RDS for MariaDB

User Guide

Issue 01

Date 2025-10-22

HUAWEI

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Trademarks and Permissions


and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei Cloud and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road
Qianzhong Avenue
Gui'an New District
Gui Zhou 550029
People's Republic of China

Website: https://www.huaweicloud.com/intl/en-us/

Contents

1 Using IAM to Grant Access to RDS

1.1 Creating a User and Granting Permissions

This chapter describes how to use **Identity and Access Management (IAM)** for fine-grained permissions management for your RDS resources. With IAM, you can:

- Create IAM users for employees based on your enterprise's organizational structure. Each IAM user will have their own security credentials for accessing RDS resources.
- Grant only the permissions required for users to perform a specific task.
- Entrust a Huawei Cloud account or cloud service to perform efficient O&M on your RDS resources.

If your Huawei Cloud account does not require individual IAM users, skip this chapter.

This section describes the procedure for granting permissions (see **Figure 1-1**).

Prerequisites

Learn about the permissions (see **Permissions**) supported by RDS and choose policies or roles according to your requirements. For the system policies of other services, see **System-defined Permissions**.

Process Flow

Figure 1-1 Process for granting RDS permissions

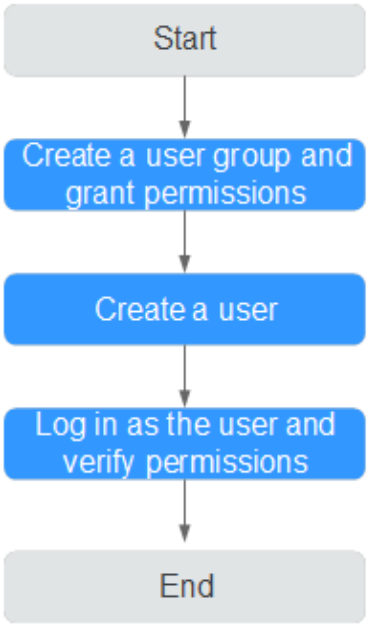

1. **Create a user group and assign permissions** to it.

   Create a user group on the IAM console, and attach the **RDS ReadOnlyAccess** policy to the group.

   NOTE

   To use some interconnected services, you also need to configure permissions of such services.

   For example, to connect to your DB instance through the console, configure the **DAS FullAccess** permission of Data Admin Service (DAS) besides **RDS ReadOnlyAccess**.

2. **Create an IAM user and add it to the user group**.

   Create a user on the IAM console and add the user to the group created in **1**.

3. **Log in** and verify permissions.

   In the authorized region, perform the following operations:

   - Choose **Relational Database Service** from the service list and click **Buy DB Instance**. If a message appears indicating that you have insufficient permissions to perform the operation, the **RDS ReadOnlyAccess** policy has already been applied.
   - Choose any other service from the service list. If a message appears indicating that you have insufficient permissions to access the service, the **RDS ReadOnlyAccess** policy has already taken effect.

1.2 Custom Policies

Custom policies can be created to supplement the system policies of RDS.

You can create custom policies in either of the following two ways:

- Visual editor: Select cloud services, actions, resources, and request conditions without the need to know policy syntax.
- JSON: Edit JSON policies from scratch or based on an existing policy.

This section contains examples of common RDS custom policies.

Example Custom Policies

- Example 1: Allowing users to create RDS DB instances

```
{
    "Version": "1.1",
    "Statement": [{
        "Effect": "Allow",
        "Action": ["rds:instance:create"]
    }]
}
```

- Example 2: Denying RDS DB instance deletion

  A policy with only "Deny" permissions must be used in conjunction with other policies to take effect. If the permissions assigned to a user include both "Allow" and "Deny", the "Deny" permissions take precedence over the "Allow" permissions.

  The following method can be used if you need to assign permissions of the **RDS FullAccess** policy to a user but you want to prevent the user from deleting RDS DB instances. Create a custom policy for denying RDS DB instance deletion, and attach both policies to the group the user belongs to. Then, the user can perform all operations on RDS DB instances except deleting RDS DB instances. The following is an example deny policy:

```
{
    "Version": "1.1",
    "Statement": [{
        "Action": ["rds:instance:delete"],
        "Effect": "Deny"
    }]
}
```


2 Buying an RDS for MariaDB Instance

Scenarios

This section describes how to buy a DB instance on the RDS console.

RDS for MariaDB supports the yearly/monthly and pay-per-use billing modes. You can tailor your compute resources and storage space to your business needs.

Procedure

Step 1 **Sign up for a HUAWEI ID and enable Huawei Cloud services**.

Step 2 Before purchasing DB instances, ensure that your account balance is sufficient. **Top up your account** if required.

Step 3 For fine-grained permissions management, create an Identity and Access Management (IAM) user and user group on the IAM console and grant the user specific operation permissions. For details, see **Creating a User and Granting Permissions**.

Step 4 Go to the **Buy DB Instance** page.

Step 5 On that page, click the **Custom Config** tab, select a billing mode, configure information about your DB instance, and click **Buy**.

- Engine Options

Figure 2-1 Basic information

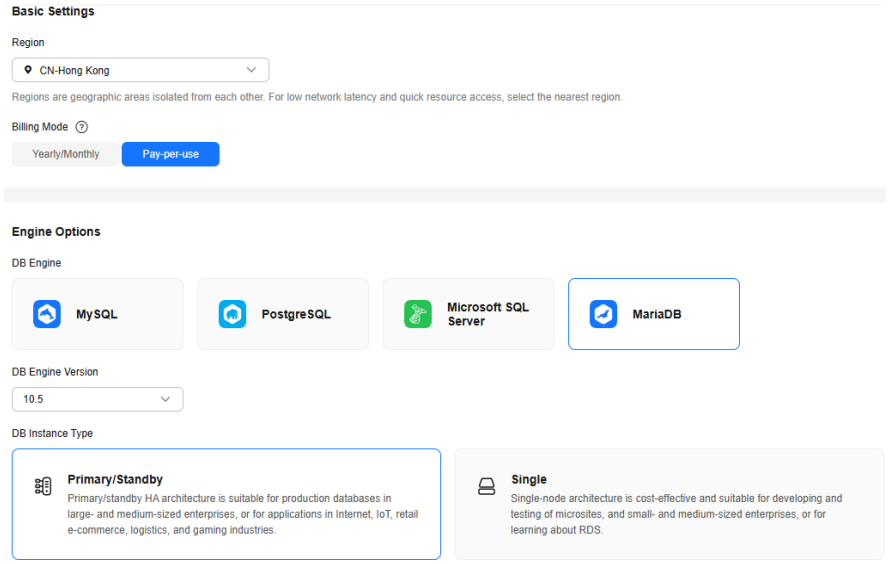

Table 2-1 Basic information

Parameter	Description
Region	The region where your resources are located. **NOTE** Products in different regions cannot communicate with each other through a private network. After a DB instance is created, the region cannot be changed. Therefore, exercise caution when selecting a region.
Billing Mode	– **Yearly/Monthly**: A prepaid billing mode in which you pay for resources before using it. Bills are settled based on the subscription period. The longer the subscription, the bigger the discount. This mode is a good option for long-term, stable services. – **Pay-per-use**: A postpaid billing mode. You pay as you go and just pay for what you use. The DB instance usage is calculated by the second but billed every hour. This mode allows you to adjust resource usage easily. You neither need to prepare for resources in advance, nor end up with excessive or insufficient preset resources.
DB Engine	MariaDB
DB Engine Version	For details, see **DB Engines and Versions**. Supported DB engine versions may vary by region. For the actual options, see them on the console.

Parameter	Description
DB Instance Type	– **Primary/Standby**: uses an HA architecture with a primary DB instance and a synchronous standby DB instance. It is suitable for production databases of large- and medium-sized enterprises in Internet, Internet of Things (IoT), retail e-commerce sales, logistics, gaming, and other sectors. When a primary instance is being created, a standby instance is provisioned along with it to provide data redundancy. The standby instance is invisible to you after being created. – **Single**: uses a single-node architecture, which is more cost-effective than primary/standby DB instances. It is only recommended for development and testing of microsites, and small and medium enterprises, or for learning about RDS.
AZ	An AZ is a physical region where resources use independent power supply and networks. AZs are physically isolated but interconnected through an internal network. Some regions support both single-AZ and multi-AZ deployment and some only support single-AZ deployment. To achieve high reliability, RDS will automatically deploy your primary and standby instances in different physical servers even if you deploy them in the same AZ. You can deploy primary and standby instances in a single AZ or across AZs to achieve failover and high availability.
Storage Type	The storage type determines the read/write speed of an instance. A higher maximum throughput enables faster I/O operations. – **Cloud SSD**: cloud disks used to decouple storage from compute. The maximum throughput is 350 MB/s. – **Extreme SSD**: The disks combine the 25GE network and RDMA technologies to provide you with up to 1,000 MB/s throughput per disk and sub-millisecond latency.

- Instance Configuration

Figure 2-2 Instance Configuration

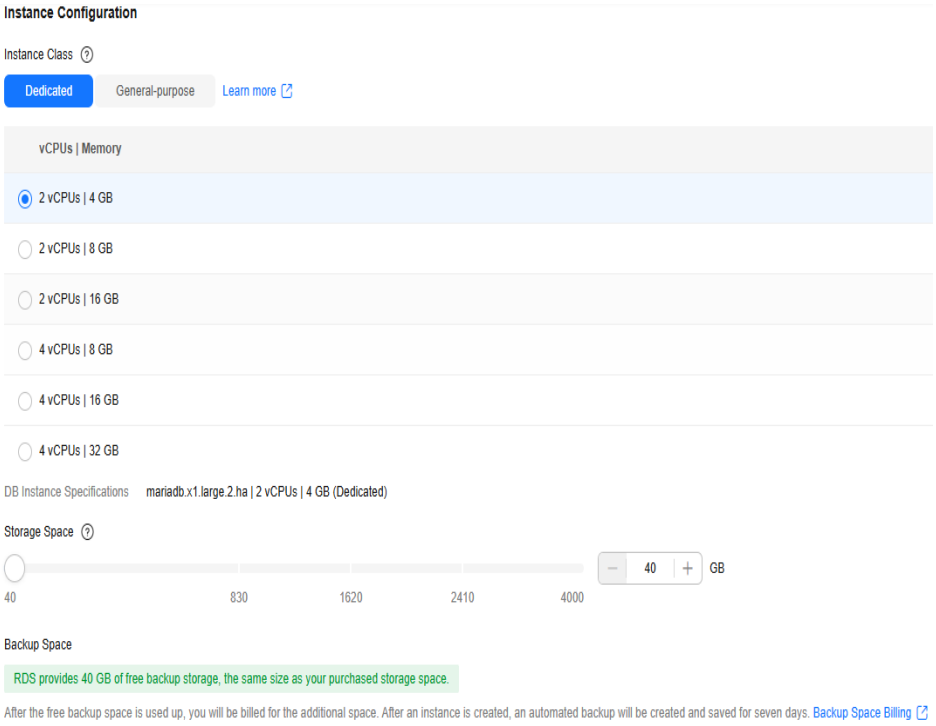

Table 2-2 Specifications and storage

Parameter	Description
Instance Class	Refers to the vCPU and memory of a DB instance. Different instance classes support different numbers of database connections and maximum IOPS. After a DB instance is created, you can **change its vCPUs and memory**.
Storage Space	Contains the system overhead required for inodes, reserved blocks, and database operation. Storage space can range in size from 40 GB to 4,000 GB and can be scaled up only by a multiple of 10 GB. After a DB instance is created, you can scale up its storage space. For details, see **Scaling Up Storage Space**.

- Basic Settings and Connectivity

Figure 2-3 Basic Settings and Connectivity

Table 2-3 Network

Parameter	Description
DB Instance Name	Must start with a letter and consist of 4 to 64 characters. Only letters (case-sensitive), digits, hyphens (-), and underscores (_) are allowed. – If you intend to buy multiple DB instances at a time, the allowed length for each instance name will change. – If you create multiple DB instances at a time, their names will include a four-digit suffix. For example, if you specify **instance** here, the names will be *instance-0001*, *instance-0002*, and so on. If existing instances' suffixes have already reached up to *0010*, the new instance names will start from *instance-0011*.
Administrator	The default login name for the database is **root**.

Parameter	Description
Administrator Password	Must consist of 8 to 32 characters and contain the following character types: uppercase letters, lowercase letters, digits, and special characters (~!@#$%^*-_=+?,()&). Enter a strong password and periodically change it for security reasons. If the password you provide is regarded as a weak password by the system, you will be prompted to enter a stronger password. Keep this password secure. The system cannot retrieve it. After a DB instance is created, you can reset this password. For details, see **Resetting the Administrator Password**.
Confirm Password	Must be the same as **Administrator Password**.
VPC	A virtual network in which your RDS DB instances are located. A VPC can isolate networks for different workloads. You can select an existing VPC or create a VPC. For details about how to create a VPC, see **Creating a VPC and Subnet**. If no VPC is available, RDS allocates a VPC to you by default. **NOTICE** After a DB instance is created, the VPC cannot be changed.
Subnet	Improves network security by providing dedicated network resources that are logically isolated from other networks. Subnets are only valid within a specific AZ. Dynamic Host Configuration Protocol (DHCP) is enabled by default for subnets where you plan to create RDS DB instances and cannot be disabled. A floating IP address is automatically assigned when you create a DB instance. You can also enter an unused IPv4 IP address in the subnet CIDR block.
Security Group	Enhances security by controlling access to your DB instance from other services. A network **access control list (ACL)** can help control inbound and outbound traffic of subnets in your VPC. Ensure that the security group you select allows the client to access the DB instance. If no security group is available or has been created, RDS allocates a security group to you by default.

