit logs and automatically dumps them to a specific server directory for long-term storage and management.

13.2.1. Features

MaxCompute is a data processing platform developed by Alibaba Group to process large amounts of data. MaxCompute provides channels for data uploads and downloads, a range of computing and analysis features including MaxCompute SQL and MaxCompute MapReduce, and comprehensive security solutions.

Data tunnel

- **Tunnel:** provides highly-concurrent offline data upload and download services. MaxCompute Tunnel enables you to upload or download large amounts of data to or from MaxCompute. MaxCompute Tunnel provides a Java API.
- **DataHub:** provides real-time data uploads and downloads. Data uploaded by using DataHub is available immediately, whereas data uploaded by using MaxCompute Tunnel is not.

Computing and analysis tasks

- **SQL:** MaxCompute stores data in tables and allows data queries by using SQL statements. MaxCompute can be used as traditional database software, and can process terabytes to petabytes of data. MaxCompute SQL does not support transactions, indexing, or operations such as UPDATE and DELETE. The SQL syntax that is used in MaxCompute is different from the SQL syntax that is used in Oracle and MySQL. SQL statements from other database engines cannot be seamlessly migrated to MaxCompute. MaxCompute SQL returns query results in minutes or seconds, not in milliseconds. MaxCompute SQL is easy to learn. You can get started with MaxCompute SQL based on your prior experience of database operations, without a deep knowledge of distributed computing.
- **MapReduce:** First proposed by Google, MapReduce is a distributed data processing model that has gained extensive attention and has been used in a wide range of business scenarios. You must have a basic knowledge of distributed computing and related programming experience before you use MapReduce. MapReduce provides a Java API.
- **Graph:** an iterative graph computing framework provided by MaxCompute. Graph computing jobs use graphs to build models. A graph is a collection of vertices and edges that have values. MaxCompute Graph performs iterations to edit graphs, enable graphs to evolve, and obtain analysis results.
- **Unstructured data access and processing in integrated computing scenarios:** MaxCompute SQL cannot directly process external data, such as unstructured data from Object Storage Service (OSS). Data must be imported to MaxCompute tables by using relevant tools before computation. The MaxCompute team introduces the unstructured data processing framework to the MaxCompute system architecture to handle this issue. MaxCompute allows you to create external tables to process data from the following data sources:
 - **Internal data sources:** OSS, Tablestore, AnalyticDB for MySQL, ApsaraDB RDS, Apsara File Storage for HDFS, and Taobao Distributed Data Layer (TDDL)
 - **External data sources:** open source Hadoop Distributed File System (HDFS), ApsaraDB for MongoDB, and ApsaraDB for HBase
- **Unstructured data access and processing in MaxCompute:** MaxCompute reads and writes volumes to store and process unstructured data, which otherwise must be stored in an external storage system.

Spark on MaxCompute

Spark on MaxCompute is a big data analysis engine that is developed by Alibaba Cloud and provides users with big data processing capabilities.

SDK

SDK is a toolkit that MaxCompute provides for developers.

Security solution

MaxCompute provides powerful security services to ensure user data security.

13.2.2. Benefits

Excellent big data cloud service and real data sharing platform in China

- MaxCompute can be used for data warehousing, mining, analysis, and sharing.
- Alibaba Group uses this centralized data processing platform in several of its own services, such as Aliloan, Data Cube, DMP (Alimama), and Yu'e Bao.

Support for a large number of clusters, users, and concurrent jobs

- A single cluster can contain more than 10,000 servers and maintain 80% linear scalability.
- A single MaxCompute system supports more than 1 million servers in multiple clusters without limits. However, linear scalability is slightly affected. It also supports multi-data-center deployment in a zone.
- A single MaxCompute system supports more than 10,000 users, more than 1,000 projects, and more than 100 departments of multiple tenants.
- A single MaxCompute system supports more than 1 million jobs (daily submitted jobs on average) and more than 20,000 concurrent jobs.

Big data computing at your fingertips

You do not need to worry about the storage difficulties and prolonged computing processes caused by the increase of the data volume. MaxCompute automatically expands the storage and computing capabilities of clusters based on the data volume. This allows you to focus on data analysis and mining to maximize your data value.

Out-of-the-box service

You do not need to worry about the creation, configuration, and O&M of clusters. Only a few simple steps are required to upload data, analyze data, and obtain analysis results in MaxCompute.

Secure and reliable data storage

MaxCompute uses multi-level data storage and access control mechanisms to protect user data against loss, leaks, and interception. These mechanisms include multi-copy technology, read and write request authentication, and application and system sandboxes.

Reliable management nodes

MaxCompute uses the multi-node cluster architecture. The management nodes of each component feature high availability. The faults that occur on O&M management nodes do not interrupt your services.

Powerful fault tolerance

MaxCompute supports automatic fault tolerance for the failures of hard disks on servers in a cluster and supports hot swapping of hard disks. In the event of a hard disk failure, services can be restored within 2 minutes.

Comprehensive storage space management

MaxCompute allows you to query information about both the storage capacity and usage of distributed file systems. It enables you to manage data lifecycles. MaxCompute also allows you to store data in different locations based on the data value or tag. For example, you can write temporary files to SSDs to accelerate I/O operations. This allows you to use cluster data more efficiently. MaxCompute also supports the self-optimizing Zstandard compression algorithm that provides the optimal compression ratio.

Comprehensive data backup

- MaxCompute allows you to perform full or incremental data backup and restore data from storage media.

- MaxCompute allows you to back up data for clusters in different data centers. This meets the requirements of mutual data backups among multiple data centers. You can use Apsara Big Data Manager (ABM) to manage the backup process in a visualized manner.
- MaxCompute allows you to back up and restore the metadata, files, and tables of key components.

Secure and reliable access control

- MaxCompute allows you to manage data access permissions. The permissions include logon permissions, permissions to create tables, read and write permissions, and whitelist-related permissions.
- MaxCompute allows you to manage administrative permissions, including administrator classification, in the Apsara Uni-manager Management Console.
- MaxCompute allows you to manage user permissions in the Apsara Uni-manager Management Console in a centralized manner. You can view and manage the permission management features of all components in the system. You can also keep permission management details from common users and simplify permission management for administrators. This improves the usability and user experience of permission management.

Multi-tenancy for multi-user collaboration

MaxCompute allows you to configure data access policies. This way, you can enable multiple data analysts in an organization to collaborate and make data accessible to users who are granted the required permissions. This ensures data security and maximizes productivity.

- **Isolation:** You can submit the tasks of multiple tenants (projects) to different queues for concurrent running. Resources are isolated among tenants.
- **Permission:** You can manage different tenants in a centralized manner and dynamically configure, manage, and isolate tenant resources. You can also collect statistics on the usage of tenant resources and manage multi-level tenants.
- **Scheduling:** MaxCompute supports multi-tenant scheduling for multiple clusters and resource pools.

Multi-region deployment

- You can specify compute clusters to efficiently use computing resources.
- Data exchanges between clusters are completed within MaxCompute, and data replication and synchronization between clusters are managed based on the configured policies. Therefore, cross-region data processing is no longer involved, which significantly reduces the waiting time for data processing.

Multi-device support

You can use CPUs, hard disks, memory, and network interface controllers with different specifications in a single-component cluster to ensure maximum compatibility with existing devices. This applies only when cluster performance is not affected.

13.2.3. Scenarios

MaxCompute is suitable for storage and computing in scenarios in which more than 100 GB of data is involved. MaxCompute can process up to exabytes of data and is widely used in Alibaba Group. MaxCompute is suitable for various big data processing scenarios, such as data warehousing and business intelligence (BI) analysis for large Internet enterprises, website log analysis, e-commerce transaction analysis, and exploration of user characteristics and interests.

Cost-effectiveness and quick data migration to the cloud

Scenario: A cloud platform for Internet big data application services is built to meet the business requirements of an information service provider that focuses on the new energy industry.

Results: The entire business system is migrated to the cloud within three months. The cloud platform reduces the data processing time to less than one third of the self-managed big data platform. Multiple security mechanisms help ensure the security of the cloud data.

