Alibaba Cloud
Apsara Stack Agility SE

Security Whitepaper

Version: 1912, Internal: V3.1.0

Issue: 20200708

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 due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall 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 the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. 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 bear any liability for any errors or financial losses incurred by any organizations, companies, or individual s arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary , incidental, 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 contact Alibaba Cloud directly if you discover any errors in this document.

Document conventions

Style	Description	Example
⛔	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.
⚠	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.
ⓘ	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.
📋	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 `cd /d C:/window` command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	`bae log list --instanceid` Instance_ID
[] or [a|b]	This format is used for an optional value, where only one item can be selected.	`ipconfig [-all|-t]`

Style	Description	Example
{} or {a|b}	This format is used for a required value, where only one item can be selected.	switch {active|stand}

Contents

1 Security of Apsara Stack products

1.1 Object Storage Service (OSS)

1.1.1 Platform security

1.1.1.1 Security isolation

OSS slices user data and discretely stores the sliced data in a distributed file system based on specific rules. The user data and its indexes are stored separately. OSS users are authenticated using symmetric AccessKeys. The signature in each HTTP request is verified. If verification is successful, OSS reassembles the distributed data. This implements data storage isolation between multiple tenants.

1.1.1.2 Authentication and access control

1.1.1.2.1 Authentication

You can create an AccessKey on Apsara Stack Management Console. An AccessKey consists of AccessKey ID and AccessKey Secret. AccessKey ID is public and used to identify a user. AccessKey Secret is private and used to authenticate a user.

Before a request is sent to OSS, a signature string must be generated in the format specified by OSS for the request. Then, the signature string is encrypted using AccessKey Secret and the HMAC algorithm, to form a verification code. The verification code carries a timestamp to prevent replay attacks. After receiving the request, OSS finds the corresponding AccessKey Secret using AccessKey ID and extracts the signature string and verification code using the same method. If the calculated verification code matches the one provided, the request is valid. Otherwise, OSS rejects the request and returns an HTTP 403 error.

1.1.1.2.2 ACL settings

Access to OSS resources is divided into access by the owner and access by third-party users. An owner owns a bucket, while third-party users are other users who access resources in the bucket. Access can be either anonymous or signature-based. If the access is initiated with an OSS request without any identification information, the access is considered to be an anonymous access. Based on the rules in the OSS API documentation, if the access is initiated with a request that contains signature information in its header or carried URL, the access is signature-based.

OSS provides access control for buckets and objects.

Three kinds of bucket access permissions are available: public read/write, public read, and private.

- Public read/write: All users (including anonymous users) can perform write (Put, Get, and Delete) operations on objects in the bucket.
- Public read: Only the bucket creator can perform write (Put, Get, and Delete) operations on objects in the bucket. All users (including anonymous users) can perform read (Get) operations on objects in the bucket.
- Private: Only the bucket creator can perform read/write operations on objects in the bucket. The other users cannot access objects in the bucket.

Note:

When you create a new bucket without configuring bucket permissions, OSS automatically sets its access permission to private.

Four kinds of object access permissions are available: public read/write, public read, private, and default.

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

Note:

If you do not configure any bucket permission when uploading an object, the object will use the default access permission set by the OSS.

1.1.1.2.3 Support for RAM and STS

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

RAM is a resource access control service provided by Alibaba Cloud. RAM allows you to create sub-accounts under a primary account. All resources belong to the primary account. The primary account can grant access permissions on resources to sub-accounts.

Alibaba Cloud STS provides temporary access credentials and short-term access permission management. STS can generate a temporary access credential for users. The access permission and expiration date of the credential are user-defined. The access credential expires automatically upon the expiration date.

1.1.1.3 Data security

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

CRC

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

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

MD5 verification

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

1.1.1.4 Data encryption

1.1.1.4.1 Server-side encryption

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

When creating an object, you only need to add the HTTP header of x-oss-server-side-encryption to the Put Object request and set its value to AES256. Then, the object can be encrypted on the server side before it is stored.

1.1.1.4.2 Client-side encryption

OSS allows you to use client-side encryption to encrypt data before the data is sent to the remote server while the data encryption key (DEK) used is kept only on the local client. Other users cannot obtain the raw data without the DEK and Enveloped Data Key (EDK), even if the data is leaked. OSS uses functions in the SDK to encrypt the data locally before the data is uploaded to the OSS bucket.