- Additional Options

Table 2-4 Additional options

Parameter	Description
Enterprise Project	If your account has been associated with an enterprise project, select the target project from the **Enterprise Project** drop-down list. For more information about enterprise projects, see ***Enterprise Management User Guide***.
Parameter Template	Contains engine configuration values that can be applied to one or more DB instances. If you intend to create a primary/standby DB pair, they use the same parameter template. You can modify the instance parameters as required after the DB instance is created. For details, see **Modifying Instance Parameters**.
Time Zone	You need to select a time zone for your instance based on the region hosting your instance. You can change the time zone after the DB instance is created.
Table Name	Specifies whether table names are case sensitive. **NOTE** The case sensitivity of table names for created instances cannot be changed.
Tag	Tags an RDS instance. This parameter is optional. Adding tags to RDS instances helps you better identify and manage the DB instances. A maximum of 20 tags can be added for each DB instance. If your organization has configured tag policies for RDS, add tags to DB instances based on the policies. If a tag does not comply with the policies, DB instance creation may fail. Contact your organization administrator to learn more about tag policies. After a DB instance is created, you can view its tag details on the **Tags** page. For details, see **Managing Tags**.

- Required Duration and Quantity

Table 2-5 Required duration and quantity

Parameter	Description
Required Duration	This option is available only for yearly/monthly DB instances. The system will automatically calculate the configuration fee based on the selected required duration. The longer the required duration is, the larger discount you will enjoy.

Parameter	Description
Auto-renew	– This parameter is available only for yearly/monthly DB instances and is not selected by default. – If you select this option, the auto-renew cycle is determined by the selected required duration.
Quantity	RDS allows you to create DB instances in batches. If you choose to create primary/standby DB instances and set **Quantity** to **1**, a primary DB instance and a standby DB instance will be created at the same time.

NOTE

- If you have any questions about the price, click **Pricing details** at the bottom of the page.
- The performance of your DB instance depends on its configurations. Hardware configuration items include the instance specifications, storage type, and storage space.

Step 6 Confirm the specifications.

Figure 2-4 Confirming specifications (pay-per-use)

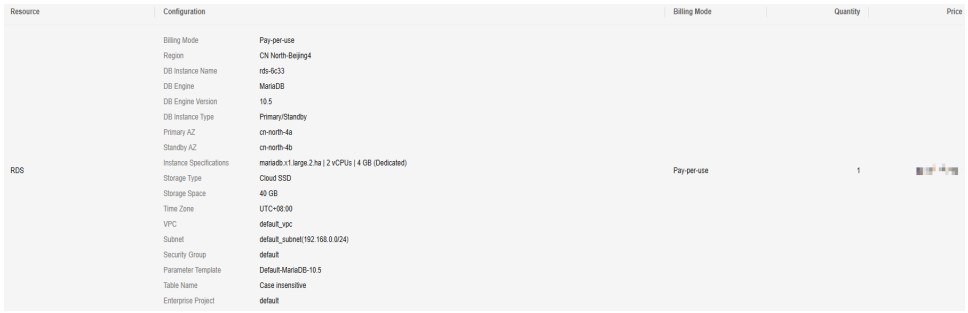

Figure 2-5 Confirming specifications (yearly/monthly)

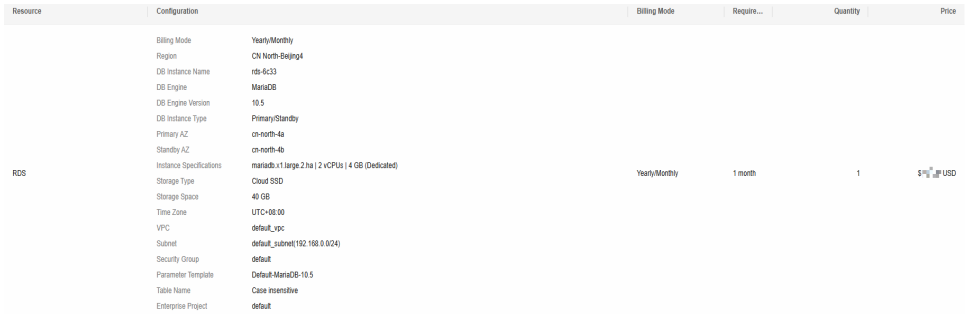

- For pay-per-use DB instances, if you do not need to modify your settings, click **Submit**.
- For yearly/monthly DB instances, if you do not need to modify your settings, click **Pay Now**.
- If you need to modify your settings, click **Previous**.

Step 7 Select a payment method and complete the payment.

NOTE

This operation applies only to the yearly/monthly billing mode.

Step 8 To view and manage your DB instance, go to the **Instances** page.

- When your DB instance is being created, the status is **Creating**. The status changes to **Available** after the instance is created. To view the detailed progress and result of the creation, go to the **Task Center** page.
- The automated backup policy is enabled by default. You can change it after the DB instance is created. An automated full backup is immediately triggered once your DB instance is created.
- After a DB instance is created, you can enter a description for it.
- The default database port is **3306**. You can change it after a DB instance is created.

  NOTE

  You are advised to change the default database port in a timely manner. For details, see **Changing a Database Port**.

----End

3 Instance Connection

3.1 Overview

You can connect to an RDS for MariaDB instance through a command-line interface (CLI), Data Admin Service (DAS), or using Java database connectivity (JDBC).

Table 3-1 Connection methods

Connection Method	Description
Connecting to an RDS for MariaDB Instance Through the MySQL CLI Client	In Linux, you need to install a MariaDB client on your device and connect to the instance through the MySQL CLI over a private or public network. • A floating IP address is provided by default. When your applications are deployed on an ECS that is in the same region and VPC as the RDS for MariaDB instance, you are advised to use a floating IP address to connect to the instance through the ECS. • If you cannot access your RDS for MariaDB instance through a floating IP address, bind an EIP to the instance and connect to the instance through the EIP.
Connecting to an RDS for MariaDB Instance Through JDBC	If you are connecting to an instance through JDBC, the SSL certificate is optional. For security reasons, you are advised to download the SSL certificate to encrypt the connection. SSL is disabled by default for RDS for MariaDB instances. You can enable SSL by referring to **Configuring an SSL Connection**. SSL encrypts connections to databases but it increases the connection response time and CPU usage. Therefore, you are advised not to enable SSL.

Connection Method	Description
Connecting to an RDS for MariaDB Instance Through DAS (Recommended)	DAS enables you to manage databases on a web-based console and provides you with database development, O&M, and intelligent diagnosis to make it easy to use and maintain your databases. The permissions required for connecting to DB instances through DAS are enabled by default.

3.2 Connecting to an RDS for MariaDB Instance Through DAS (Recommended)

Scenarios

Data Admin Service (DAS) enables you to connect to and manage DB instances with ease on a web-based console. The permission required for connecting to DB instances through DAS has been enabled for you by default. Using DAS to connect to your DB instance is recommended, which is more secure and convenient.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the DB instance and click **Log In** in the **Operation** column.

Figure 3-1 Login page

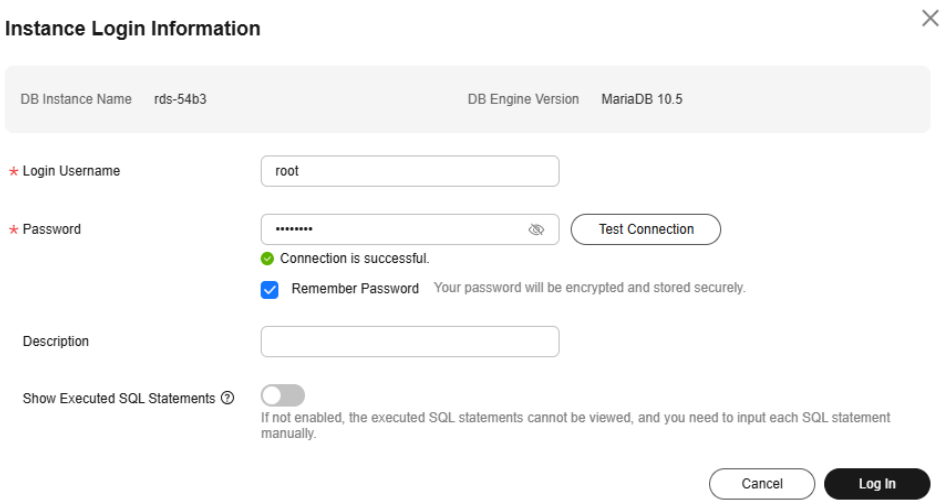

Step 5 Enter the database username and password and click **Test Connection**.

Step 6 After the connection test is successful, click **Log In**.

For details about how to manage databases using DAS, see **RDS for MariaDB Database Management** in the *Data Admin Service User Guide*.

----End

3.3 Connecting to an RDS for MariaDB Instance Through the MySQL CLI Client

3.3.1 Using MySQL CLI to Connect to an Instance Through a Private Network

If your applications are deployed on an ECS that is in the same region and VPC as your RDS for MariaDB instance, you are advised to connect to the DB instance through a floating IP address using the ECS.

This section describes how to connect a Linux ECS to a DB instance with SSL enabled or disabled through a floating IP address. SSL encrypts connections to the DB instance, making in-transit data more secure.

Prerequisites

1. You have logged in to an ECS.
   - For details on how to create and log in to an ECS, see **Purchasing an ECS** and **Logging In to an ECS** in *Elastic Cloud Server Getting Started*.
   - To connect to a DB instance through an ECS, you must ensure that:
     - The ECS and DB instance are in the same VPC.
     - The ECS is allowed by the security group to access the DB instance.
       - If the security group with which the DB instance is associated is the default security group, you do not need to configure security group rules.
       - If the security group with which the DB instance is associated is not the default security group, check whether the security group rules allow the ECS to connect to the DB instance.

         If the security group rules allow the access from the ECS, you can connect to the DB instance through the ECS.

         If the security group rules do not allow the access from the ECS, you need to add a security group rule, allowing the ECS to access the DB instance.
2. You have installed a database client to connect to DB instances.

   In Linux, install a **MariaDB client** on a device that can access RDS. It is recommended that you download a MariaDB client running a version later than that of the DB instance.

Connecting to a DB Instance Using a CLI (SSL Connection)

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target instance name to go to the **Overview** page.

Step 5 Under **SSL**, click **Enable** and then click **OK**.

Step 6 Click **Download** under **SSL** to download **Certificate Download.zip**, and extract the root certificate **ca.pem** and bundle **ca-bundle.pem** from the package.

Step 7 Import the root certificate **ca.pem** to the Windows or Linux ECS.

Step 8 Connect to the RDS for MariaDB instance. In Linux, for example, run the following command:

mysql -h <*host*> **-P** <*port*> **-u** <*userName*> **-p --ssl-ca=**<*caName*>

Example:

mysql -h 172.16.0.31 -P 3306 -u root -p --ssl-ca=ca.pem

Table 3-2 Parameter description

Parameter	Description
<*host*>	Floating IP address. To obtain this parameter value, go to the **Overview** page of the target instance and find **Floating IP Address**.
<*port*>	Database port. By default, the value is **3306**. To obtain this parameter value, go to the **Overview** page of the target instance and find **Database Port**.
<*userName*>	Username of the database account used for logging in to the DB instance. The default value is **root**.
<*caName*>	Name of the CA certificate. The certificate should be stored in the directory where the command is executed.

Step 9 Enter the password of the database account if the following information is displayed:

```
Enter password:
```


Figure 3-2 Connection example

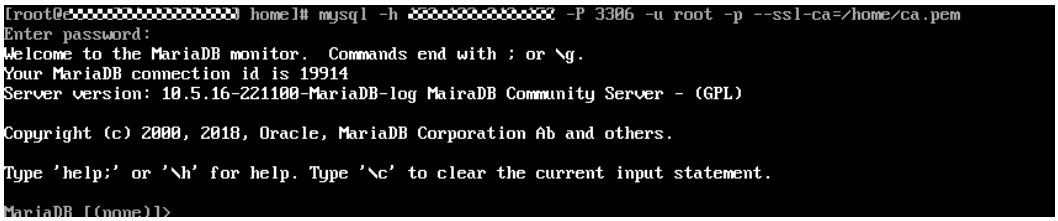

----End

Connecting to a DB Instance Using a CLI (Non-SSL Connection)

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target instance name to go to the **Overview** page.

Step 5 Under **SSL**, click **Disable** and then click **OK**.

Step 6 Connect to the RDS for MariaDB instance. In Linux, for example, run the following command:

mysql -h <*host*> **-P** <*port*> **-u** <*userName*> **-p**

Example:

mysql -h 172.16.0.31 -P 3306 -u root -p

Table 3-3 Parameter description

Parameter	Description
<*host*>	Floating IP address. To obtain this parameter value, go to the **Overview** page of the target instance and find **Floating IP Address**.
<*port*>	Database port. By default, the value is **3306**. To obtain this parameter value, go to the **Overview** page of the target instance and find **Database Port**.
<*userName*>	Username of the database account used for logging in to the DB instance. The default value is **root**.

Step 7 Enter the password of the database account if the following information is displayed:

```
Enter password:
```


Figure 3-3 Non-SSL connection example

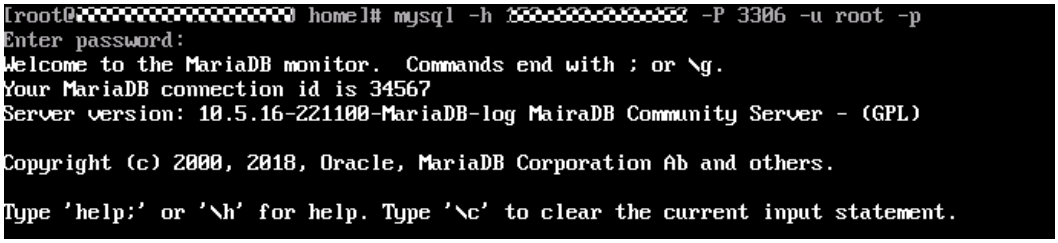

----End

3.3.2 Using MySQL CLI to Connect to an Instance Through a Public Network

If you cannot access your DB instance through a floating IP address, bind an EIP to the DB instance and connect to it through the EIP.