Customer benefits:

- **More focus on core business:** The entire business system is migrated to the cloud within three months. This allows the enterprise to use various cloud resources to improve the business.
- **Low investment and O&M costs:** Compared with a self-managed big data platform, the cloud platform significantly helps reduce the costs of infrastructure construction, O&M personnel, and R&D.
- **High security and stability:** The comprehensive service capabilities and stable performance of Alibaba Cloud ensure data security on the cloud.

Fine-grained business operation for millions of users based on large amounts of data

Scenario: A big data platform is built to meet the business requirements of a community-oriented vertical e-commerce app that focuses on the manicure industry. The big data platform provides business monitoring, business analysis, fine-grained operation, and recommendation for the app.

Results: MaxCompute helps achieve fine-grained operation for millions of users. The cloud platform helps the e-commerce app be more agile, intelligent, and insightful and quickly respond to the data and analytics requirements of new business.

Customer benefits:

- **Improved business insights:** MaxCompute helps achieve fine-grained operation for millions of users.
- **Data-driven business:** The platform improves the business data analysis capability and monitors business data in an effective manner to improve business efficiency.
- **Quick response to business requirements:** The MaxCompute ecosystem can quickly respond to the business requirements for data analysis.

Precision marketing based on big data

Scenario: A core big data precision marketing platform is built for an Internet enterprise that focuses on precision marketing and advertising technologies and services.

Results: A core big data precision marketing platform is built. All log data is stored in MaxCompute. Batch scheduling and analysis are performed by using DataWorks.

Customer benefits:

- **Cost-effective analysis of large amounts of data:** Statistical analysis of large amounts of data can reduce costs by half while equivalent business requirements are met. This helps reduce costs in an effective manner and helps startup enterprises grow rapidly.
- **Real-time data query and analysis:** MaxCompute provides technical advantages for enterprises, eliminates the technical bottleneck in processing and analysis of large amounts of data, and queries and analyzes data in real time. MaxCompute collects, analyzes, and stores more than 2 billion data records of visitor activities per day. MaxCompute can respond to data queries from a table that contains hundreds of millions of log entries within milliseconds based on business requirements.
- **Easy use of Machine learning platform:** The performance of algorithm models directly affects the revenue of the customer. The machine learning platform of MaxCompute is easy to use and is a good choice to meet the business requirements of the customer.

13.3. DataWorks

DataWorks is an end-to-end big data development and governance platform based on compute engines such as MaxCompute and E-MapReduce. DataWorks integrates all processes from data collection to data display and from data analysis to application running. DataWorks helps you quickly complete the entire research and development (R&D) process, including data integration, data development, data governance, data service provisioning, data quality control, and data security assurance.

13.3.1. Features

DataWorks provides an end-to-end solution from data collection to data display and from data analysis to application driving. DataWorks supports batch processing, batch analysis, and mining of large amounts of data. DataWorks integrates various services, such as Data Integration, DataStudio, Operation Center, DataAnalysis, Data Asset Management, Data Quality, Data Protection, and DataService Studio. DataWorks can also work with Machine Learning Platform for AI.

Data Integration

Data Integration is a stable, efficient, and scalable data synchronization service. This service is used to efficiently and stably migrate and synchronize data between heterogeneous data sources in complex network environments.

Data Integration can read and monitor the data of your database. This service provides an overview of all data sources, such as relational databases, NoSQL databases, big data databases, and FTP servers, and allows you to use the following method to synchronize data: batch synchronization, full synchronization, and incremental synchronization. Data can be synchronized at intervals of minutes, hours, days, weeks, or months.

Supported data sources

- **Metadata:** Data Integration can collect metadata from more than 20 types of common data sources, such as MySQL databases, SQL Server databases, Oracle databases, and MaxCompute projects. Data Integration provides a comprehensive overview of all data assets from the collected metadata to help manage data assets and synchronize core data.
- **Relational databases:** Data Integration allows you to read data from and write data to relational databases, such as MySQL, SQL Server, Oracle, PostgreSQL, Db2, and PolarDB-X 1.0.
- **NoSQL databases:** Data Integration allows you to read data from and write data to NoSQL databases, such as HBase, MongoDB, and Tablestore.
- **Massively parallel processing (MPP) databases:** Data Integration allows you to read data from and write data to MPP databases, such as HybridDB for MySQL and HybridDB for PostgreSQL.
- **Big data databases:** Data Integration allows you to read data from and write data to MaxCompute projects and Hadoop Distributed File System (HDFS). Data Integration also allows you to write data to AnalyticDB databases.
- **Unstructured data sources:** Data Integration allows you to read data from and write data to Object Storage Service (OSS) objects and FTP servers.

Real-time synchronization

Data Integration allows you to synchronize data from a data source, such as MySQL, Oracle, Db2, SQL Server, OceanBase, DataHub, LogHub, or Kafka, to MaxCompute, Hologres, Kafka, or DataHub in real time.

Inbound data control

Data Integration supports conversion between different data types and accurately identifies, filters, collects, and displays dirty data to facilitate inbound data control. In addition, Data Integration provides statistics such as data volume, data throughput, and job duration and detects dirty data in each job.

Fast transmission

Data Integration uses the network interface controller (NIC) on each server and adopts a distributed architecture to transmit gigabytes or terabytes of data within a short period of time.

Accurate throttling

Data Integration implements accurate throttling on channels, record streams, and byte streams. Data Integration also supports fault tolerance and allows you to rerun specific or all threads, processes, and jobs.

Synchronization plug-ins

Data Integration provides synchronization plug-ins that can be used to connect to the servers of data sources and collect data.

Cross-network transmission

Data Integration supports data transmission in complex network environments. For example, DataWorks enables data transmission across local private networks or VPCs.

Data Modeling

DataWorks provides the Datablau feature based on the Datablau Data Modeler (DDM) client. You can use Datablau to design and manage data models and add and implement data standards for data development in DataWorks. This allows you to obtain more value from data and efficiently manage the lifecycle of data.

- Visualized model design: Data Modeling provides professional local clients and lightweight online clients that you can use to build visualized models in different scenarios.
- Collaborative design: Data Modeling allows multiple users to collaborate to build a model.
- Customization of data standards: Data Modeling allows you to define data standards, code specifications, and naming conventions.
- Intelligent reference of data standards: Data Modeling allows you to set data standards before you build a data model. This way, you can reference the pre-set standards when you build data models.
- Forward and reverse modeling based on DDL statements: Data Modeling allows you to execute DDL statements to publish a model to a compute engine instance. Data Modeling also allows you to export a model from a compute engine instance, edit the model, and then republish the model to the compute engine instance. This way, you do not need to perform complex operations in traditional modeling, such as manual import and export of data models.
- Model library: Data Modeling allows you to import data models to or export data models from a model library for centralized management.
- Model monitoring: Data Modeling helps you monitor the data models that are published to a compute engine instance. This way, you can easily identify the inconsistency between the schema of a table and the schema of a physical model.
- Integration with DataStudio: Data Modeling allows you to publish a model to a compute engine instance that is associated with a DataWorks workspace for subsequent data development and analysis. This way, you can design and publish models in a standard manner.

DataStudio

DataStudio is an end-to-end integrated development environment (IDE) that enables you to develop the extract, transform, and load (ETL) process and data mining algorithms. DataStudio also allows you to build data warehouses in DataWorks.

After underlying data is aggregated, the data is disordered and cannot be directly used by upper-layer algorithms or data intelligence applications. In this case, you must process the aggregated data. You can build data centers in the data resource platform to process data.

Workflows

DataStudio allows you to organize different types of nodes in a workflow for subsequent data development. DataStudio provides you with a directed acyclic graph (DAG) for nodes in each workflow. You can use related tools to develop and manage the nodes in the DAG in a convenient and intelligent manner.

A workflow can contain the following types of nodes: ODPS SQL, ODPS MR, Shell, machine learning, data synchronization, PyODPS, SQL component, and zero load nodes. You can configure dependencies between nodes within a workflow or across workflows. You can also schedule a workflow or specific nodes to run at specific points in time based on your business requirements.