1.1.1.4.3 KMS-based encryption

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

1.1.2 Tenant security

1.1.2.1 Log audit

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

1.1.2.2 Hotlink protection

To prevent your data in OSS from being leeched, you can configure hotlink protection through the following parameters:

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

For example, for a bucket named oss-example, you can add **http://www.aliyun.com/** to the Referer whitelist. Then, requests with a Referer of **http://www.aliyun.com/** can access objects in this bucket.

1.2 ApsaraDB for RDS

1.2.1 Platform security

1.2.1.1 Secure isolation

Tenant isolation

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

1.2.1.2 Authentication

ApsaraDB for RDS secures data through authentication.

Identity authentication

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

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

Permission control

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

Access control

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

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

1.2.1.3 Data security

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

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

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

1.2.1.4 Data encryption

SSL

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

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

1.2.1.5 DDoS attack prevention

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

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

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

- **Traffic scrubbing**

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

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

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

- **Black hole filtering**

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

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

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

1.2.2 Tenant security

1.2.2.1 Log audit

ApsaraDB for RDS can audit logs to identify security issues.

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

1.2.2.2 IP address whitelist

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

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

1.2.2.3 Software update

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

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

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

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

1.3 AnalyticDB for PostgreSQL

1.3.1 Platform security

1.3.1.1 Security isolation

Network isolation

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

A VPC is a private network environment that you can set in Apsara Stack to strictly isolate network packets at the network layer by using network protocols and control access.

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

- Access requests from ECS EIPs.
- Access requests from the Internet egress of your on-premises IDC.

Note:

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

Tenant isolation

AnalyticDB for PostgreSQL uses virtualization to isolate tenants and grants each tenant their own database permissions. Alibaba Cloud also hardens security for database servers. For example, to prevent other users from accessing your data, users cannot use a database to read or write operating system files.

1.3.1.2 Authentication

AnalyticDB for PostgreSQL instances created by your Apsara Stack tenant account are also owned by the account. Alibaba Cloud tenant accounts have full access permissions on their resources.

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

1.3.1.3 Primary and secondary nodes

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

1.3.2 Tenant security

1.3.2.1 Database account

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

1.3.2.2 IP address whitelist

AnalyticDB for PostgreSQL instances cannot be accessed from any IP addresses by default. The default whitelist contains `127.0.0.1`. You can add IP addresses to a whitelist on the Security Controls page of the console or by using an API operation. Updating the IP address whitelist does not require an instance to restart nor affect operations on the instance.
You can configure multiple IP address whitelists. Each whitelist can contain up to 1,000 IP addresses or CIDR block entries.

1.3.2.3 SQL audit

AnalyticDB for PostgreSQL allows you to view SQL details. You can audit SQL operations on a regular basis to locate problems in a timely manner. The Proxy module records the information of all SQL statements executed in AnalyticDB for PostgreSQL, including the IP address, the name of the accessed database, the account that executed the statement, the

SQL statement, the execution duration, the number of returned records, and the execution time point.

1.3.2.4 Backup and restoration

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

1.3.2.5 Software update

- New versions of database software are provided by AnalyticDB for PostgreSQL on a regular basis.
- Software updates are optional and only implemented when you request them.
- If the current database version that you are using contains critical security risks, the AnalyticDB for PostgreSQL team will notify you and recommend that you schedule an update. The AnalyticDB for PostgreSQL team can provide full support throughout the update process.
- AnalyticDB for PostgreSQL updates are usually completed within five minutes. During updates, instances may be disconnected several times and will become read-only for about a minute. There is minimal interruption to services if the database reconnection settings or connection pool are properly configured for your applications.

1.4 Data Transmission Service (DTS)

1.4.1 Platform security

1.4.1.1 Security isolation

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

1.4.1.2 Authentication

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

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

1.4.1.3 Transmission security

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

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

DTS also supports HTTPS to effectively improve access security.

1.4.1.4 Data security

When you use DTS to subscribe to incremental data, a large portion of incremental data is stored on the DTS server. The incremental data is serialized and stored based on the storage format defined in DTS. The DTS-defined storage format provides enhanced data security.

Note:

Data written to the DTS server is automatically deleted after it is stored for seven days.