This section describes how to connect a Linux ECS to a DB instance with SSL enabled or disabled through an EIP. SSL encrypts connections to the DB instance, making in-transit data more secure.

Prerequisites

1. You have bound an EIP to the target DB instance and configured security group rules.
   a. Bind an EIP to the target DB instance.
   b. Obtain the IP address of the ECS you use to connect to the DB instance.
   c. Configure security group rules.
      Add the IP address obtained in **1.b** and the instance port to the inbound rule of the security group.
   d. Run the **ping** command to ping the EIP bound in **1.a** to ensure that the EIP is accessible through the ECS.
2. You have installed a database client to connect to DB instances.
   In Linux, you need to install a **MariaDB client** on your device. It is recommended that you download a MariaDB client running a version later than that of the DB instance.

Connecting to a DB Instance Using a CLI (SSL Connection)

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target instance name to go to the **Overview** page.

Step 5 Under **SSL**, click **Enable** and then click **OK**.

Step 6 Click **Download** under **SSL** to download **Certificate Download.zip**, and extract the root certificate **ca.pem** and bundle **ca-bundle.pem** from the package.

Step 7 Import the root certificate **ca.pem** to the Windows or Linux ECS.

Step 8 Connect to the RDS for MariaDB instance. In Linux, for example, run the following command:

mysql -h <*host*> **-P** <*port*> **-u** <*userName*> **-p --ssl-ca=**<*caName*>

Example:

mysql -h 172.16.0.31 -P 3306-u root -p --ssl-ca=ca.pem

Table 3-4 Parameter description

Parameter	Description
<*host*>	EIP of the DB instance to be connected.
<*port*>	Port of the DB instance to be connected.
<*userName*>	Username of the database account used for logging in to the DB instance. The default value is **root**.
<*caName*>	Name of the CA certificate. The certificate should be stored in the directory where the command is executed.

Step 9 Enter the password of the database account if the following information is displayed:

```
Enter password:
```

Figure 3-4 Connection example

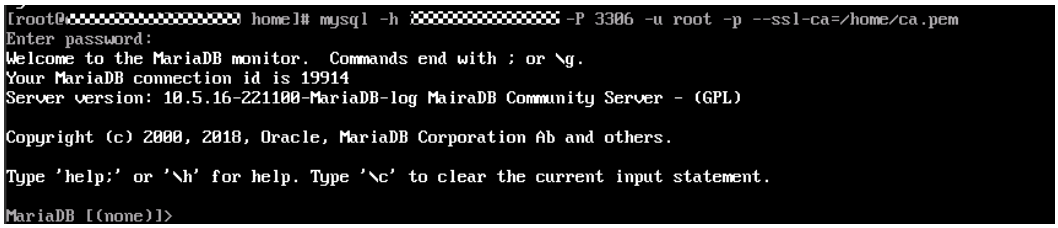

NOTE

If the connection fails, ensure that all **prerequisites** are correctly configured and try again.

----End

Connecting to a DB Instance Using a CLI (Non-SSL Connection)

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target instance name to go to the **Overview** page.

Step 5 Under **SSL**, click **Disable** and then click **OK**.

Step 6 Connect to the RDS DB instance. In Linux, for example, run the following command:

mysql -h <*host*> **-P** <*port*> **-u** <*userName*> **-p**

Example:

mysql -h 172.16.0.31 -P 3306 -u root -p

Table 3-5 Parameter description

Parameter	Description
<*host*>	EIP of the DB instance to be connected.
<*port*>	Port of the DB instance to be connected.
<*userName*>	Username of the database account. The default administrator is **root**.

Step 7 Enter the password of the database account if the following information is displayed:

```
Enter password:
```

Figure 3-5 Non-SSL connection example

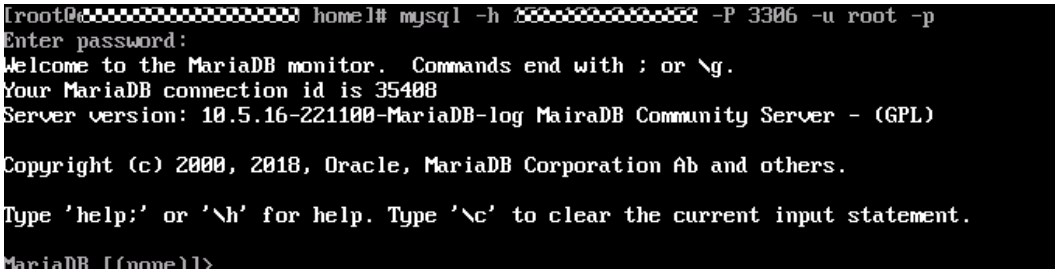

NOTE

If the connection fails, ensure that preparations have been correctly made in **Prerequisites** and try again.

----End

3.4 Connecting to an RDS for MariaDB Instance Through JDBC

If you are connecting to an instance through JDBC, an SSL certificate is optional, but using an SSL certificate can improve the security of your data. SSL is disabled

by default for RDS for MariaDB instances. It encrypts connections to databases but increases the connection response time and CPU usage. For this reason, you are advised not to enable SSL.

Prerequisites

You are familiar with:

- Computer basics.
- Java.
- JDBC.

Connection with the SSL Certificate

NOTE

Download the SSL certificate and verify it before connecting to your instance.

Step 1 Download the CA certificate or certificate bundle. On the **Instances** page, click the instance name to go to the **Overview** page. Under **SSL**, click **Download**.

Step 2 Use keytool to generate a truststore file using the CA certificate.

```
<keytool_installation_path> ./keytool.exe -importcert -alias <MariaDBCACert> -file <ca.pem> -keystore <truststore_file> -storepass <password>
```

Table 3-6 Parameter description

Parameter	Description
<keytool installation path>	Bin directory in the JDK or JRE installation path, for example, **C:\Program Files (x86)\Java\jdk11.0.7\bin**.
<MariaDBCACert>	Name of the truststore file. Set it to a name specific to the service for future identification.
<ca.pem>	Name of the CA certificate downloaded and decompressed in **Step 1**, for example, ca.pem.
<truststore_file>	Path for storing the truststore file.
<password>	Password of the truststore file.

Code example (using keytool in the JDK installation path to generate the truststore file):

```
Owner:  CN=MySQL_Server_5.7.17_Auto_Generated_CA_Certificate
Issuer: CN=MySQL_Server_5.7.17_Auto_Generated_CA_Certificate
Serial number: 1
Valid from: Thu Feb 16 11:42:43 EST 2017 until: Sun Feb 14 11:42:43 EST 2027
Certificate fingerprints:
    MD5: 18:87:97:37:EA:CB:0B:5A:24:AB:27:76:45:A4:78:C1
    SHA1: 2B:0D:D9:69:2C:99:BF:1E:2A:25:4E:8D:2D:38:B8:70:66:47:FA:ED
    SHA256:C3:29:67:1B:E5:37:06:F7:A9:93:DF:C7:B3:27:5E:09:C7:FD:EE:2D:18:86:F4:9C:40:D8:26:CB:DA:95:
A0:24
    Signature algorithm name: SHA256withRSA Subject Public Key Algorithm: 2048-bit RSA key
    Version: 1
```



```
Trust this certificate? [no]: y
Certificate was added to keystore
```

Step 3 Connect to your RDS for MariaDB instance through JDBC.

```
jdbc:mysql://<instance_ip>:<instance_port>/<database_name>?param1=value1&param2=value2
```

Table 3-7 Parameter description

Parameter	Description
<instance_ip>	IP address of the DB instance. **NOTE** • If you are accessing the DB instance through an ECS, *instance_ip* is the floating IP address of the instance. You can obtain this IP address on the **Connectivity & Security** page. • If you are accessing the DB instance through a public network, *instance_ip* is the EIP that has been bound to the instance. You can obtain this IP address on the **Connectivity & Security** page.
<instance_port>	Database port of the DB instance. The default port is **3306**. **NOTE** You can obtain this port number on the **Connectivity & Security** page.
<database_name>	Database name used for connecting to the DB instance. The default value is **MariaDB**.
<param1>	**requireSSL**, indicating whether the server supports SSL. Its value can be either of the following: • **true**: The server supports SSL. • **false**: The server does not support SSL. **NOTE** For details about the relationship between **requireSSL** and **sslmode**, see **Table 3-8**.
<param2>	**useSSL**, indicating whether the client uses SSL to connect to the server. Its value can be either of the following: • **true**: The client uses SSL to connect to the server. • **false**: The client does not use SSL to connect to the server. **NOTE** For details about the relationship between **useSSL** and **sslmode**, see **Table 3-8**.
<param3>	**verifyServerCertificate**, indicating whether the client verifies the server certificate. Its value can be either of the following: • **true**: The client verifies the server certificate. • **false**: The client does not verify the server certificate. **NOTE** For details about the relationship between **verifyServerCertificate** and **sslmode**, see **Table 3-8**.

Parameter	Description
<param4>	**trustCertificateKeyStoreUrl**. Its value is **file:<truststore_file>**. **<truststore_file>** is the path for storing the truststore file set in **Step 2**.
<param5>	**trustCertificateKeyStorePassword**. Its value is the password of the truststore file set in **Step 2**.

Table 3-8 Relationship between connection parameters and sslmode

useSSL	requireSSL	verifyServerCer-tificate	sslMode
false	N/A	N/A	DISABLED
true	false	false	PREFERRED
true	true	false	REQUIRED
true	N/A	true	VERIFY_CA

Code example (Java code for connecting to an RDS for MariaDB instance):

```
import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.ResultSet;
import java.sql.Statement;
import java.sql.SQLException;

// There will be security risks if the username and password used for authentication are directly written into code. Store them in ciphertext in the configuration file or environment variables.
// In this example, the username and password are stored in the environment variables EXAMPLE_USERNAME_ENV and EXAMPLE_PASSWORD_ENV as needed.
public class JDBCTest {
    String USER = System.getenv("EXAMPLE_USERNAME_ENV");
    String PASS = System.getenv("EXAMPLE_PASSWORD_ENV");

    public static void main(String[] args) {
        Connection conn = null;
        Statement stmt = null;
      // Set the required parameters in the URL based on the site requirements.
        String url = "jdbc:mysql://<instance_ip>:<instance_port>/<database_name>?
param1=value1&param2=value2";

        try {

            Class.forName("com.MariaDB.cj.jdbc.Driver");
            conn = DriverManager.getConnection(url, USER, PASS);

            stmt = conn.createStatement();
            String sql = "show status like 'ssl%'";
            ResultSet rs = stmt.executeQuery(sql);

            int columns = rs.getMetaData().getColumnCount();
            for (int i = 1; i <= columns; i++) {
                System.out.print(rs.getMetaData().getColumnName(i));
```



```
            System.out.print("\t");
        }

        while (rs.next()) {
            System.out.println();
            for (int i = 1; i <= columns; i++) {
                System.out.print(rs.getObject(i));
                System.out.print("\t");
            }
        }
        rs.close();
        stmt.close();
        conn.close();
    } catch (SQLException se) {
        se.printStackTrace();
    } catch (Exception e) {
        e.printStackTrace();
    } finally {
        // release resource ....
    }
  }
}
```

----End

Connection Without the SSL Certificate

NOTE

You do not need to download the SSL certificate because certificate verification on the server is not required.

Connect to the RDS for MariaDB instance through JDBC.

jdbc:mysql://*<instance_ip>*:*<instance_port>*/*<database_name>*?**useSSL**=false

Table 3-9 Parameter description

Parameter	Description
<instance_ip>	IP address of the DB instance. **NOTE** • If you are accessing the DB instance through an ECS, *instance_ip* is the floating IP address of the instance. You can obtain this IP address on the **Connectivity & Security** page. • If you are accessing the DB instance through a public network, *instance_ip* is the EIP that has been bound to the instance. You can obtain this IP address on the **Connectivity & Security** page.
<instance_port>	Database port of the DB instance. The default port is **3306**. **NOTE** You can obtain this port number on the **Connectivity & Security** page.
<database_name>	Database name used for connecting to the DB instance. The default value is **MariaDB**.