- **Nodes:** The following types of nodes are supported: ODPS SQL, ODPS MR, Shell, machine learning, data synchronization, PyODPS, SQL component, and zero load nodes. You can configure dependencies between nodes within a workflow or across workflows. You can schedule a node based on your business requirements.
- **Node properties:** You can double-click the name of a node to configure the node. For example, you can write SQL statements for an SQL node or configure data synchronization rules for a data synchronization node. You can also click the tabs in the right-side navigation pane to view the version information or modify the configurations of a node, such as the scheduling properties, lineage, and schema.
- **Historical versions:** You can view the versions of nodes, such as ODPS SQL, ODPS MR, and Shell nodes. You can also roll back a node to an earlier version.
- **Node deployment:** In workspaces that are in standard mode, you can deploy nodes that passed tests to the production environment.

Solutions

DataStudio allows you to create solutions and workflows for data development. You can add one or more workflows to one solution. This way, you can develop data across workflows. You can also add a workflow to multiple solutions. This way, you can develop data based on solutions.

Code editor

DataStudio provides a code editor in which you can perform SQL and MapReduce (MR) programming, upload files as resources, register user-defined functions (UDFs), and write Shell scripts.

- **SQL programming:** The web-based code editor allows you to write SQL statements and supports various features, such as automatic SQL statement completion, code formatting, code highlighting, and code debugging.
- **MapReduce programming:** When you configure an ODPS MR node in the code editor, you can upload a Java Archive (JAR) file that contains MapReduce code as a JAR resource and then reference the file in the node.
- **Resource files:** You can upload files, such as JAR files, Python files, Shell scripts with custom parameters, XML configuration files, and TXT configuration files, as resources. The code editor can identify the types of files based on the file name extensions, such as .zip, .tgz, .tar.gz, .tar, and .jar.

- Registration of UDFs: The code editor allows you to register Java or Python UDFs. Before you register UDFs, you must upload JAR or Python files as resources. You can use the UDFs to develop data.
- Shell script programming: The code editor provides an online environment in which you can write and debug a Shell script.

Code repository and team collaboration

DataStudio allows multiple users to simultaneously work on the same workflow. This improves development efficiency.

DataStudio provides a lock mechanism that allows you to lock workflows or nodes. If you lock a workflow or node, other users cannot edit the workflow or node. If you want to edit a workflow or node that is locked by another user, you can try to obtain a lock on the workflow or node. After you initiate the request, the system notifies the user who locked the workflow or node.

DataWorks records the version of a node or a workflow each time you commit the node or the workflow. You can compare two versions of a node or a workflow.

Operation Center

Operation Center is an end-to-end service that provides O&M capabilities for data developers and O&M personnel. Operation Center allows you to specify the priorities of nodes and control and monitor the execution of nodes.

To handle the volume, diversity, and complexity of data that is used in DataWorks and meet the requirements of various data processing tasks, a scheduling system that supports high concurrency, multiple cycles, and various data processing procedures is required.

Operation Center allows you to monitor the status of nodes, view alerts of failed nodes or instances, and view the daily O&M statistics.

Overview page

On the Workbench Overview tab of the Overview page, you can view the statistics about nodes and instances in the following sections: Node Completion, Running Status Distribution, Runtime Ranking, Number of Auto Triggered Nodes, Error Ranking in Recent Month, Node Types, and Scheduling Resource Allocation.

Node O&M

Operation Center allows you to view the DAG of nodes. DAGs help simplify the maintenance and management of online nodes.

- You can monitor the status of nodes by using alert rules. You can rerun, stop, or suspend nodes, and set the status of nodes to successful.
- You can view nodes in a node list or in the DAG.
- You can view the status of auto triggered nodes, test nodes, and manually triggered nodes.
- You can view the run logs, code, and property settings of nodes.

Intelligent monitoring

The intelligent monitoring feature can be used to monitor and analyze nodes. The intelligent monitoring system monitors the status of nodes and sends alert notifications based on the intervals, notification methods, and notification recipients that are specified in alert rules. When the conditions in an alert rule are met, the intelligent monitoring system automatically sends an alert notification to the specified recipient by using the specified notification method.

The intelligent monitoring system provides comprehensive monitoring and alerting logic. You need to only specify the names of important nodes in your business. Then, the intelligent monitoring system automatically monitors the entire process of your nodes and generates standard alert rules for the nodes. You can also customize alert rules based on your business requirements.

Engine O&M

DataWorks allows you to view the resource usage of compute engines on the Engine Maintenance page. For example, you can view the usage details of computing and storage resources and the usage details of the resources that are occupied by jobs. You can use the engine O&M feature to view the details of a node and identify and remove the nodes that fail to run. This prevents failed nodes from affecting the execution of the descendant nodes.

DataAnalysis

DataAnalysis provides the ad hoc query and table management capabilities. DataAnalysis accelerates the analysis process by using the data collection tools of MaxCompute in near real-time mode.

Data Map

You can use Data Map to manage the metadata and data assets of your business. You can also use Data Map to globally search for data, view the details of metadata, preview data, view data lineage, and manage data categories. Data Map helps you search, understand, and use data.

Security Center

Security Center allows you to easily manage permissions and submit and handle requests on a visual interface. You can audit and manage the permissions.

- Self-service permission request: Users can select the tables whose permissions they want to request and submit permission requests online. This online request mode is more efficient than the original mode in which users need to contact administrators offline.
- Permission management: Administrators can view the users who have permissions on database tables. If the permissions of the users do not meet specific requirements, administrators can revoke the permissions from the users. Users can also manually remove the permissions that they no longer require.
- Permission request handling: Administrators cannot directly grant permissions to users and can only handle the permission requests that are submitted by users. Security Center uses a visual and process-oriented mechanism to manage permissions. This way, users can monitor the status of permission requests.

DataService Studio

DataService Studio allows you to create APIs based on data tables. You can register existing APIs in DataService Studio. This way, you can manage and publish APIs in a centralized manner. DataService Studio is integrated with API Gateway. This allows you to publish APIs to API Gateway by performing simple operations and provides a secure, stable, cost-effective, and easy-to-use data sharing service. DataService Studio is built based on a serverless architecture. You can focus on the query logic of APIs without the need to consider the infrastructure, such as the runtime environment. DataService Studio prepares computing resources for you, supports elastic scaling, and requires zero O&M costs.

DataService Studio can be used by public service sectors as a secure, flexible, and reliable platform for data sharing across departments and networks. DataService Studio also enables the public service sectors to share data with the public.

API generation

DataService Studio allows you to create APIs on the codeless user interface (UI) based on tables in relational databases, NoSQL databases such as Tablestore, and analytical databases such as AnalyticDB. You can quickly create APIs without the need to write code and call the API in the codeless UI. DataService Studio also allows you to create an API in the code editor. You can write SQL statements to customize the query logic of the API. In the code editor, you can specify multi-table join queries, complex query criteria, and aggregate functions.

API registration

DataService Studio allows you to register your RESTful APIs. This way, you can manage your RESTful APIs and the APIs that you create in DataService Studio based on tables in a centralized manner. You can register the following types of RESTful APIs: GET, POST, PUT, and DELETE. The following data formats are supported: tables, JSON, and XML.

API Gateway

API Gateway provides API lifecycle management services, including API publishing, management, O&M, and sale. API Gateway provides a simple, fast, cost-effective, and low-risk service that you can use to aggregate microservices, separate the frontend from the backend, integrate systems. API Gateway allows you to share features and data with partners and developers. DataService Studio allows you to publish APIs to API Gateway in a convenient manner. You can publish the APIs that you create based on data tables and the APIs that you register in DataService Studio to API Gateway to implement centralized management. API Gateway can be used for authorization, authentication, throttling, and billing.

Migration Assistant

Migration Assistant allows you to migrate data objects across different DataWorks editions, Alibaba Cloud accounts, clusters, regions, and workspaces. Migration Assistant also allows you to perform custom migration on your data objects. The data objects include auto triggered nodes, manually triggered nodes, resources, functions, data sources, script templates, ad hoc queries, table metadata (including the DDL statements that are used to create tables), and objects in DataService Studio.