1.5 KVStore for Redis

1.5.1 Platform security

1.5.1.1 Security isolation

Tenant isolation

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

Network isolation

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

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

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

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

Notice:

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

1.5.1.2 Authentication

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

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

1.5.1.3 Transmission encryption

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

Note:

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

1.5.2 Tenant security

1.5.2.1 Database account

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

1.5.2.2 IP address whitelist

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

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

1.5.2.3 Backup and recovery

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

1.5.2.4 Software upgrade

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

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

1.6 Distributed Relational Database Service (DRDS)

1.6.1 Platform security

1.6.1.1 Isolation

Network isolation

Distributed Relational Database Service (DRDS) supports advanced network access control by using a Virtual Private Cloud (VPC).

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

1.6.1.2 Authentication

Distributed Relational Database Service (DRDS) supports a MySQL-like account and permission system and supports commands and functions such as GRANT, REVOKE, SHOW GRANTS, CREATE USER, DROP USER, and SET PASSWORD.

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

- Permissions can be granted at the database and table levels. Currently, global permissions and column permissions are not supported.
- Eight associated basic permissions are supported: CREATE, DROP, ALTER, INDEX, INSERT, DELETE, UPDATE, and SELECT.
- The `user@'host'` format can be used to match and verify access to a host.

Note:

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

1.6.2 Tenant security

1.6.2.1 IP address whitelist

Distributed Relational Database Service (DRDS) provides an IP address whitelist to ensure secure access. Each DRDS database can be configured with an IP address whitelist.

By default, DRDS instances are set to be accessible from any IP address. You can add IP addresses to the DRDS whitelist on the **Whitelist settings** page in the console. Updating the IP address whitelist does not require restart of the DRDS instance, and does not affect operations on the instance. You can also set IP addresses or CIDR blocks in the IP address whitelist.

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

1.6.2.2 Protection against high-risk SQL misoperations

Distributed Relational Database Service (DRDS) prohibits high-risk operations such as full table deletion and update by default. You can temporarily skip this restriction by adding a hint. The following statements are prohibited by default:

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

The actual effect is as follows:

```
mysql> delete from tt;
ERR-CODE: [TDDL-4620][ERR_FORBID_EXECUTE_DML_ALL] Forbid execute DELETE ALL or
 UPDATE ALL sql. More: [http://middleware.alibaba-inc.com/faq/faqByFaqCode.html?
faqCode=TDDL-4620]
```

After the hint is added, the statements are successfully executed.

```
mysql> /*TDDL:FORBID_EXECUTE_DML_ALL=false*/delete from tt;
```

```
Query OK, 10 row affected (0.21 sec)
```

1.6.2.3 Slow SQL audit

In the Distributed Relational Database Service (DRDS) console, you can query the slow SQL statements sent by the client to DRDS. Slow SQL statements increase the response time (RT) of the entire link and reduce the DRDS throughput.

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

1.6.2.4 Monitoring information

The Distributed Relational Database Service (DRDS) console provides monitoring metrics in different dimensions. You can perform related operations based on the monitoring information.

DRDS monitoring information can be classified into two types:

- Resource monitoring information, including the CPU, memory, and network.
- Engine monitoring information, including the logical queries per second (QPS), physical QPS, logical response time (RT) (in ms), physical RT (in ms), number of connections, and number of active threads.

The QPS and CPU performance of a DRDS instance are in positive correlation. When DRDS encounters a performance bottleneck, the CPU usage of the DRDS instance remains high. If the CPU usage exceeds 90% or remains higher than 80%, the DRDS instance encounters a performance bottleneck. If there is no bottleneck in the DRDS instance, the current DRDS instance type cannot meet the QPS performance requirements of the business. In this case, upgrade the DRDS instance.

1.7 AnalyticDB for MySQL

1.7.1 Platform security

1.7.1.1 Security isolation

In AnalyticDB for MySQL, databases are the basic unit of tenant isolation. The Apsara Stack tenant account used to create a database is the owner of the database. The database owner must grant access permissions before other users can access the database. Each database runs on an exclusive process for each user, isolating databases at the process level.

Each AnalyticDB for MySQL database uses a multi-tenancy mechanism to completely isolate each process. Physical resources, such as CPU, memory, and storage space, are isolated between databases.

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

1.7.1.2 Authentication

The following figure shows the identity authentication and access control mechanisms of AnalyticDB for MySQL.