Code example (Java code for connecting to an RDS for MariaDB instance):

```
import java.sql.Connection;
import java.sql.DriverManager;
```



```
import java.sql.ResultSet;
import java.sql.Statement;
// There will be security risks if the username and password used for authentication are directly written into code. Store them in ciphertext in the configuration file or environment variables.
// In this example, the username and password are stored in the environment variables. Before running this example, set environment variables EXAMPLE_USERNAME_ENV and EXAMPLE_PASSWORD_ENV as needed.

public class MyConnTest {
    final public static void main(String[] args) {
        Connection conn = null;
                // Set the required parameters in the URL based on the site requirements.
                String url = "jdbc:mysql://<instance_ip>:<instance_port>/<database_name>?
param1=value1&param2=value2";
                String USER = System.getenv("EXAMPLE_USERNAME_ENV");
                String PASS = System.getenv("EXAMPLE_PASSWORD_ENV");
        try {
            Class.forName("com.MariaDB.jdbc.Driver");
            conn = DriverManager.getConnection(url,USER,PASS);
            System.out.println("Database connected");

            Statement stmt = conn.createStatement();
            ResultSet rs = stmt.executeQuery("SELECT * FROM mytable WHERE columnfoo = 500");
            while (rs.next()) {
                System.out.println(rs.getString(1));
            }
            rs.close();
            stmt.close();
            conn.close();
        } catch (Exception e) {
            e.printStackTrace();
            System.out.println("Test failed");
        } finally {
            // release resource ....
        }
    }
}
```

Related Issues

- Symptom

  When you use JDK 8.0 or a later version to connect to an RDS for MariaDB instance with an SSL certificate downloaded, an error similar to the following is reported:

```
javax.net.ssl.SSLHandshakeException: No appropriate protocol (protocol is disabled or
cipher suites are inappropriate)
    at sun.security.ssl.HandshakeContext.<init>(HandshakeContext.java:171) ~[na:1.8.0_292]
    at sun.security.ssl.ClientHandshakeContext.<init>(ClientHandshakeContext.java:98) ~
[na:1.8.0_292]
    at sun.security.ssl.TransportContext.kickstart(TransportContext.java:220) ~
[na:1.8.0_292]
    at sun.security.ssl.SSLSocketImpl.startHandshake(SSLSocketImpl.java:428) ~
[na:1.8.0_292]
    at
com.MariaDB.cj.protocol.ExportControlled.performTlsHandshake(ExportControlled.java:316) ~
[MariaDB-connector-java-8.0.17.jar:8.0.17]
    at
com.MariaDB.cj.protocol.StandardSocketFactory.performTlsHandshake(StandardSocketFactory.java
:188) ~[MariaDB-connector-java8.0.17.jar:8.0.17]
    at
com.MariaDB.cj.protocol.a.NativeSocketConnection.performTlsHandshake(NativeSocketConnection.
java:99) ~[MariaDB-connector-java8.0.17.jar:8.0.17]
    at
com.MariaDB.cj.protocol.a.NativeProtocol.negotiateSSLConnection(NativeProtocol.java:331) ~
[MariaDB-connector-java8.0.17.jar:8.0.17]
... 68 common frames omitted
```


- Solution

  Specify the corresponding parameter values in the code link of **Step 3** based on the JAR package used by the client. Example:

  - MariaDB-connector-java-5.1.*xx*.jar
    In the database connection URL jdbc:**mysql**://*<instance_ip>*:*<instance_port>*/*<database_name>*?param1=value1&param2=value2, replace **param1=value1** with **enabledTLSProtocols=TLSv1.2**.
  - MariaDB-connector-java-8.0.*xx*.jar
    In the database connection URL jdbc:**mysql**://*<instance_ip>*:*<instance_port>*/*<database_name>*?param1=value1&param2=value2, replace **param1=value1** with **tlsVersions=TLSv1.2**.

3.5 Connection Management

3.5.1 Changing a Floating IP Address

You can change floating IP addresses after migrating on-premises databases or other cloud databases to RDS.

Constraints

Changing a floating IP address will interrupt the database connection. You are advised to change a floating IP address during off-peak hours.

Procedure

When you buy a DB instance, select a VPC and subnet on the **Buy DB Instance** page. Then, a floating IP address will be automatically assigned to your instance.

After the instance is created, you can change its floating IP address.

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 Under **Floating IP Address**, click **Configure**.

Alternatively, in the navigation pane on the left, choose **Connectivity & Security**. In the **Connection Information** area, click **Change** next to the **Floating IP Address** field.

Figure 3-6 Floating IP address

Step 6 In the displayed dialog box, check the number of in-use IP addresses. If the in-use IP addresses are less than 254, there are unused floating IP addresses.

Figure 3-7 Changing a floating IP address

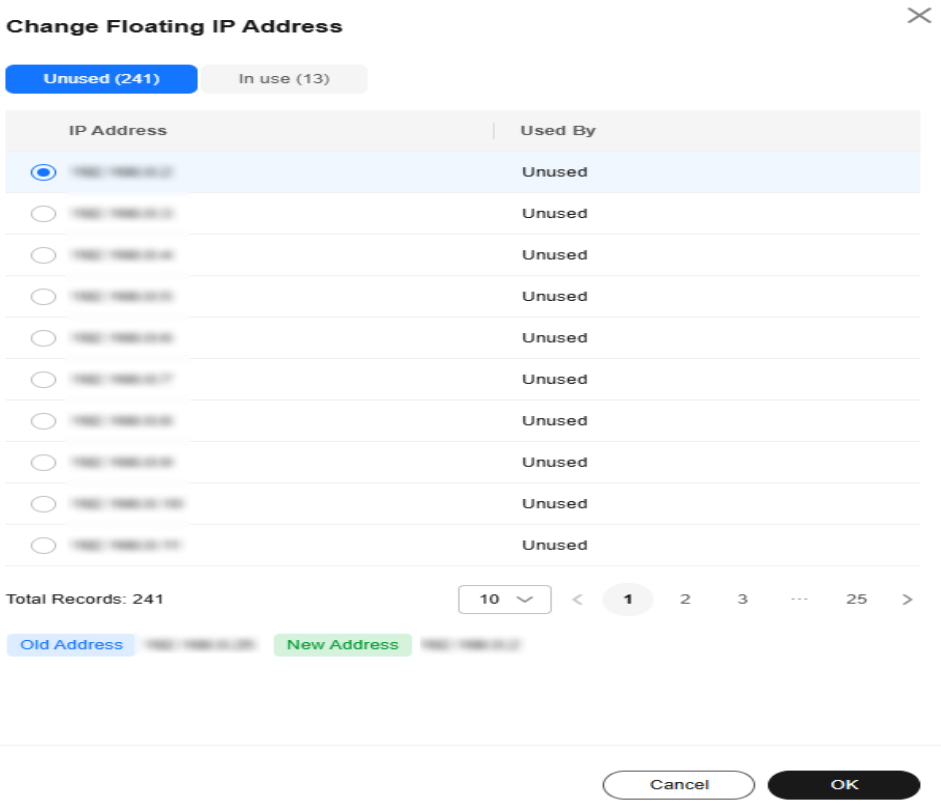

Step 7 Enter an available IP address and click **OK**.

An in-use IP address cannot be used as the new floating IP address of the DB instance.

Step 8 If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account and cloud product. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

----End

3.5.2 Binding and Unbinding an EIP

Scenarios

You can bind an EIP to a DB instance for public accessibility, and you can unbind the EIP from the DB instance later if needed.

NOTICE

To ensure that the DB instance is accessible, the security group associated with the instance must allow access over the database port. For example, if the database port is 8635, ensure that the security group allows access over the 8635 port.

Precautions

- To bind an EIP to a DB instance, **submit a service ticket** to apply for the required permissions.
- You need to configure security groups and enable specific IP addresses and ports to access the target DB instance. Before accessing the DB instance, add an individual IP address or an IP address range that will access the DB instance to the inbound rule. For details, see **Configuring Security Group Rules**.

Prerequisites

- You can bind an EIP to a primary DB instance or a read replica only.
- If a DB instance has already an EIP bound, you must unbind the EIP from the DB instance first before binding a new one to it.

Binding an EIP

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Connectivity & Security**. In the **Connection Information** area, click **Bind** next to the **EIP** field.

Alternatively, in the **Connection Topology** area, click **Public Connection** and then **Bind** in the connection topology.

Step 6 In the displayed dialog box, all unbound EIPs are listed. Select the EIP to be bound and click **Yes**.

Step 7 On the **Connectivity & Security** page, view the EIP that has been bound to the DB instance.

You can also view the progress and result of binding an EIP to a DB instance on the **Task Center** page.

To unbind the EIP from the DB instance, see **Unbinding an EIP**.

----End

Unbinding an EIP

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the DB instance that has an EIP bound to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Connectivity & Security**. In the **Connection Information** area, click **Unbind** next to the **EIP** field. In the displayed dialog box, click **Yes**.

Alternatively, in the **Connection Topology** area, click **Public Connection** and then **Unbind** in the connection topology. In the displayed dialog box, click **Yes**.

Step 6 (**Optional**) If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

Step 7 On the **Connectivity & Security** page, check the results.

You can also view the progress or the results of unbinding an EIP from a DB instance on the **Task Center** page.

To bind an EIP to the DB instance again, see **Binding an EIP**.

----End

3.5.3 Changing a Database Port

Scenarios

This section describes how to change the database port of a primary DB instance or a read replica. For primary/standby DB instances, changing the database port of the primary DB instance will cause the database port of the standby DB instance to also be changed.

If specific security group rules have been configured for a DB instance, you need to change the inbound rules of the security group to which the DB instance belongs after changing the database port.

Constraints

When the database port of a DB instance is being changed, you cannot:

- Bind an EIP to the DB instance.
- Delete the DB instance.

- Create a backup for the DB instance.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target DB instance or click first and then click the target read replica.

Step 5 On the **Overview** page, find **Database Port** and click **Configure** under it.

Alternatively, choose **Connectivity & Security** in the navigation pane on the left. On the displayed page, click **Change** next to the **Database Port** field.

NOTE

RDS for MariaDB instances can use database ports 1024 to 65535, excluding 12017 and 33071, which are reserved for RDS system use.

Step 6 In the displayed dialog box, enter a new port and click **Yes**.

If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

- If you change the database port of the primary DB instance, that of the standby DB instance will also be changed and both DB instances will be rebooted.
- If you change the database port of a read replica, the change will not affect other DB instances. Only the read replica will be rebooted.
- This process takes 1 to 5 minutes.

Step 7 Check the results on the **Overview** page.

----End

3.5.4 Downloading a Certificate

RDS for MariaDB allows you to download a certificate.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target instance name to go to the **Overview** page.

Step 5 Under **SSL**, click **Download** to download the root certificate and certificate bundle.

Alternatively, choose **Connectivity & Security** from the navigation pane. In the **Connection Information** area, click ⬇ next to the **SSL** field to download the root certificate and certificate bundle.

----End

3.5.5 Configuring Security Group Rules

Scenarios

A security group is a collection of access control rules for ECSs and RDS DB instances that have the same security protection requirements and are mutually trusted in a VPC.

To ensure database security and reliability, you need to configure security group rules to allow only specific IP addresses and ports to access your DB instance.

- When you attempt to connect to an RDS DB instance through an EIP, you need to configure an **inbound rule** for the security group associated with the RDS DB instance.
- When you attempt to connect to an RDS DB instance through a private network, check whether the ECS and DB instance are in the same security group.
  - If the ECS and RDS DB instance are in the same security group, they can communicate with each other by default. No security group rules need to be configured. **Connect to the RDS for MariaDB instance through a private network**.
  - If they are in different security groups, configure security group rules for them, separately.
    - RDS instance: Configure an **inbound rule** for the security group with which the RDS instance is associated.
    - ECS: The default security group rule allows all outgoing data packets. In this case, you do not need to configure a security group rule for the ECS. If not all outbound traffic is allowed in the security group, you need to configure an **outbound rule** for the ECS.

For details about the requirements of security group rules, see **Adding a Security Group Rule** in the *Virtual Private Cloud User Guide*.

Constraints

The default security group rule allows all outgoing data packets. ECSs and RDS DB instances can access each other if they are in the same security group. After a

security group is created, you can configure security group rules to control access from and to the DB instances in the security group.

- By default, you can create a maximum of 100 security groups in your cloud account.
- By default, you can add up to 50 security group rules to a security group.
- One RDS DB instance can be associated with multiple security groups, and one security group can be associated with multiple RDS DB instances.
- Too many security group rules will increase the first packet latency. You are advised to create no more than 50 rules for a security group.
- To access a DB instance from resources outside the security group, you need to configure an **inbound rule** for the security group associated with the DB instance.

NOTE

To ensure data and instance security, use permissions properly. You are advised to use the minimum access permission, change the default database port **3306**, and set the accessible IP address to the remote server's address or the remote server's minimum subnet address to control the access scope of the remote server.

The default value of **Source** is **0.0.0.0/0**, indicating that RDS DB instances in the security group can be accessed from any IP address.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target DB instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **Connectivity & Security**. In the **Security Group Rules** area, click the security group name to view the security group rules.

Figure 3-8 Security Group Rules

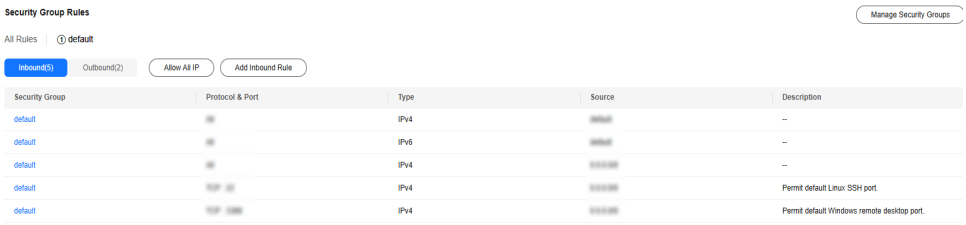

Step 6 Click **Add Inbound Rule** or **Allow All IP** to configure security group rules.

To add more inbound rules, click .

NOTE

Allow All IP allows all IP addresses to access RDS DB instances in the security group, which poses high security risks. Exercise caution when performing this operation.