Platform management

DataWorks provides the platform management capability to help administrators manage DataWorks users. Workspaces are organizational units for code, member, role, and permission management in DataWorks. Workspaces are isolated from each other. You can view and modify code of a node in a workspace only if you are a member of the workspace and are granted the required permissions.

Organization management

The Organizations page displays the account, AccessKey ID, and AccessKey secret of the owner of the current organization. You can manage all members of the organization on this page.

Workspace management

Workspaces are displayed as a list. On the Workspaces page in the DataWorks console, you can create, configure, enable, and disable workspaces.

Member management

The User Management page displays members in a list. You can view the information about each member of a workspace, such as the member name, account, and roles, on this page.

Permission management

You can manage the permissions of roles and users in a centralized manner.

Data Asset Management

Data Asset Management is a tenant-level feature. To use this feature, you must obtain the required permissions.

This feature allows you to manage data assets in your business system and DataWorks, such as tables and APIs. Before you use this feature, you must use Data Integration to synchronize data and then use DataStudio to process the data. This way, you can manage data in a centralized manner. You can also have a command of the core data assets and formulate standards for the management of the data assets.

Data Security Guard

Data Security Guard is a data security management platform that you can use to identify data assets, detect sensitive data, classify data based on categories and sensitivity levels, mask data, monitor data access behavior, report alerts, and audit risks.

Rule configurations

Data Security Guard allows you to configure risk identification rules as a data security administrator to identify sensitive data.

Data discovery

On the next day after you configure and enable sensitive data identification rules as a data security administrator, you can view the overall data distribution, hierarchical data distribution, and field details.

Data access

Data activities include data access and data export activities. On the next day after you configure sensitive data identification rules as a data security administrator, you can view the data usage and the data that is exported from MaxCompute.

Data masking

On the Data Masking tab of the Data Masking page, you can create, modify, delete, or test a data masking rule. You can configure a data masking method for each data masking rule. If some data does not need to be masked, you can configure a masking rule whitelist.

Sensitivity level management

On the Data classification and level page, you can configure the sensitivity levels if the existing configuration cannot meet your business requirements.

Manual data correction

Users can manually modify incorrect identification results. For example, users can delete data that is incorrectly identified, change the sensitive field type of identified data, and process multiple data records at a time.

Data risks

Data activities are audited manually or based on the risk identification rules and AI-based identification rules. The Data Risks page displays data activities that are considered risky. You can add comments to audit results based on your business requirements.

Risk identification rule management

You can configure risk identification rules or enable AI-based identification rules to identify risks when users access your data. The Data Risks page displays the identified data risks and allows users to tag the data risks as secure or risky. In the View Details dialog box that appears after you click View Details in the Actions column of a data access record on the Data Activities page, you can view the risk identification rules that correspond to the identified data risks.

Data audit

You can view the risk handling results and distribution of your data risks from multiple dimensions.

Data Security Chain

Data Security Chain enables data transfer without transferring data control ownership to another user. This ensures data security during data transfer and data sharing, and facilitates data sharing.

- Data Security Chain provides a secure environment for data transfer and ensures that the transferred data is encrypted based on the business requirements of the data owner.
- Data Security Chain provides the cryptographic access control feature that allows encrypted data to be accessed only by users who are granted the required permissions. This prevents data from being sold.
- The cryptographic access control feature also allows the data owner to revoke related permissions from a user. After the permissions are revoked from a user, the user cannot use the transferred encrypted data.
- Data Security Chain provides a component that can be used to decrypt and query ciphertext. Users can use this component only to process ciphertext. The decryption or query results that are generated by the component are masked based on the data masking rules that are specified by the data owner.
- The component can record all data usage logs. This allows the data owner to audit the operations that are performed on the data.
- All data sharing, request, processing, and usage records are stored on blockchains. This ensures that data is traceable and cannot be tampered with.

Data Security Chain provides the following capabilities: directory management after data sharing, request progress management, permission management, protection policy management, and data usage audit. This ensures data security during data sharing.

13.3.2. Benefits

This topic describes the benefits of DataWorks.

Powerful computing capabilities

DataWorks can be integrated with compute engines to process large amounts of data.

- DataWorks supports join operations for trillions of data records, millions of concurrent nodes, and I/O throughput of petabytes per day.
- The offline scheduling system can run millions of concurrent nodes. You can also configure rules and alerts to monitor the execution of nodes in real time.
- DataWorks provides efficient and easy-to-use SQL and MapReduce engines, and supports most standard SQL syntax.
- DataWorks protects your data from loss, breach, or theft by using multi-layer data storage and access security mechanisms, including triplicate backups, read/write request authentication, application sandboxes, and system sandboxes.

End-to-end platform

DataWorks provides a graphical user interface (GUI) that allows multiple users to collaborate on a workspace.

- DataWorks integrates all processes from data integration, processing, management, and monitoring to output.
- DataWorks allows you to create and edit workflows on the GUI.

- DataWorks provides a collaborative development environment. Users can be added to a workspace and assigned different roles to perform operations, such as node development, online scheduling, O&M, and data permission management.

Integration of various heterogeneous data sources

DataWorks allows you to schedule batch synchronization nodes by minute, hour, day, week, or month to synchronize data between data sources. More than 400 pairs of heterogeneous data sources are supported.

Web-based software

DataWorks is an out-of-the-box service. You can use DataWorks over an internal network or the Internet.

Multitenancy

DataWorks uses multitenancy to isolate data among tenants. Each tenant separately manages permissions, data, resources, and members.

Intelligent monitoring and alerting

DataWorks allows you to configure monitoring baselines to control the entire process of all nodes and monitor the status of each node.

Easy-to-use SQL editor

The SQL editor supports automatic code and metadata completion, code formatting and folding, and pre-compilation. The SQL editor also offers two editor themes. These features ensure good user experience.

Comprehensive data quality monitoring

DataWorks supports quality check, notification, and management for heterogeneous data sources, batch data, and real-time data.

Convenient API development and management

The API Gateway service and DataService Studio allow you to easily develop and publish APIs for data sharing.

Secure data sharing

DataWorks allows you to mask sensitive data before sharing the data to other tenants. This ensures the security of your big data assets and maximizes the value of the assets.

13.3.3. Scenarios

DataWorks is widely used in cloud-based data warehouse building and data-driven operations scenarios.

Cloud-based data warehouse building

Enterprises can use DataWorks in Apsara Stack to build large data warehouses.

- **Massive data storage:** allows you to build a data warehouse that stores petabytes or exabytes of data and supports linear expansion of storage capacity.
- **Data integration:** supports data synchronization and integration across heterogeneous data sources to eliminate data silos.

- **Data development:** provides big data development capabilities based on MaxCompute, supports programming frameworks such as SQL and MapReduce, and provides visualized workflows for specific business scenarios.
- **Data management:** provides a data resource management view and a procedure for processing data permission requests based on a unified metadata service.
- **Batch scheduling:** supports periodic scheduling of nodes at different intervals, scheduling of millions of concurrent threads per day, and real-time node monitoring and alerting.

Data-driven operations

- **Innovative business:** supports data mining, data modeling, and real-time decision making to help you use big data analytics results in your business system.
- **Small and medium-sized enterprises:** supports efficient processing and analysis of data to help enterprises make business decisions.

13.4. DataHub

DataHub is a platform designed to process streaming data. You can publish and subscribe to streaming data in DataHub and distribute the data to other platforms. DataHub allows you to analyze streaming data and build applications based on the streaming data.

DataHub is seamlessly integrated with Realtime Compute, enabling you to use SQL to analyze streaming data.

13.4.1. Features

DataHub collects, stores, and processes streaming data from mobile devices, applications, website services, and sensors. You can use your own applications or Apsara Stack Realtime Compute to process streaming data in DataHub, such as real-time website access logs, application logs, and events. The processed data is updated in real time and delivered to you in the form of alerts and statistics (visualized as graphs or tables).

Data queue

DataHub preserves the order of data within each shard. The performance of a single topic is horizontally scalable by adjusting the number of shards.