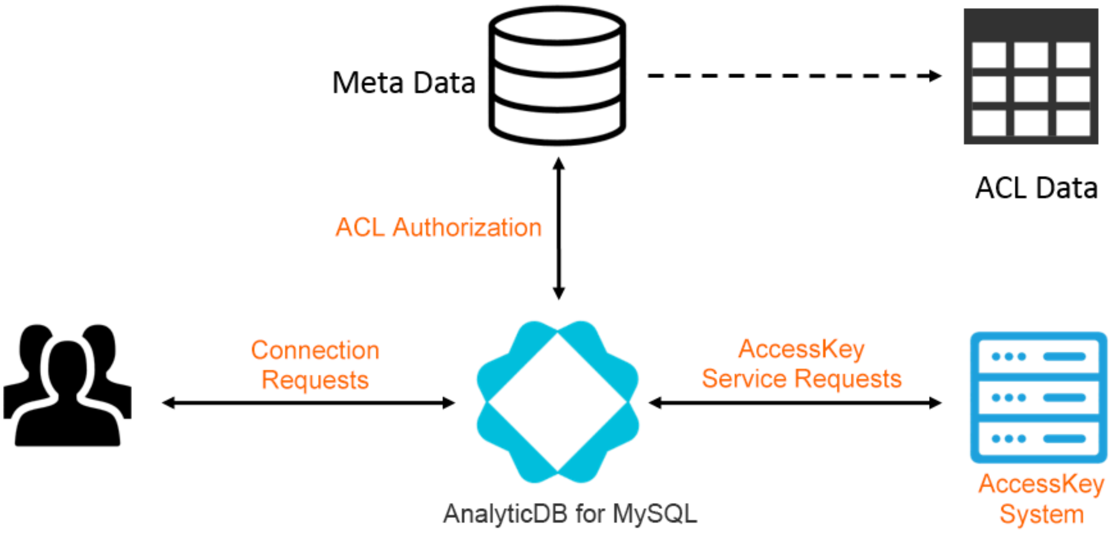

Identity authentication

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

Like other Alibaba Cloud services, AnalyticDB for MySQL uses the AccessKey mechanism to implement identity authentication. You can connect to AnalyticDB for MySQL by using an AccessKey pair or a driver such as Connector/J or Connector/ODBC.

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

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

Access control

AnalyticDB for MySQL uses access control list (ACL) rules to provide table-level permission management similar to those of MySQL. However, unlike MySQL, AnalyticDB for MySQL ACL does not provide host-based authorization.

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

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

AnalyticDB for MySQL provides the following authorization objects:

- Database: specifies a database or all tables in a database, such as `db_name.*` or `*` (default database).
- Table: specifies a table, such as `db_name.table_name` or `table_name`.
- Column: specifies a column in a specified table. It is composed of `column_list` and `Table`.

Access control

AnalyticDB for MySQL supports resource access management (RAM), but not security token service (STS).

RAM allows you to create RAM users by using an Apsara Stack tenant account and grant resource access permissions to RAM users. All RAM users created by you are affiliated with your Apsara Stack tenant account, which means that all of their resources will also belong to your Apsara Stack tenant account.

1.7.1.3 Data security

Multi-tenancy

AnalyticDB for MySQL provides tenant isolation. Resources (such as CPU, memory, disks, and network bandwidth) of different databases are completely isolated from each other to ensure tenant isolation.

Data reliability

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

Data consistency

AnalyticDB for MySQL uses a multiversion concurrency control (MVCC) mechanism to store changes to real-time table data resulting from INSERT and DELETE operations. This ensures that query results returned during concurrent data updates are consistent with the data version at the time the query was initiated.

Note:

You can clear outdated data versions at regular intervals.

1.7.2 Tenant security

1.7.2.1 Log audit

You can enable log audit to record all SQL operation information generated in AnalyticDB for MySQL. The information includes:

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

You can then use SQL statements to query historical data.

Example of the audit log format:

```
[2017-10-10 13:37:57,351] INFO [pool-31-thread-22] c.a.c.a.f.l.AccessLog.info - Client=
127.0.0.1 Total_time=1044 Exec_time=1043 Queue_time=1 - [2017-10-10 13:37:56 308] 1
SQL Statement \;process=2017101013375601000316310809999838042\;CLUSTER=ayads
-bjyz
```