Figure 3-9 Adding an inbound rule

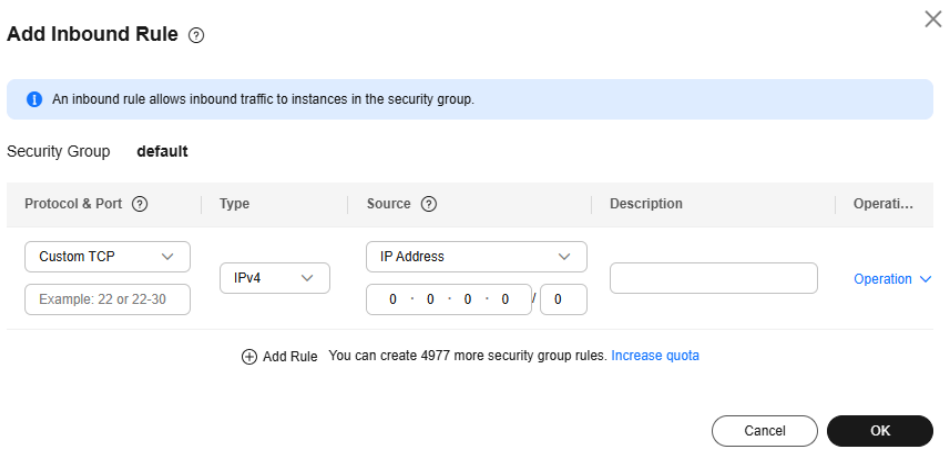

Table 3-10 Inbound rule parameter description

Parameter	Description	Example Value
Protocol & Port	**Protocol**: network protocol. Available options: **All ports**, **Custom TCP**, **Custom UDP**, **ICMP**, and **GRE**.	Custom TCP
	Port: the port over which the traffic can reach your DB instance. RDS for MariaDB instances can use database ports 1024 to 65535, excluding 12017 and 33071, which are reserved for RDS system use.	3306
Type	Supported source IP address type. Its value can be: • IPv4 • IPv6	IPv4

Parameter	Description	Example Value
Source	The source in an inbound rule is used to match the IP address or address range of an external request. Examples: • Single IP address: 192.168.10.10/32 (IPv4 address) • IP address segment: 192.168.1.0/24 (IPv4 address segment) • All IP addresses: 0.0.0.0/0 (any IPv4 address) • Security group: sg-abc • IP address group: ipGroup-test	0.0.0.0/0
Description	Supplementary information about the security group rule. This parameter is optional. The description can contain a maximum of 255 characters and cannot contain angle brackets (<) or (>).	N/A

Step 7 Click **OK**.

----End

4 Database Usage

4.1 Suggestions on Using RDS for MariaDB

4.1.1 Instance Usage Suggestions

DB Instances

DB Instance Types

- Primary/Standby
  - A primary/standby pair provides an HA architecture. It is suitable for production databases of large and medium enterprises in the Internet, Internet of Things (IoT), retail e-commerce sales, logistics, gaming, and other sectors.
  - When a primary instance is being created, a standby instance is provisioned along with it to provide data redundancy. The standby instance is invisible to you after being created.
  - If a failover occurs due to a primary instance failure, your database client will be disconnected for a short period of time. The client needs to be able to reconnect to the instance.
- Single
  - A single-node architecture is more cost-effective than primary/standby pairs.
  - It is only recommended for development and testing of microsites, and small and medium enterprises, or for learning about RDS.
  - If a fault occurs on a single instance, the instance cannot recover in a timely manner.
- Read replica

  RDS for MariaDB supports single read replicas.

  NOTE

  If the replication between a read replica and the DB instance is abnormal, it can take a long time to rebuild and restore the read replica (depending on the data volume).

Instance Classes

- Dedicated

  The instance has dedicated CPU and memory resources to ensure stable performance. The performance of a dedicated instance is never affected by other instances on the same physical machine. This instance class is good when performance stability is important.

- General-purpose

  CPU resources are shared with other general-purpose DB instances on the same physical machine. CPU usage is maximized through resource overcommitment. This instance class is a cost-effective option and suitable for scenarios where performance stability is not critical.

Database Connection

- Configure RDS for MariaDB parameters for your workloads.
- Keep an appropriate number of active connections.
- Periodically release persistent connections because maintaining them may generate a large cache and use up memory.

Reliability and Availability

- Select primary/standby DB instances for production databases.
- Deploy primary and standby instances in different AZs.
- Create read replicas and enable read/write splitting for workloads involving frequent read/write operations.
- Change instance classes during off-peak hours.
- Select an instance class and storage space appropriate to your workloads.
- After scaling up your primary DB instance, scale up its read replicas in a timely manner to prevent service exceptions caused by insufficient storage of read replicas.

Backup and Restoration

- Perform manual backups during off-peak hours and change the backup time window (default setting: 01:00-02:00 (GMT+08:00)) for automated backup as required.
- Set the backup cycle to **All** for DB instances that process many write requests every day.
- Configure a backup retention period suited to your service demands. The default value is 7 days.
- If a DB instance is deleted, its automated full backups and binlog backups are also deleted. Perform manual backup for all data before deleting a DB instance.
- Configure a custom recycling policy to ensure that any instances that are deleted by mistake can be rebuilt.

Routine O&M

- Periodically check slow query logs and error logs to identify problems in advance.

- Periodically check the resource usage of DB instances. If the resources are insufficient, scale up the resources in a timely manner.
- Monitor instance metrics. If any metric is beyond its expected range, address related issues as soon as possible.
- Run the **SELECT** statement before deleting or modifying a record.

Security

- Prevent your database from being accessed from the Internet. If you want to allow the access from the Internet, bind an EIP to your DB instance and configure security group rules.
- Use SSL to connect to your DB instance.

4.1.2 Database Usage Suggestions

Database Naming

- The names of database objects like databases, tables, and columns should be in lowercase. Different words in the name are separated with underscores (_).
- Reserved words and keywords cannot be used to name database objects in RDS for MariaDB.
- Each database object name must be explainable and contain a maximum of 32 characters.
- Each temporary table in databases is prefixed with **tmp** and suffixed with a date.
- Each backup table in databases is prefixed with **bak** and suffixed with a date.
- All columns storing the same data in different databases or tables must have the same name and be of the same type.

Database Design

- All tables use the InnoDB storage engine unless otherwise specified. InnoDB supports transactions and row locks. It delivers excellent performance, making it easy to restore data.
- Databases and tables all use the UTF8 character set to avoid characters getting garbled by character set conversion.
- All tables and fields require comments that can be added using the COMMENT clause to maintain the data dictionary from the beginning of the design.
- To avoid cross-partition queries, RDS for MariaDB partitioned tables are not recommended. Cross-partition queries will decrease the query efficiency. A partitioned table is logically a single table, but the data is actually stored in multiple different files. If you use partitioned tables for storage, store files from different partitions on different disk arrays.
- Do not create too many columns in one table. Store cold and warm data separately to reduce the width of a table. In doing so, more rows of data can be stored in each memory page, decreasing disk I/O and making more efficient use of the cache.
- Columns that are frequently used together should be in the same table to avoid JOIN operations.

- Do not create reserved fields in a table. Otherwise, modifying the column type will lock the table, which has a greater impact than adding a field.
- Do not store binary data such as images and files in databases.

Field Design

- Select a small data type for each column as much as possible. Numeric data is preferred, followed by dates or binary data, and the least preferred is characters. The larger the column data type, the more the space required for creating indexes. As a result, there are fewer indexes on a page and more I/O operations required, so database performance deteriorates.
- If the integer type is used as the database field type, select the shortest column type. If the value is a non-negative number, it must be the unsigned type.
- Ensure that each column has the NOT NULL attribute.
- Do not use the ENUM type. Instead, use the TINYINT type. Change ENUM values using ALTER. The ORDER BY operations on ENUM values are inefficient and require extra operations.

  If you have specified that ENUM values cannot be numeric, other data types (such as CHAR) can be used.
- If the numeric data type is required, use DECIMAL instead of FLOAT or DOUBLE. FLOAT and DOUBLE data cannot be stored precisely, and value comparison results may be incorrect.
- When you want to record a date or specific time, use the DATETIME or TIMESTAMP type instead of the string type.
- Store IP addresses using the INT UNSIGNED type. You can convert IP addresses into numeric data using function inet_aton or inet_ntoa.
- The VARCHAR data should be as short as possible. Although the VARCHAR data varies in length dynamically on disks, it occupies the maximum length in memory.
- Use VARBINARY to store variable-length character strings that are case-sensitive. VARBINARY is case-sensitive by default and quick to process because no character sets are involved.

Index Design

- Use no more than 5 indexes in a single table. Indexes speed up queries, but too many indexes may slow down writes. Inappropriate indexes sometimes reduce query efficiency.
- Do not create an independent index for each column in a table. A well-designed composite index is much more efficient than a separate index on each column.
- Create a primary key for each InnoDB table. Neither use a frequently-updated column as the primary key nor a multi-column primary key. Do not use the UUID, MD5, or character string column as the primary key. Use a column whose value can increment continuously as the primary key. So, the auto-increment ID column is recommended.
- Create an index on the following columns:
  - Columns specified in the WHERE clause of SELECT, UPDATE, or DELETE statements

  - Columns specified in ORDER BY, GROUP BY, or DISTINCT
  - Columns used to join multiple tables
- The index column order is as follows:
  - Put the column with the highest selectivity on the far left when creating a composite index. Selectivity = Different values in a column/Total rows in the column
  - Put the column with the smallest field length on the far left of the composite index. The smaller length a field has, the more data one page stores, and the better the I/O performance is.
  - Put the most frequently used column on the left of the composite index, so you can create fewer indexes.
- Avoid using redundant indexes, such as primary key (id), index (id), and unique index (id).
- Avoid using duplicate indexes, such as index(a,b,c), index(a,b), and index(a). Duplicate and redundant indexes may slow down queries because the RDS for MariaDB query optimizer does not know which index it should use.
- When creating an index on the VARCHAR field, specify the index length based on selectivity. Do not index the entire field.

  If an index with the length of 20 bytes is the string type, its selectivity can reach 90% or above. In this case, use **count(distinct left(column name, index length))/count(*)** to check index selectivity.
- Use covering indexes for frequent queries.

  A covering index is a special type of index where all required fields for a query are included in the index. The index itself contains columns specified in WHERE and GROUP BY clauses, but also column combinations queried in SELECT, without having to execute additional queries on InnoDB tables.
- Constraints on foreign keys are as follows:

  The character sets of the columns for which a foreign key relationship is established must be the same, or the character sets of the parent and child tables for which a foreign key relationship is established must be the same.

SQL Statement Development

- Use prepared statements to perform database operations in programs. Prepared statements can be executed multiple times in a program once they are written. They are more efficient than SQL statements.
- Avoid implicit conversions because they may cause indexes to become invalid.

  Do not perform function conversions or math calculations on columns in the WHERE clause. Otherwise, the index becomes invalid.
- Do not use double percent signs (%%) or place % before a query condition, or the index cannot be used.
- Do not use **SELECT *** for queries because using **SELECT ***:
  - Consumes more CPUs, IP addresses, and bandwidth.
  - Causes covering indexes to become unavailable.
  - Increases the impact of table structure changes on code.

- Do not use subqueries. Subqueries generate temporary tables that do not have any indexes. If there is a lot of data, the query efficiency is severely affected. Convert subqueries into join queries.
- Minimize the use of JOIN operations for more than 5 tables. Use the same data type for the fields that require JOIN operations.

  Each JOIN operation on a table occupies extra memory (controlled by **join_buffer_size**) and requires temporary table operations, affecting query efficiency.
- Reduce interactions with the same database as much as possible. The database is more suitable for processing batch operations.
- Replace OR clauses with IN clauses because IN clauses can effectively use indexes. Specify no more than 500 values for an IN clause.
- Do not perform reverse queries, for example, NOT IN and NOT LIKE.
- Do not use ORDER BY RAND() for random sorting.

  This operation loads all data that meets the conditions from the table to the memory for sorting, consuming more CPUs, I/O, and memory resources.

  Obtain a random value from the program and retrieve data from the involved database based on the value.
- If deduplication is not required, use UNION ALL instead of UNION.

  UNION ALL does not sort out result sets.
- Combine multiple operations and perform them in batches. The database is good for batch processing.

  This reduces interactions with the same database.
- If there are more than 1 million rows of write operations, perform them in multiple batches.

  A large number of batch writes may result in excessive primary/standby latency.
- If ORDER BY is used, use the order of indexes.
  - The last field of ORDER BY is a part of a composite index and is placed at the end of the composite index order.
  - Avoid file_sort to speed up queries.

  Correct example: in **WHERE a=? AND b=? ORDER BY c;**, index: **a_b_c**

  Wrong example: If an index supports range search, the index order cannot be used. For example, **WHERE a>10 ORDER BY b;**, index: **a_b** (sorting is not allowed)

4.2 Database Management

4.2.1 Creating a Database

Scenarios

After a DB instance is created, you can create databases on it.

Constraints

- Databases cannot be created for DB instances that are in the process of being restored.
- Database names must be unique.
- After a database is created on the RDS console, its name cannot be changed.

Creating a Database Through RDS

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Databases and Accounts** and then click the **Databases** tab.

Step 6 Click **Create Database**. In the displayed dialog box, enter a database name, select a character set, and authorize permissions for users. Then, click **OK**.

Figure 4-1 Creating a database

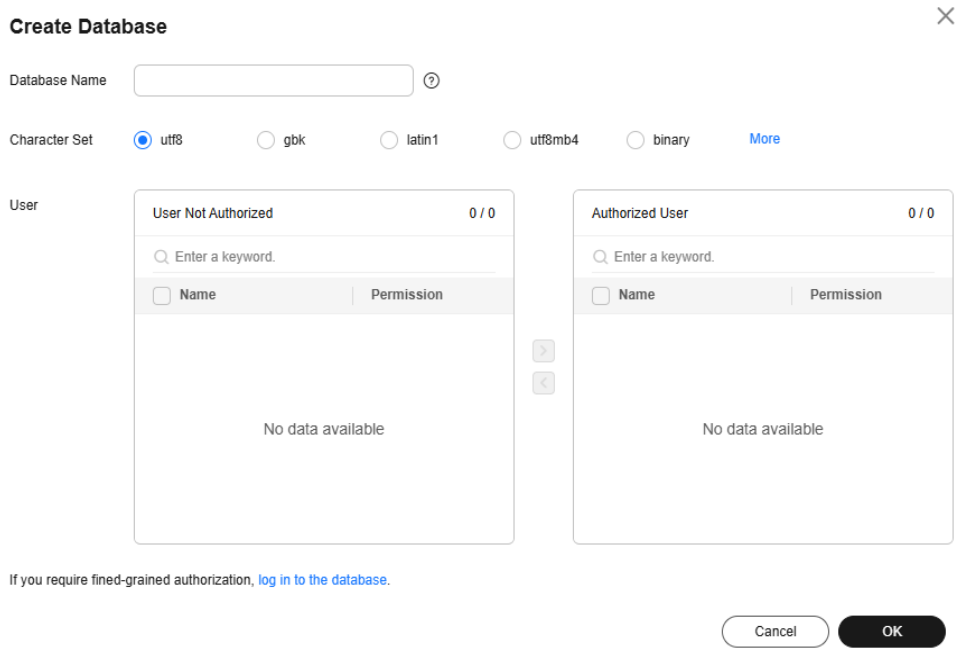

- The database name can contain 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed. The total number of hyphens (-) cannot exceed 10.
- The default character set is **utf8**. You can click **More** to view more character sets.