Checkpoint-based data restoration

DataHub provides the application checkpointing feature. If applications fail, you can restore data from any one of the saved checkpoints.

Data synchronization

Data in DataHub is automatically synchronized to other Alibaba Cloud services. You can create a DataConnector to synchronize data in DataHub to other Alibaba Cloud services in real time or near real time. These Alibaba Cloud services include MaxCompute, Object Storage Service (OSS), Elasticsearch, ApsaraDB RDS for MySQL, AnalyticDB, and Tablestore.

After you configure DataConnector, data you write to DataHub becomes available to other Alibaba Cloud services. This allows you use the data in other Alibaba Cloud Services. The synchronization process applies at-least-once semantics to ensure data is not lost when errors occur during synchronization. However, this may result in duplicate records.

Scalability

DataHub can scale flexibly to increase or decrease the number of shards in a cluster. This allows you to handle traffic surges or save resource costs based on actual business conditions.

For example, if the topic has insufficient throughput to cope with the surge in service load during the Double 11 Shopping Festival, you can use the SplitShard operation to increase the number of shards to a maximum of 256 shards in response to the traffic spike. This expands the throughput capacity up to 256 MB/s. When traffic returns to normal after the Double 11 Shopping Festival and you no longer require high throughput capacity, you can use the MergeShard operation to merge every two shards into one.

13.4.2. Benefits

Built on the Apsara system of Alibaba Cloud, DataHub features elastic scalability, high availability, high throughput, and low latency.

High throughput

You can write terabytes (TB) of data into a topic and up to 80 million records into a shard every day.

Low latency

DataHub makes it easy to collect and process various streaming data in real time so you can react quickly to new information.

High usability

- DataHub provides a variety of SDKs in different programming language, including C++, Java, Python, Ruby, and Go.
- In addition to SDKs, DataHub provides RESTful APIs so that you can manage DataHub by using existing protocols.
- You can use collection tools such as Fluentd, Logstash, and Oracle GoldenGate to write streaming data to DataHub.
- DataHub supports structured and unstructured data. You can write unstructured data to DataHub, or create a schema for the data before it is written to the system.

High availability

- No service interruption occurs during auto-scaling.
- DataHub automatically backs up multiple copies of data to implement data redundancy.

Dynamic scaling

The throughput of each topic can be increased or decreased dynamically, with the maximum throughput of a topic reaching 256,000 records per second.

High security

- DataHub provides enterprise-level security mechanisms and tenant-level isolation.
- It supports multiple authentication and authorization methods. For example, you can configure an allowlist or grant permissions by using Resource Access Management (RAM).

13.4.3. Scenarios

As a streaming data processing platform, DataHub can be used with various Alibaba Cloud products to provide one-stop data processing services.

Data uploading

DataHub can be connected to other Alibaba Cloud services. This allows you to use data uploaded to DataHub in other Alibaba Cloud services without the need to repeatedly upload the data.

Data collection

DataHub provides a variety of data collection tools for you to write your data into DataHub. DataHub supports log collection from Logstash and Fluentd, and binary log collection from Data Transmission Service (DTS) and Oracle GoldenGate (OGG). DataHub also supports the collection of surveillance videos based on the GB/T 28181 standard.

Data utilization

You can build an application to consume the data in DataHub, process the data in real time, and output the process results.

You can also use another application to process the streaming data output from the previous application to form a directed acyclic graph (DAG)-based data processing procedure.

Data archiving

You can create a DataConnector to periodically archive data in DataHub to MaxCompute.

14.Security services

14.1. Apsara Stack Security

Apsara Stack Security is a security solution that is designed to protect cloud assets in a comprehensive manner, including security operations, host security, application security, network security, data security, content security, O&M audit, and security service.

14.1.1. What is Apsara Stack Security?

Apsara Stack Security consists of Apsara Stack Security Standard Edition and optional security services. This topic describes the details of Apsara Stack Security Standard Edition and optional security services.

Overview

- Apsara Stack Security Standard Edition provides the following services: Network Detection and Response, Server Guard, Security Audit, Web Application Firewall (WAF), Threat Detection Service (TDS), Security Operations Center (SOC), and on-premises security operations service.
- Apsara Stack Security also provides various optional security services such as Bastionhost.

Network Detection and Response

Network Detection and Response detects and responds to network attacks at the Internet border and internal network borders based on a variety of threat detection engines and threat intelligence data.

Feature	Description
Traffic collection and analysis	Allows users to view statistics and details of traffic that flows through the Internet firewall or an internal firewall to a virtual private cloud (VPC).
Threat detection	Detects network intrusions such as web attacks, remote code execution, injections, scan attacks, brute-force attacks, and webshell uploads.
Threat alerting	Generates alerts for detected attacks. You can view the details of an alert to obtain the information about an attack, such as the source IP address, attack type, threat level, payload, and request and response content. You can also download raw Packet Capture (PCAP) packets for further analysis.
Threat response	Allows you to handle alerts in an efficient manner. For example, you can block alerts or add alerts to the whitelist with a few clicks.
Attacker profile	Analyzes and displays the basic information, threat intelligence, attack methods, attack targets, and attack processes about the attacks that are initiated from an IP address in a centralized manner.
Policy management	Allows you to manage network-layer and application-layer protection policies.

Behavior analysis	Allows you to analyze domain names that are generated by using domain generation algorithms (DGA), Domain Name System (DNS) tunneling, and the behavior of encrypted traffic.
Log retrieval	Allows you to view and retrieve traffic logs and security event logs that are generated within a specified period of time.

Server Guard

Server Guard protects Elastic Compute Service (ECS) instances against intrusions and malicious code.

Feature	Description
Baseline check	Performs security baseline checks for ECS instances. The check items include accounts, weak passwords, and at-risk configuration items. The baseline checks ensure that the ECS instances comply with the security standards for enterprise servers.
Vulnerability management	<ul style="list-style-type: none">Scans ECS instances for software vulnerabilities and provides suggestions on vulnerability fixes.Provides quick fixes for high-risk vulnerabilities in applications and operating systems on your ECS instances.
Webshell detection and removal	Detects and removes webshells based on specified rules and allows you to manually quarantine webshells.
Brute-force attack blocking	Detects and blocks brute-force attacks in real time.
Unusual logon alerting	Detects unusual logons based on the approved logon settings and generates alerts.
Suspicious server detection	Detects suspicious activities such as reverse shells, Java processes that run CMD commands, and unusual file downloads by using Bash.
Asset fingerprint	Collects up-to-date information about the servers, such as ports, accounts, processes, and applications, to perform event tracking.
Log retrieval	Centrally manages server logs of processes, networks, and system logons. This helps you locate the cause of an issue based on the logs.

Security Audit

Security Audit summarizes and analyzes logs so that security auditors can detect and eliminate risks in time.

Feature	Description
Raw log collection	Collects the following types of logs: <ul style="list-style-type: none">• Database and server logs• Operation logs of the user console and the IT administrator console• Network device logs
Audit log query	Allows you to query audit logs by audit type, audit object, operation type, operation risk level, alert, or creation time. Full-text search is supported.
Policy configuration	Allows you to configure audit rules based on the following parameters: initiator, object, command, result, and cause. This feature identifies risky operations in raw logs and generates alerts.

WAF

WAF protects web applications against attacks and ensures that mobile and PC users can securely access web applications over the Internet.

Feature	Description
Protection overview	Provides the following capabilities: <ul style="list-style-type: none">• Protection 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 protection report: allows you to export daily reports, weekly reports, and reports of scheduled tasks.• Attack detection statistics: provides statistics on attack detection.
Protection logs	Provides the following capabilities: <ul style="list-style-type: none">• 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 the log details of each attack.• HTTP flood protection logs: provides HTTP flood protection logs. The log list displays logs for matched HTTP flood protection rules, including the request URLs, the names of the matched rules, 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 logs: provides system operation logs, including usernames, operations, and IP addresses.• Access logs: provides access logs, including the access addresses, destination IP addresses, source IP addresses, request methods, and HTTP status codes of requests.

Protection configuration	<p>Provides the following capabilities:</p> <ul style="list-style-type: none">• Protection site management: allows you to create, delete, modify, enable, and disable feature forwarding proxies of a protected website.• Custom rules: allows you to create, delete, enable, and disable custom rules. This implements fine-grained HTTP access control for websites.• Website protection policies:<ul style="list-style-type: none">◦ Supports decoding methods, such as URL decoding, JSON parsing, Base64 decoding, hexadecimal conversion, backslash unescape, XML parsing, PHP deserialization, and UTF-7 decoding.◦ 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.◦ 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, separately enable or disable each attack detection module, and configure the detection granularity. WAF also allows you to specify the Block Status Code parameter.◦ Allows you to enable HTTP response detection and configure the length of the response body in detection rules.◦ Allows you to configure the length of the request body in detection rules.◦ 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. WAF restricts the access frequency of known IP addresses or sessions or blocks these IP addresses or sessions. WAF supports the whitelist feature for HTTP flood protection. HTTP flood protection rules are not applicable to IP addresses or sessions in a whitelist.
System management	<p>Displays the workload, network, and detection status of a node. You can configure syslog to send logs and also configure the service- and system-related alert thresholds.</p>

TDS

TDS monitors traffic and overall security status to audit and centrally manage assets.

Feature	Description
---------	-------------

Threat monitoring		Displays the security situation of Apsara Stack tenants. It is used to monitor network security in the Apsara Stack environment.
Risk analysis	Security monitoring	Displays security alert data in three dimensions in a centralized manner: network attacks, abnormal behavior, and vulnerabilities. It is used for security monitoring of Apsara Stack tenants.
	Cloud service check	Checks whether the security configurations of Apsara Stack products and services have security risks. It is used to scan and manage the security compliance of Apsara Stack products.
	Traffic analytics	Displays the statistics of network traffic. It is used for risk analysis of network traffic.
Assets		Displays the list of ECS instance assets or virtualized assets, as well as basic asset information, network attacks, abnormal behavior, and vulnerabilities in a centralized manner. It is used to manage asset risks.

SOC

SOC manages the overall security operations of the Apsara Stack environment. You can build a closed-loop security operations system that provides risk prediction and discovery, defense control, detection and analysis, and response management. SOC is a cloud-native solution.

Feature		Description
Operations Center		Collects and displays the real-time security alerts of various security products of Apsara Stack Security in a centralized manner. It supports security operations in scenarios such as security inspection, security monitoring, and attack-defense confrontation.
Threat Monitoring		Displays the global security situation, asset risks, and risk threat situation of the Apsara Stack platform and all Apsara Stack tenants in a centralized manner. It is used to monitor network security in the Apsara Stack environment.
Risk Analysis	Threat Events	Analyzes the correlation of standardized logs based on predefined or custom security analysis rules to generate indicators of compromise (IOC). Then, the paths and evidence of attacks can be found to identify attack sources.
	Threat Tracing	Automates intelligent threat tracing based on abnormal host behavior, network attacks, and threat events to restore attack paths in a visual manner and locate attack sources and targets.

	Vulnerability Analytics	Monitors and analyzes the security vulnerabilities and security configuration baselines of the ECS instance assets or virtualized assets of the Apsara Stack platform and tenants in a centralized manner.
	Cloud Service Check	Checks the security configurations of Apsara Stack tenant products and services for security risks. It is used to scan and manage the security compliance of Apsara Stack products.
	Traffic Analysis	Displays the global network traffic statistics of the Apsara Stack platform and tenants to help with risk analysis of network traffic.
Threat Intelligence	Intelligence Overview	Displays the statistics of threat intelligence such as highly active IP addresses, domain names, and URLs in Apsara Stack in a centralized manner. It is used to manage threat intelligence.
	Intelligence Query	Queries and confirms threat intelligence by IP address, domain name, or file MD5. It is used to identify malicious IOC.
	Intelligence Management	Displays the authorization of threat intelligence and the statistics of intelligence data application in a centralized manner. It is used for the consumption, configuration, and management of threat intelligence.
Assets	Platform Assets	Displays the list of platform-side ECS instance assets, as well as basic asset information and abnormal behavior in a centralized manner. It is used to manage the risks on the assets on the Apsara Stack platform.
	Tenant Assets	Displays the list of tenant-side ECS instance assets or virtualized assets, as well as basic asset information, network attacks, abnormal behavior, and vulnerabilities in a centralized manner. It is used to manage risks on tenant-side assets.
Process Task		Allows you to analyze and handle global security risks in a streamlined process. This facilitates security operations and management.
Logs	Overview	Displays predefined and custom statistical views of access logs in a centralized manner. It is used for statistical analysis of access logs.
	Log Query	Allows you to query the global correlations of standardized logs and export the query results.
	Platform Log	Allows you to query Apsara Stack platform-side log source logs.
	Tenant Log	Allows you to query tenant-side log source logs.

	Log Configurations	<ul style="list-style-type: none"> Allows you to configure and manage the log sources of Apsara Stack Security and third-party network elements. Allows you to configure standardized log parsing templates, access log parsing, and standardized processing. Allows you to configure the forwarding of logs to a third-party server.
Reports		Allows security operations administrators to export reports and create daily, weekly, and monthly reports. After the reports are created, the feature sends the reports to the specified email addresses.
Rules	Analytic Rules	Allows you to manage built-in and custom association analysis rules. The analytical computing engine can be called to implement advanced threat detection.
	Alert Rule	Allows you to manage custom alert rules and provides a semi-automated alert push for security operations.
	Block Rules	Blocks requests from specified IP addresses with a few clicks to shield security alerts.
Operations	Security Audit	Supports log management, log audit, log query, and audit policy management for ECS instances, network devices, databases, user operations, and O&M operations.
	Storage Management	Collects statistics on and manages storage space usage.
	IP Address Library	Allows you to customize the location of IP addresses.
	Organization Authorization	Supports service-linked role (SLR) management for global organizations and single organizations.

On-premises security service

The on-premises security service provides operations management on tenant systems at premises to ensure the security of tenant systems.

Category	Feature	Description
----------	---------	-------------

Security operations of tenants	Asset research	Performs authorized research for your business systems in Apsara Stack at regular intervals and develops a business list that contains information such as the name of the business system, ECS information, Relational Database Service (RDS) information, IP address, domain name, and owner.
	New system assessment	<ul style="list-style-type: none"> • Detects system and application vulnerabilities in a new business system by using automation tools and manual operations before you migrate the system to Apsara Stack. • Provides suggestions and verification on vulnerability fixes.
	Periodic security assessment	<ul style="list-style-type: none"> • Periodically uses automation tools to detect system vulnerabilities, application vulnerabilities, and security risks in running business systems. • Provides suggestions on handling detected risks, including but not limited to security policy settings, patch updates, and application vulnerability handling.
	Access control management	Provides inspection and guidance on access control policies when you migrate a new business system to Apsara Stack.
	Access control inspection	Periodically checks for access control risks of your business systems.
	Security risk inspection	Monitors and inspects security events in Apsara Stack Security, informs you of the verified events, and provides suggestions on event handling.
Security operations of Apsara Stack Security	Rule update	Periodically updates the rules repository of Apsara Stack Security.
	Service integration	<ul style="list-style-type: none"> • Provides support for integrating Apsara Stack Security with your business systems. • Helps you customize and optimize security policies.
Emergency response	Security notices	Synchronizes security notices from Alibaba Cloud and helps you handle risks.
	Event handling	Handles urgent events such as attacker intrusions.

SDDP

SDDP prevents data leaks and helps your business system meet compliance requirements.

Feature	Description
Security situation overview	Allows you to view the overall security status of sensitive data.
Detection and processing of suspicious activities	Detects suspicious activities related to sensitive data and allows you to confirm or exclude the activities after manual verification.
Sensitive data detection	Detects sensitive data in services such as MaxCompute, Tablestore, Object Storage Service (OSS), AnalyticDB for MySQL, and ApsaraDB RDS.
Static data masking	Uses data masking algorithms to mask sensitive data at rest.
Intelligent audit	Allows you to create audit rules to intelligently perform audits on services such as OSS, MaxCompute, and ApsaraDB RDS.
Data permission management	Displays departments and users in a hierarchical structure, displays users and accounts by type, and allows you to query and manage detailed permissions of accounts.
Dataflow monitoring	Allows you to view the dataflow details of DataHub and Data Integration.
Rule configuration	Allows you to configure detection rules, risk levels, and abnormal output rules to detect sensitive data.
Access authorization	Supports department-specific authorization and protects the data assets of authorized departments.