- Select unauthorized users and click  to authorize permissions or select authorized users and click  to revoke permissions.

  If there are no unauthorized users, you can create one by referring to **Creating a Database Account**.

Step 7 After the database is created, manage it on the **Databases** page.

----End

Creating a Database Through DAS

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the instance you want to log in and click **Log In** in the **Operation** column.

Step 5 On the displayed login page, enter the username and password and click **Log In**.

Step 6 On the top menu bar, choose **SQL Operations** > **SQL Query**.

Step 7 Run the following command to create a database:

create database *database_name*;

Step 8 Run the following command to view the database:

show databases;

----End

4.2.2 Granting Database Permissions

Scenarios

You can grant permissions to database users you have created to use specific databases or revoke permissions from specific database users.

Constraints

Permissions cannot be granted to database users for a DB instance that is in the process of being restored.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Databases and Accounts** and then click the **Databases** tab.

Step 6 Locate the target database and click **Authorize** in the **Operation** column.

Step 7 In the displayed dialog box, select unauthorized users and click › to authorize them or select authorized users and click ‹ to revoke permissions.

If no users are available, you can create one by referring to **Creating a Database Account**.

Step 8 In the displayed dialog box, click **OK**.

----End

4.2.3 Deleting a Database

Scenarios

You can delete databases that you have created.

NOTICE

Deleted databases cannot be recovered. Exercise caution when performing this operation.

Constraints

Custom databases cannot be deleted from DB instances that are in the process of being restored.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Databases and Accounts** and then click the **Databases** tab.

Step 6 Locate the target database and click **Delete** in the **Operation** column. In the displayed dialog box, enter **DELETE** and click **OK**.

Step 7 (**Optional**) If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

----End

4.3 Account Management (Non-Administrator)

4.3.1 Creating a Database Account

Scenarios

When you create a DB instance, account **root** is created at the same time by default. You can create other database accounts as needed.

You can create a database account using RDS or DAS:

- **RDS**: RDS is easy to use. There are no special commands to remember.
- **DAS**: DAS is a powerful platform that offers more flexibility, but you need to be familiar with the creation commands. The process requires a bit more expertise.

Account Type

Table 4-1 Account description

Account Type	Description
Administrator account **root**	Only the administrator account **root** is provided on the instance creation page. For details about the supported permissions, see **RDS for MariaDB Constraints**. **NOTE** Running **revoke**, **drop user**, or **rename user** on **root** may cause service interruption. Exercise caution when running any of these statements.

Account Type	Description
System accounts	To provide O&M services, the system automatically creates system accounts when you create RDS for MariaDB instances. These system accounts are unavailable to you. • **mariadb.sys**: used to create views. • **rdsAdmin**: a management account with the highest permission. It is used to query and modify instance information, rectify faults, migrate data, and restore data. • **rdsRepl**: a replication account, used to synchronize data from the primary instance to the standby instance or read replicas. • **rdsBackup**: a backup account, used for backend backup. • **rdsMetric**: a metric monitoring account used by watchdog to collect database status data. • **dsc_readonly**: used to anonymize data.
Other accounts	Accounts created through the console, APIs, or SQL statements After an account is created, you can assign permissions to it as required. For details, see **Changing Permissions for a Database Account**.

Constraints

Database accounts cannot be created for DB instances that are in the process of being restored.

Creating a Database Account Through RDS

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Databases and Accounts** and then click the **Accounts** tab.

Step 6 On the displayed page, click **Create Account**. In the displayed dialog box, specify **Username** and **Host IP Address**, authorize permissions for databases, enter a password, and confirm the password. Then, click **OK**.

Figure 4-2 Creating a database account

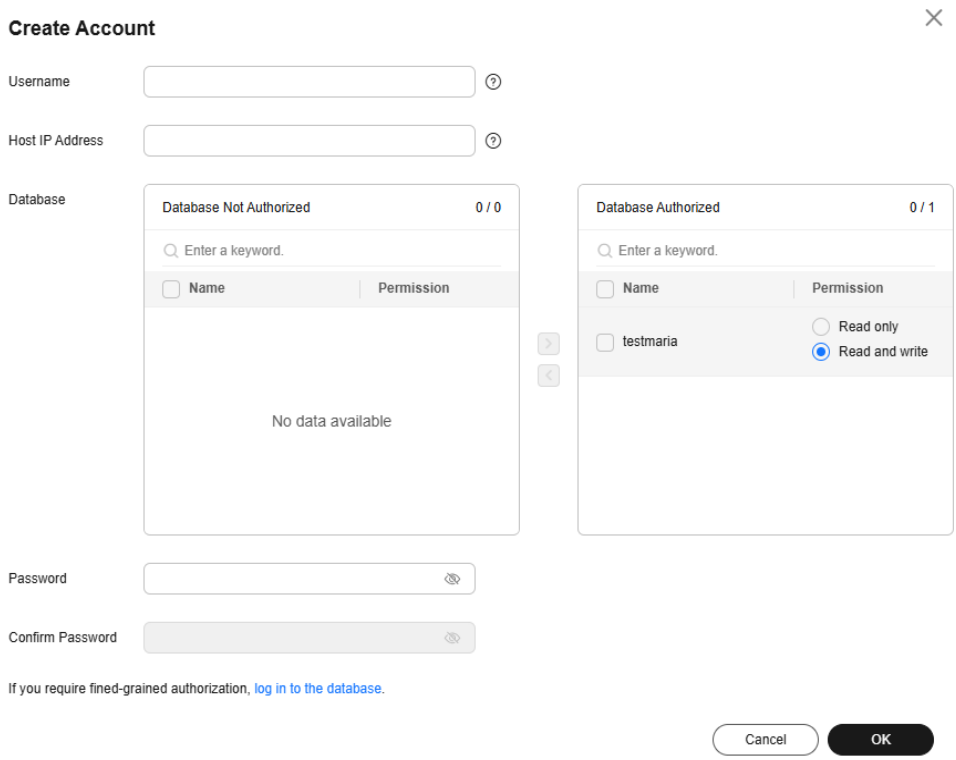

- The username consists of 1 to 32 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.
- Select unauthorized databases and click to authorize them or select authorized databases and click to revoke permissions.

  If there are no unauthorized databases, you can create one by referring to **Creating a Database**. You can also modify the permissions after the account creation by referring to **Changing Permissions for a Database Account**.
- The password must consist of 8 to 32 characters and contain all types of the following characters: uppercase letters, lowercase letters, digits, and special characters (~!@#$%^*-_=+?,()&).
- You can specify IP addresses that are allowed to access your DB instance.
  - To enable all IP addresses to access your instance, enter **%** for **Host IP Address**.
  - To enable all IP addresses in the subnet 10.10.10.*X* to access your instance, enter **10.10.10.%** for **Host IP Address**.
  - To specify multiple IP addresses, separate them with commas (,), for example, **192.168.0.1,172.16.213.9** (no spaces before or after the comma).

Step 7 After the account is created, you can manage it on the **Accounts** page of the DB instance.

----End

Creating a Database Account Through DAS

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the instance you want to log in and click **Log In** in the **Operation** column.

Step 5 On the displayed login page, enter the username and password and click **Log In**.

Step 6 Create an account.

- On the top menu bar, choose **Account Management** > **User Management**. On the displayed page, click **Create User**. Then, configure basic information, advanced settings, global permissions, and object permissions, and click **Save**. In the displayed dialog box, click **OK**.

  For details about how to set permissions, see **Creating a User**.
- You can also choose **SQL Operations** > **SQL Query** from the top menu bar and run the following command to create an account:

  create user *username*;

----End

4.3.2 Resetting a Password for a Database Account

Scenarios

You can reset passwords for the accounts you have created. To protect against brute force hacking attempts and ensure system security, change your password periodically, such as every three or six months.

Constraints

Passwords cannot be reset for DB instances that are in the process of being restored.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Databases and Accounts** and then click the **Accounts** tab.

Step 6 Locate the target account and click **Reset Password** in the **Operation** column.

Step 7 In the displayed **Reset Password** dialog box, enter and confirm a new password, and click **OK**.

Figure 4-3 Resetting a password

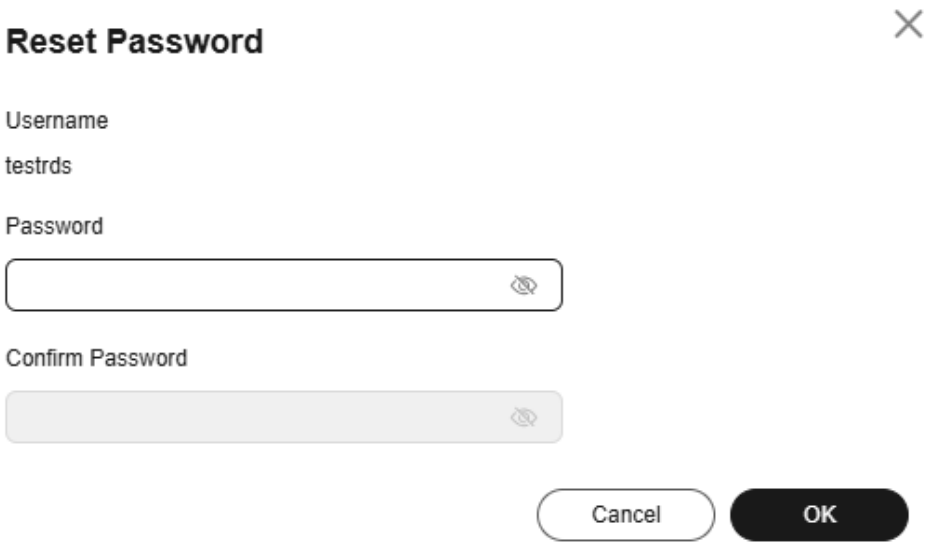

- The password must consist of 8 to 32 characters and contain all types of the following characters: uppercase letters, lowercase letters, digits, and special characters (~!@#$%^*-_=+?,()&).
- The password must be different from the username or username spelled backwards.
- You are advised to enter a strong password to improve security and prevent security risks such as brute force cracking.
- After the password is reset, the database will not be rebooted and permissions will not be changed.
- You can use Cloud Trace Service (CTS) to query the password reset records. For details, see **Viewing Traces**.

Step 8 (**Optional**) If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

----End

4.3.3 Changing Permissions for a Database Account

Scenarios

You can authorize database users you have created to specific databases or revoke permissions from authorized database users.

Constraints

Permissions cannot be changed for DB instances that are in the process of being restored.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target DB instance name to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Databases and Accounts** and then click the **Accounts** tab.

Step 6 Locate the target account and click **Change Permission** in the **Operation** column.

Step 7 In the displayed dialog box, select unauthorized databases and click to authorize them. You can also select authorized databases and click to revoke permissions.

Figure 4-4 Changing permissions

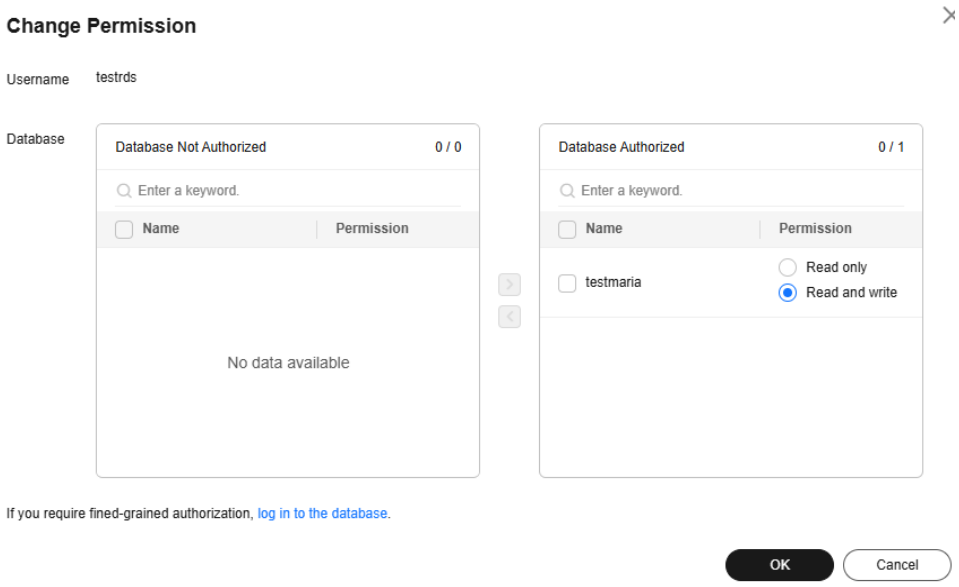

If there are no unauthorized databases, you can create one by referring to Creating a Database.

Step 8 Click **OK**.

----End

4.3.4 Modifying Host IP Addresses

Scenarios

You can change the host IP addresses that are allowed to access your instance as needed.

Constraints

This operation cannot be performed for DB instances that are being restored.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target DB instance name to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Databases and Accounts** and then click the **Accounts** tab.

Step 6 Locate the target account and choose **More** > **Modify Host IP Address** in the **Operation** column.

Step 7 In the displayed dialog box, enter the new IP addresses.

Figure 4-5 Modifying host IP addresses

- To enable all IP addresses to access your instance, enter **%** for **Host IP Address**.

- To enable all IP addresses in the subnet 10.10.10.*X* to access your instance, enter **10.10.10.%** for **Host IP Address**.
- To specify multiple IP addresses, separate them with commas (,), for example, **192.168.0.1,172.16.213.9** (no spaces before or after the comma).

Step 8 Click **OK**.

Step 9 (**Optional**) If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

----End

4.3.5 Deleting a Database Account

Scenarios

You can delete database accounts you have created.

NOTICE

Deleted database accounts cannot be restored. Exercise caution when deleting an account.

Constraints

Accounts cannot be deleted from DB instances that are in the process of being restored.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target DB instance name to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Databases and Accounts** and then click the **Accounts** tab.

Step 6 Locate the target account and choose **More** > **Delete** in the **Operation** column.

Step 7 In the displayed dialog box, enter **DELETE** and click **OK**.

Step 8 (**Optional**) If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

----End

5 Instance Management

5.1 Rebooting DB Instances or Read Replicas

Scenarios

You may need to reboot a DB instance during maintenance. For example, after you modify some parameters, a reboot is required for the modifications to take effect. You can reboot a primary DB instance or a read replica on the management console.

Constraints

- If the DB instance status is **Abnormal**, the reboot may fail.
- Rebooting DB instances will cause service interruptions. During the reboot process, the DB instance status is **Rebooting**.
- Rebooting DB instances will cause instance unavailability and clear cached memory. To prevent traffic congestion during peak hours, you are advised to reboot DB instances during off-peak hours.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the target DB instance, or click and then locate the target read replica. Choose **More** > **Reboot** in the **Operation** column.

For primary/standby DB instances, if you reboot the primary DB instance, the standby DB instance is also rebooted automatically.

Step 5 In the displayed dialog box, click **OK**.

Figure 5-1 Rebooting a DB instance

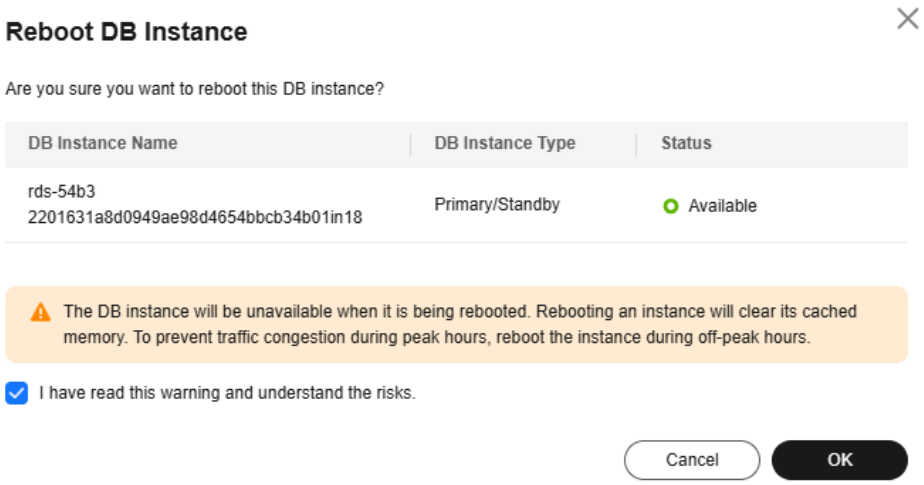

Step 6 (**Optional**) If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

Step 7 Refresh the DB instance list and view the status of the DB instance. If its status is **Available**, it has rebooted successfully.

----End

5.2 Selecting Displayed Items

Scenarios

You can customize which instance items are displayed on the **Instances** page.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click above the instance list, select desired items from the custom columns, and click **OK**.

- **Table Text Wrapping**: If you enable this function, excess text will move down to the next line.
- **Operation Column**: If you enable this function, the **Operation** column is always fixed at the rightmost position of the table.

- The following items can be displayed: **Name/ID**, **Description**, **DB Instance Type**, **DB Engine Version**, **Status**, **Billing Mode**, **Floating IP Address**, **Private Domain Name**, **IPv6 Address**, **Enterprise Project**, **Created**, **Database Port**, **Storage Type**, and **Operation**.

----End

5.3 Exporting DB Instance Information

Scenarios

You can export information about all or selected DB instances to view and analyze DB instance information.

Constraints

A tenant can export a maximum of 3,000 instances at a time. The time required for the export depends on the number of instances.

Exporting Information About All DB Instances

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click **Export** above the DB instance list. By default, information about all DB instances are exported. In the displayed dialog box, you can select the items to be exported and click **OK**.

Step 5 Find a .csv file locally after the export task is completed.

----End

Exporting Information About Selected DB Instances

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, filter DB instances by DB engine, DB instance name, DB instance ID, DB instance tag, enterprise project, or floating IP address, or select the DB instances to be exported, and click **Export** above the DB instance list. In the displayed dialog box, select the items to be exported and click **OK**.

Step 5 Find a .csv file locally after the export task is completed.

----End

5.4 Deleting a Pay-per-Use DB Instance or Read Replica

Scenarios

To release resources, you can delete DB instances or read replicas billed on the pay-per-use basis as required on the **Instances** page.

Constraints

- DB instances cannot be deleted when operations are being performed on them. They can be deleted only after the operations are complete.
- A maximum of 50 pay-per-use DB instances can be deleted at a time.
- If you delete a pay-per-use DB instance, its automated backups will also be deleted and you will no longer be billed for them. Manual backups, however, will be retained and generate additional costs.

NOTICE

- If you delete a primary DB instance, its standby DB instance and read replicas (if any) are also deleted automatically. Exercise caution when performing this operation.
- Deleted DB instances cannot be recovered and resources are released. Exercise caution when performing this operation. If you want to retain data, **create a manual backup** first before deleting the DB instance.
- You can use a manual backup to restore a DB instance. For details, see **Restoring a DB Instance from a Backup**.

Deleting a Pay-per-Use DB Instance

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the primary DB instance to be deleted and choose **More** > **Delete** in the **Operation** column.

Figure 5-2 Deleting a DB instance

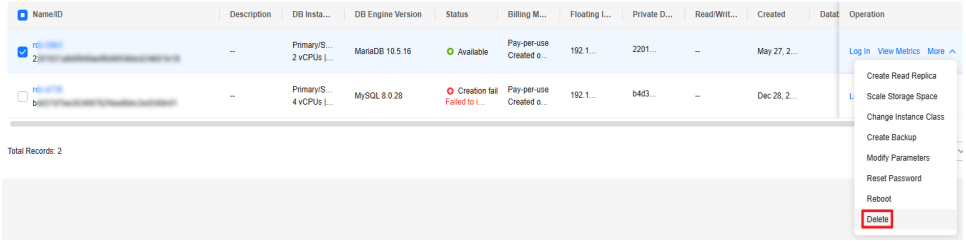

Step 5 In the displayed dialog box, enter **DELETE** and then click **OK**.

Step 6 (**Optional**) If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

Step 7 Refresh the DB instance list later to confirm that the deletion was successful.

----End

Deleting a Pay-per-Use Read Replica

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and click . All the read replicas created for the DB instance are displayed.

Step 5 Locate the read replica to be deleted and choose **More** > **Delete** in the **Operation** column.

Step 6 In the displayed dialog box, click **OK**.

Step 7 (**Optional**) If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

Step 8 Refresh the DB instance list later to check that the deletion is successful.

----End

5.5 Modifying Recycling Policy

RDS allows you to move unsubscribed yearly/monthly DB instances and deleted pay-per-use DB instances to the recycle bin. You can rebuild a DB instance that was deleted up to 7 days ago from the recycle bin.

Constraints

- Read replicas cannot be moved to the recycle bin.
- A stopped instance will not be moved to the recycle bin after being deleted.
- The recycle bin is enabled by default and cannot be disabled.

Precautions

- The recycle bin is enabled by default and cannot be disabled. This function is free of charge.
- Instances in the recycle bin are retained for 7 days by default. A new recycling policy only applies to DB instances that were put in the recycle bin after the new policy was put into effect. For DB instances that were in the recycle bin before the modification, the original recycling policy takes effect.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 In the navigation pane on the left, choose **Recycle Bin**.

Step 5 Click **Modify Recycling Policy** and set the retention period of deleted instances. The value ranges from 1 to 7 days.

Figure 5-3 Setting the retention period

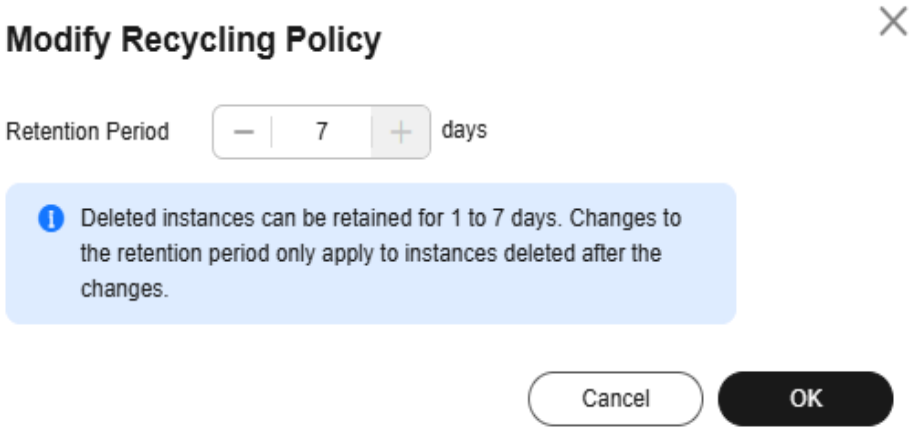

Step 6 Then, click **OK**.

----End

5.6 Rebuilding a DB Instance

You can rebuild DB instances that were deleted up to 7 days ago from the recycle bin. This section describes how to rebuild a DB instance.

Precautions

- Only primary/standby or single-node DB instances can be rebuilt.
- You can only rebuild DB instances within the retention period.

- After an instance is deleted, the system keeps the most recent automated full backup of the previous day. If there is no automated full backup generated on that day, it retains the latest available one. A full backup is also created. You can choose either backup to rebuild the instance.
- If resources are not renewed after expiration, you can rebuild DB instances from the recycle bin to restore data.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 In the navigation pane on the left, choose **Recycle Bin**.

Step 5 In the DB instance list, locate the target DB instance and click **Rebuild** in the **Operation** column.

Step 6 In the displayed dialog box, locate a backup and click **Rebuild** in the **Operation** column.

Figure 5-4 Selecting a backup to rebuild the instance

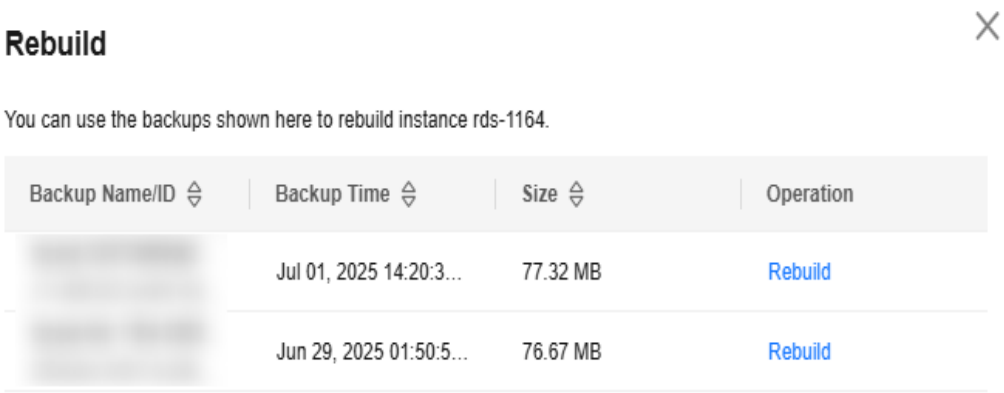

Step 7 On the displayed page, set the required parameters and submit the task. For details, see Restoring to a New Instance.

----End

6 Instance Modifications

6.1 Upgrading a Minor Version

RDS for MariaDB supports minor version upgrades. Upgrading a minor version not only fixes historical issues, but also enriches user experience. This section describes how to upgrade a minor version.

Precautions

- To upgrade a minor version, **submit a service ticket** to apply for the required permissions.
- When any new minor version is released for addressing issues and vulnerabilities from the open source community, perform a minor version upgrade for your instance.
- The upgrade will cause the DB instance to reboot and interrupt services intermittently. To limit the impact of the upgrade, perform the upgrade during off-peak hours, or ensure that your applications support automatic reconnection.
- If your RDS instance is involved in a DRS task, upgrading the minor version may cause the DRS task to fail.

  You are advised to check the retention period of RDS instance binlogs before upgrading the minor version.

  - If the binlogs are within the retention period, the DRS task will automatically restart after the minor version is upgraded.
  - If the binlogs are beyond the retention period, you need to reconfigure or recreate a DRS task.
- A minor version upgrade cannot be rolled back after the upgrade is complete. If the upgrade fails, the DB instance will be automatically rolled back to the source version.
- DDL operations on events, such as CREATE EVENT, DROP EVENT, and ALTER EVENT, are not allowed during a minor version upgrade.