14.1.2. Benefits

Apsara Stack Security is a leading service in the cloud security industry and has received various certifications from relevant authorities. Apsara Stack Security fully protects the security of the Apsara Stack environment based on advanced security systems and security technologies.

Leading service in the cloud security industry

The Apsara Stack Security team accumulated a large amount of security experience by protecting all internal business systems of Alibaba Group from 2005. After the service was released in 2011, Apsara Stack Security became a leading service that provides comprehensive protection to ensure cloud security.

Mature systems and advanced technologies

Apsara Stack Security is a service that is developed based on ten years of protection experience. After the service protected the internal business systems of Alibaba Group for more than 10 years, Alibaba Group obtained a large number of security research achievements, security data, and security operations methods, and built a professional cloud security team. Apsara Stack Security is developed based on the rich experience of the security team to help enterprises build sophisticated systems that enhance security for cloud computing platforms. This service can protect the cloud platform, cloud network environments, and cloud business systems of Apsara Stack users.

Comparison with traditional security services

Benefit	Traditional security service	Apsara Stack Security
Comprehensive industry-leading security capabilities among Internet enterprises	A traditional security service provider offers only limited services and features, and cannot provide a comprehensive security protection system.	Alibaba Group accumulated a large number of intelligence sources based on multiple years of attack prevention experience. As a result, the service can detect common Internet attacks including zero-day exploits, and provide comprehensive security capabilities.
Early risk detection	Traditional security services cannot detect risks because they do not have experience in monitoring a wide array of services.	Apsara Stack Security can detect and quickly respond to critical vulnerabilities and security events to prevent security issues.
Big data modeling and analysis	Traditional security services cannot detect threats by scanning signatures. The traditional log analysis feature can be used to only collect data and report analysis results. You cannot use this feature to perform data modeling and analysis.	Apsara Stack Security implements big data modeling and analysis to detect threats in the entire network and display the security data. More than 30 algorithmic models are used to analyze the historical data, network data, and server data. This way, Apsara Stack Security provides additional insights into the security of the system.
Scalability and decoupling from hardware	Traditional security services are developed based on existing hardware devices. These security services require virtual hosts that are created on virtualization platforms.	<ul style="list-style-type: none">• Hardware and software decoupling: All modules are developed based on the cloud computing architecture and the common x86 architecture for hardware. As a result, the modules do not require specific hardware.• Scalability: You can scale out hardware devices to increase performance without the need to change the network architecture.

Comprehensive protection system that supports unified detection and response	Traditional security services add devices to improve security capabilities. The devices can collect only device logs and status data, and display the data on the management platform. You cannot integrate the devices to use additional features.	Apsara Stack Security provides comprehensive Internet protection to ensure the security of networks, servers, applications, data, and user identities. The security modules automatically interact with each other to respond to security issues and share intelligence.
Compatibility with all data center environments and decoupling from specific cloud platforms	Most traditional security services are provided as hardware appliances. Due to this reason, the services are incompatible with the cloud platforms that adopt Software Defined Network (SDN) technology.	Based on the interactions between servers and the operating system, Apsara Stack Security performs data analysis to detect threats at the network perimeter. This helps ensure that Apsara Stack Security is compatible with all data center environments and prevents the complex network topology in the data centers.

14.1.3. Scenarios

Apsara Stack Security provides flexible and scalable security solutions for users of different scales and from various sectors, such as industry, agriculture, transportation, public service, finance, and education.

Apsara Stack for classified protection compliance

Apsara Stack can help deploy the security modules provided by Apsara Stack Security on the computing platforms of users to implement in-depth defense. The modules include server security, network security, application security, data security, O&M audit, and centralized management modules. This helps detect security risks on the platforms and services in real time. This also helps meet the requirements of Multi-Level Protection Scheme (MLPS) 2.0 level 3 and MLPS 2.0 level 4. Apsara Stack also provides on-premises security operations and MLPS compliance consultation to help users ensure the security and stability of business and meet the requirements of MLPS 2.0 level 3 and MLPS 2.0 level 4.

Apsara Stack for public services

Apsara Stack for public services allows users from the public service sector to build a centralized service platform for public services in the Apsara Stack environment. This way, the users can work on intelligent platforms and provide intelligent convenience services for the public. The work efficiency of the users is greatly improved, and the public receives better services. Users can deploy Anti-DDoS Service and Web Application Firewall (WAF) on the cloud to establish the first and outermost-level line of defense, and use Apsara Stack Security to ensure the operational security, data security, and O&M security of public service systems. Apsara Stack Security can also provide security services to help build an organizational structure for security management, improve management policies, and foster an awareness of security operations.

Apsara Stack for finance

Apsara Stack for finance can help users build an in-depth security protection system. Users can deploy Network Detection and Response, Anti-DDoS Service, and Cloud Firewall on the egress of Apsara Stack to perform in-depth detection and protection against attacks at the network layer, deploy WAF to filter attacks at the application layer, and deploy the Server Guard agent on hosts to ensure terminal security. Apsara Stack Security collects network-wide security logs to perform centralized big data modeling and analysis and eliminate security silos. Users can deploy security protection systems on the cloud based on the x86 architecture, eliminating the need of traditional security hardware.

14.2. Key Management Service

Key Management Service (KMS) is an end-to-end service platform for key management and data encryption. KMS provides simple, reliable, secure, and standard-compliant capabilities to encrypt and protect data. KMS greatly reduces your costs of procurement, O&M, and research and development (R&D) on cryptographic infrastructure and data encryption products. This way, you can focus more on your business.

14.2.1. Service description

The features of Key Management Service (KMS) include customer master keys (CMKs), Bring Your Own Key (BYOK), fully managed hardware security modules (HSMs), rotation of encryption keys, CMK aliases, and resource tags.

CMK

Key Management is a core component of KMS. CMKs are used as encryption keys. CMKs can be classified into the following types:


- Symmetric keys: mainly used to encrypt or decrypt data. If you do not specify the KeySpec parameter during key creation, KMS creates a symmetric key. You can call the Encrypt or Decrypt operation to encrypt or decrypt data without the need to obtain the plaintext key material of a symmetric key.
- Asymmetric keys: mainly used for data encryption and digital signatures. An asymmetric CMK in KMS consists of a public key and a private key, which are cryptographically related to each other. The public key can be made available for anyone to use, but the private key must be kept secure. KMS ensures the security of the private key and does not provide an API to export the private key of an asymmetric key. You can use the private key to perform data decryption or digital signature by calling the API for private key calculation. Anyone with a public key can use it to encrypt data or verify the signature generated by the corresponding private key.

BYOK

KMS allows you to encrypt your data in the cloud by using the BYOK feature. This feature helps you meet stringent security and compliance requirements. We recommend that you use a managed HSM to protect your keys. You can import your keys into a CMK whose protection level is HSM. Keys in a managed HSM can only be destroyed, and their plaintext cannot be exported.

Fully managed HSMs

KMS uses qualified managed HSMs to securely generate keys, store keys, and perform cryptographic calculations. You can import external keys to a managed HSM to enable or disable keys, manage the lifecycle of keys, set key aliases or tags, and call cryptographic operations. This helps protect the most sensitive computing tasks and assets.

 **Note** To use this feature, you must purchase an HSM and the KMS license of the advanced edition.

Rotation of encryption keys

You can periodically rotate keys and configure key versioning to enhance the security of your keys, and implement security policies and best practices for data protection. KMS supports automatic key rotation and manual key rotation.

- Automatic key rotation: A CMK in KMS can have multiple key versions. Each version represents an independently generated key and does not have any relation with other versions. KMS automatically rotates encryption keys. This helps you implement the best security practices and comply with audit requirements.
- Manual key rotation: If your CMKs do not support version-based automatic rotation, you can manually rotate the CMKs based on your scenario.

CMK aliases

KMS allows you to create CMK aliases, which facilitate CMK usage.

An alias is optionally used to identify a CMK. It must be unique within an organization account and a region. An organization account can have the same alias in different regions. An alias can be bound to only one CMK in a region, but a CMK can have multiple aliases.

Resource tags

You can use resource tags to manage key resources in KMS.

14.2.2. Benefits

Compared with key management infrastructure (KMI), Key Management Service (KMS) features multi-service integration and ease of use.

Multi-service integration

KMS is integrated with multiple Apsara Stack services, such as Elastic Compute Service (ECS), ApsaraDB RDS, and Object Storage Service (OSS). You can easily use customer master keys (CMKs) in KMS to encrypt and control the data stored in these services and manage the cloud computing and storage environments.

Ease of use

Product value	Description
Easy encryption	KMS simplifies cryptographic concepts and provides cryptographic API operations that allow you to easily encrypt and decrypt data. For applications that require a key hierarchy, KMS provides convenient envelope encryption that implements a key hierarchy. For example, KMS generates a data key (DK) and uses a CMK as a key encryption key (KEK) to protect the DK.
Centralized key hosting	KMS provides centralized key hosting and control. You can import keys to KMS from KMI or from hardware security modules (HSMs) of Data Encryption Service. No matter whether keys are imported from external sources or created in KMS, their confidential information or sensitive data is used by other Apsara Stack services for data encryption and protection.
BYOK	KMS supports the Bring Your Own Key (BYOK) feature. You can import your own keys to KMS to encrypt data in the cloud. This facilitates key management.

Custom key rotation policies	<p>KMS supports the automatic rotation of symmetric encryption keys based on security policies. You can use key rotation to achieve the following objectives:</p> <ul style="list-style-type: none">• Reduce the amount of data that is encrypted by using a key• Respond in advance to security events• Provide logical isolation of data• Reduce the time window during which keys can be cracked• Meet the requirements of regulatory compliance
------------------------------	---

14.2.3. Scenarios

You can use Key Management Service (KMS) to encrypt sensitive data in your cloud to ensure the security of the data. For example, KMS provides encryption key protection for application developers. This ensures the security of sensitive data in application systems.

Envelope encryption

Your CMKs are stored in KMS. You only need to deploy enveloped data keys (EDKs). You can call the Decrypt API operation to decrypt the EDKs and use the returned plaintext DKs to encrypt or decrypt your local business data.

Direct encryption

You can call the Encrypt or Decrypt API operation of KMS to directly encrypt or decrypt sensitive data by using CMKs.

Server-side encryption

You can use the server-side encryption feature of Apsara Stack services to encrypt and protect data in a simple and effective way. For example, you can use the server-side encryption feature of Object Storage Service (OSS) to protect buckets that store sensitive data or use transparent data encryption (TDE) to protect tables that contain sensitive data.

15. Application Services

15.1. API Gateway

API Gateway helps enterprises quickly build an API-centric system architecture for scenarios such as technology introduction, system integration, and capability packaging. In addition, API Gateway provides mechanisms to secure APIs and reduce the risks arising from APIs. These mechanisms include anti-replay protection, request encryption, identity authentication, permission management, and throttling. API Gateway automatically generates SDK references and API references. This improves the efficiency of API management and iteration. API Gateway improves the reusability of different capabilities. This accelerates business innovation inside enterprises.

15.1.1. Features

API Gateway is a fully hosted service for APIs. You can use it to design, develop, test, publish, sell, operate, maintain, monitor, supervise, and unpublish APIs.

API lifecycle management

- Manages APIs throughout their entire lifecycle, including publishing, testing, and unpublishing APIs.
- Supports API maintenance features such as routine management, version management, and quick rollback.
- Allows APIs in different environments to be accessed by using different domain names or headers.
- Supports the API diff feature. When you release a new version of an API, you can check the differences between the new version and an earlier version of the API.

Comprehensive security protection

- Supports multiple authentication methods, including anonymous access, simple authentication, signature authentication, and JSON Web Token (JWT) authentication.
- Supports HTTPS and SSL encryption.
- Provides multiple security mechanisms to prevent injections, replay attacks, and tampering.
- Supports backend signature authentication. Your backend service can authenticate API Gateway based on signatures.

Flexible access control

- Manages the permissions of apps to make API calls.
- Allows only authorized apps to call specific APIs.
- Allows API owners to authorize apps to call specific APIs.

Various plug-ins

- Allows you to use various plug-ins that can be bound or unbound to expand the features of APIs.
- Provides the following plug-in types: throttling, IP address-based access control, backend signature, JWT authentication, cross-origin resource sharing (CORS), caching, routing, RAM, circuit breaker, and error mapping.

Request verification

Supports the verification of parameter types and values. The value ranges, enumerated values, and regular expressions in values can be verified. API Gateway denies the requests with invalid parameter types or values. This reduces backend resources that are wasted on invalid requests and significantly reduces the costs of processing requests for the backend services.

Data conversion

Supports data conversion for frontend requests. Allows you to configure rules for mappings between the frontend data and backend data.

Integration

Supports big data platforms as backend services. This allows you to use APIs to provide data services.

Automation tools

- Automatically generates API documentation.
- Provides SDK samples in multiple programming languages.
- Provides graphical debugging tools for quick testing and deployment.

Monitoring and alerting

- Provides a graphical panel for real-time API monitoring that displays information such as the number of API calls, response time, and error rate.
- Allows you to configure alerts to track the status of each API in real time.
- Delivers the logs about HTTP requests and responses of APIs to Log Service (SLS) for full log query and analysis.

15.1.2. Benefits

API Gateway features easy maintenance and high performance, stability, and security.

Easy maintenance

After you create APIs, API Gateway does all the complicated chores for you, such as documentation maintenance, version management, and SDK maintenance. This significantly reduces routine maintenance workloads.

High performance

API Gateway supports efficient access over HTTP/2 and maintains persistent connections by supporting the binary protocol WebSocket. This improves the performance of the connections between clients and API Gateway. API Gateway adopts distributed deployment and automatically scales out to handle a large number of API requests with low latency. API Gateway offers reliable and efficient features for your backend services.

Stability

Since its official commercial release on Alibaba Cloud in 2016, API Gateway has been running stably under various unfavorable conditions in Alibaba Cloud and Apsara Stack, such as large packets, unstable backend services, and indefinite response time.

Security

API Gateway implements SSL encryption in the full link of communication to protect all data against eavesdropping during transmission. API Gateway implements signature verification in the full link of communication to prevent data tampering during transmission. To ensure that your services are secure, stable, and controllable, API Gateway also provides a set of API security features. The features include strict permission management, replay attack prevention, parameter cleansing, IP address-based access control, precise throttling, and integration with Web Application Firewall (WAF) of Alibaba Cloud.

15.1.3. Scenarios

API Gateway can be used in scenarios such as business mid-end API management, multi-terminal compatibility, and system integration.

API management hub for business mid-ends

API Gateway can manage APIs of various systems in a centralized manner by leveraging its interconnection and integration capabilities. Centralized API management, including throttling, permission management, and monitoring, facilitates O&M and allows you to configure a single API that can be called by multiple systems. This significantly improves operational efficiency.

Multi-terminal compatibility

As mobile networks and IoT develop, APIs need to support more types of terminals in more business scenarios. However, this increases system complexity.

- In API Gateway, enterprises can manage and maintain APIs in a single service system and adapt APIs to different types of terminals, such as apps, devices, and web clients, only by changing API definitions.
- Enterprises can develop and manage a single API for multiple scenarios, multiple types of terminals and users, and multi-tier services. This reduces the costs and complexity of O&M.

System integration

- API Gateway helps standardize APIs of different systems. This way, you can integrate systems with standard APIs.
- API Gateway helps integrate and manage resources with efficiency and prevents resource redundancy and waste caused by fast development. This way, you can focus on business development.