Notes

- If the primary and standby instances are deployed in the same AZ, upgrading the minor version will trigger a primary/standby switchover. If they are deployed in different AZs, upgrading the minor version will trigger two switchovers.
- When you upgrade the minor version of a primary instance, the minor versions of read replicas (if any) will also be upgraded automatically. Read replicas cannot be upgraded separately.
- A minor version can be upgraded in minutes.
- For primary/standby DB instances, the standby DB instance is upgraded first and then the primary DB instance is upgraded afterwards.

Constraints

- If the replication delay between primary and standby DB instances is longer than 300 seconds, the minor version cannot be upgraded.
- Minor versions cannot be upgraded for DB instances with abnormal nodes.
- RDS for MariaDB DB instances with the event scheduler enabled do not support minor version upgrades. If you want to perform a minor version upgrade, disable the event scheduler first. For details, see **Enabling or Disabling Event Scheduler**.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 Under **DB Engine Version**, click **Upgrade Minor Version**.

If the minor version of your DB instance is already the latest, there is no need to upgrade the minor version.

Step 6 In the displayed dialog box, select a scheduled time and click **OK**.

- **Upon submission**: The system upgrades the minor version immediately after you have submitted your upgrade request.
- **In maintenance window**: (To use this function, **submit a service ticket** to apply for required permissions.) The system will upgrade the minor version during the maintenance window you set.

----End

6.2 Changing a DB Instance Name

You can change the name of a primary DB instance or read replica.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and click next to it to edit the DB instance name. Then, click **OK**.

Alternatively, click the instance name to go to the **Overview** page. Under **DB Instance Name**, click to edit the instance name.

The instance name must start with a letter and consist of 4 to 64 characters. Only letters (case-sensitive), digits, hyphens (-), and underscores (_) are allowed.

- To submit the change, click .
- To cancel the change, click .

Step 5 Check the results on the **Overview** page.

----End

6.3 Changing a DB Instance Description

Scenarios

After a DB instance is created, you can add a description.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the DB instance you wish to edit the description for and click in the **Description** column to make your modification. When you are finished, click **OK**.

Alternatively, click the instance name to go to the **Overview** page. Under **Description**, click to edit the instance description.

NOTE

The instance description can contain a maximum of 64 characters.

- To submit the change, click .

- To cancel the change, click ✕.

Step 5 Check the results on the **Overview** page.

----End

6.4 Changing the Replication Mode

Scenarios

You can change the replication mode for primary/standby DB instances to **Asynchronous** or **Semi-synchronous**.

- **Asynchronous**:

  When applications update data, the primary DB instance responds to the applications immediately after data is updated. This mode provides better performance than the semi-synchronous mode.

- **Semi-synchronous** (default value):
  - When applications update data, the primary DB instance responds to the applications only after the standby DB instance receives logs, which affects database performance.
  - If the standby DB instance is abnormal, the primary DB instance waits for the response of the standby DB instance for several seconds and does not respond to write operations during this period.
    - If the standby DB instance is recovered during the waiting period, the primary DB instance starts to respond to write operations normally.
    - If the standby DB instance is not recovered during the waiting period, the replication mode is automatically switched to asynchronous. After the switchover is complete, the primary DB instance starts to respond to write operations.

Procedure

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the primary instance name.

Step 5 On the displayed **Overview** page, find **Replication Mode** and click **Configure** under it. In the displayed dialog box, select a mode and click **OK**.

Step 6 Check the results on the **Overview** page.

----End

6.5 Changing the Failover Priority

Scenarios

RDS gives you control over the failover priority of your primary/standby DB instance. You can set it to **Reliability** or **Availability**.

- **Reliability** (default setting): Data consistency is preferentially ensured during a primary/standby failover. This is recommended for applications whose highest priority is data consistency.
- **Availability**: Database availability is preferentially ensured during a primary/standby failover. This is recommended for applications that require databases to provide uninterrupted online services.

Constraints

The failover priority cannot be changed when the DB instance is stopped or its instance class is being changed.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the primary instance name.

Step 5 On the displayed **Overview** page, find **Failover Priority** and click **Configure** under it. In the displayed dialog box, select a priority and click **OK**.

Step 6 Check the results on the **Overview** page.

----End

6.6 Enabling or Disabling Event Scheduler

Scenarios

Event scheduler manages the scheduling and execution of events. The built-in event scheduler cannot guarantee the consistency of event statuses between primary and standby DB instances. If a failover or switchover occurs, events will not be scheduled. RDS for MariaDB resolves this issue. With RDS for MariaDB, even if there is a failover or switchover, the events will still be properly scheduled. You can simply enable or disable the event scheduler on the RDS console.

Notes

- By default, the event scheduler is disabled after a DB instance is created.
- After a primary/standby failover or switchover is performed, the event scheduler status remains unchanged. The **event_scheduler** is **on** for the original primary DB instance and **off** for the original standby DB instance.
- After a restoration to a new DB instance, the event scheduler status is the same as that of the original DB instance.
- After a single-node DB instance is changed to a primary/standby DB instance, the event scheduler status is the same as that of the primary DB instance.

Constraints

Read replicas do not support this function.

Enabling Event Scheduler

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 On the **Overview** page, click **Enable** under **Event Scheduler**.

NOTICE

After the event scheduler is enabled, reactivate the previously created events to ensure that the event statuses on the primary and standby instances are the same.

----End

Disabling Event Scheduler

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the DB instance name.

Step 5 On the **Overview** page, click **Disable** under **Event Scheduler**.

----End

FAQs

What Should I Do If I Do Not Have the Permission to Change the event_scheduler Settings?

Answer: You can only enable or disable the event scheduler on the console. For details, see this section.

6.7 Changing a DB Instance Class

Scenarios

You can change the instance class (vCPU or memory) of a DB instance as required. If the status of a DB instance changes from **Changing instance class** to **Available**, the change is successful.

Constraints

- You can change the DB instance class only when your account balance is greater than or equal to $0 USD.
- An instance cannot be deleted while its instance class is being changed.
- The following operations cannot be performed on an instance whose instance class is being changed: rebooting the instance, scaling up storage space, modifying the parameter template, creating a manual backup, creating a database account, and creating a database.
- You can scale up or down your RDS for MariaDB instance specifications.
- If there are any large transactions being processed during an instance class change, the change may fail.
- If the primary/standby replication delay of a DB instance is longer than 5 minutes, the instance class change will fail.
- Changing an instance class will interrupt services. Ensure that your applications support automatic reconnection. Perform this operation during off-peak hours because changing an instance class during peak hours takes much more time.
- Changing the instance class takes 5 to 15 minutes (during off-peak hours). If more time is required, submit a service ticket.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and choose **More** > **Change Instance Class** in the **Operation** column.

Step 5 On the displayed page, specify the new instance class and click **Next**.

Figure 6-1 Changing a DB instance class

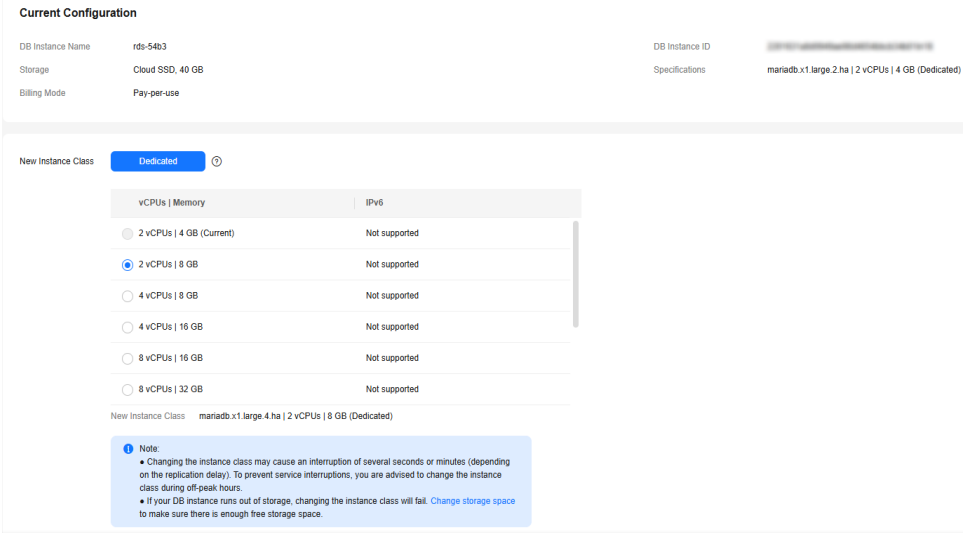

Step 6 Confirm the specifications.

- If you need to modify your settings, click **Previous**.
- Click **Submit** for pay-per-use DB instances.

  To view the cost incurred by the DB instance class change, choose **Billing & Costs** > **Bills** in the upper right corner.
- For yearly/monthly DB instances:
  - If you intend to scale down the DB instance class, click **Submit**.

    The refund is automatically returned to your account. You can click **Billing** in the upper right corner and then choose **Orders** > **My Orders** in the navigation pane on the left to view the details.
  - If you intend to scale up the DB instance class, click **Pay Now**. The scaling starts only after the payment is successful.

Step 7 View the DB instance class change result.

Return to the **Instances** page and view the instance status. During the change period, the instance status is **Changing instance class**. You can view the execution progress of **Changing a MariaDB DB instance class** on the **Task Center** page. After a few minutes, view the DB instance class on the **Overview** page to check that the change is successful.

----End

6.8 Scaling Up Storage Space

Scenarios

If the original storage space is insufficient as your services grow, scale up storage space of your DB instance. **The backup space increases with the instance scale-up.**

The DB instance needs to preserve at least 13% of its capacity to work properly. The new minimum storage space required to make this instance available has been automatically calculated for you.

You are advised to set alarm rules for the storage space usage by referring to **Setting Alarm Rules**.

RDS allows you to scale up storage space of DB instances but you cannot change the storage type. **During the scale-up period, services are not interrupted.**

Constraints

- You can scale up storage space only when your account balance is greater than or equal to $0 USD.
- You can scale up storage space only when your instance status is **Available** or **Storage full**.
- The maximum allowed storage is 4,000 GB. If you want to increase the storage upper limit to 10 TB, **submit a service ticket**.
- When storage space is being scaled up, the DB instance is in **Scaling up** state and the backup tasks of the instance are not affected.
- For primary/standby DB instances, scaling up the primary DB instance will cause the standby DB instance to also be scaled up accordingly.
- Reboot is not required during scale-up.
- You cannot reboot or delete a DB instance that is being scaled up.
- Storage space can only be scaled up, not down.
- If you scale up a DB instance with the disk encrypted, the expanded storage space will also be encrypted using the original encryption key.

Scaling Up a Primary DB Instance

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and choose **More** > **Scale Storage Space** in the **Operation** column.

You can also perform the following operations to scale up storage space:

- Click the target instance name to enter the **Overview** page. In the **Storage & Backup** area, click **Scale Storage Space**.
- If the storage space is full, locate the DB instance on the **Instances** page and click **Scale** in the **Status** column.

Step 5 On the displayed page, specify the new storage space and click **Next**.

The minimum increment for each scaling is 10 GB.

Step 6 Confirm specifications.

- If you need to modify your settings, click **Previous**.
- If you do not need to modify the settings, click **Submit** for a pay-per-use instance or click **Pay Now** for a yearly/monthly instance.

Step 7 View the scale-up result.

Scaling up storage space takes 3-5 minutes. During this time period, the status of the DB instance on the **Instances** page will be **Scaling up**. After a while, click the instance name and view the new storage space on the displayed **Overview** page to verify that the scale-up is successful.

For RDS for MariaDB instances, you can view the detailed progress of the task on the **Task Center** page. For details, see **Task Center**.

----End

Scaling Up a Read Replica

Scaling up the storage space of a read replica does not affect that of the primary DB instance. Therefore, you can separately scale read replicas to meet service requirements. New storage space of read replicas after scaling up must be greater than or equal to that of the primary DB instance.

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and click in front of it. Locate the read replica to be scaled and choose **More** > **Scale Storage Space** in the **Operation** column.

You can also perform the following operations to scale up storage space:

- Click the read replica name to enter the **Overview** page. In the **Storage & Backup** area, click **Scale Storage Space**.
- If the storage space is full, locate the read replica on the **Instances** page and click **Scale** in the **Status** column.

Step 5 On the displayed page, specify the new storage space and click **Next**.

The minimum increment for each scaling is 10 GB.

Step 6 Confirm specifications.

- If you need to modify your settings, click **Previous**.
- If you do not need to modify your settings and the read replica uses pay-per-use billing, click **Submit**.

Step 7 View the scale-up result.

Scaling up storage space takes 3-5 minutes. During the time period, the status of the read replica on the **Instances** page will be **Scaling up**. After a while, click the read replica name and view the new storage space on the displayed **Overview** page to verify that the scale-up is successful.

For RDS for MariaDB read replicas, you can view the detailed progress of the task on the **Task Center** page. For details, see **Task Center**.

----End

6.9 Configuring Storage Autoscaling

With storage autoscaling enabled, when RDS detects that you are running out of database space, it automatically scales up your storage.

Autoscaling up the storage of a read replica does not affect that of the primary DB instance. Therefore, you can separately autoscale read replicas to meet service requirements. New storage space of read replicas after autoscaling up must be greater than or equal to that of the primary DB instance.

Constraints

- You can enable storage autoscaling only when your account balance is greater than or equal to $0 USD.
- The storage space can be scaled up automatically only when your instance status is **Available** or **Storage full**.
- To enable storage autoscaling, **submit a service ticket** to request permissions.
- The maximum allowed storage is 4,000 GB.
- For a primary/standby DB instance, autoscaling the storage for the primary node will also autoscale the storage for the standby node.
- Storage autoscaling is unavailable when the DB instance is in any of the following statuses: changing instance class, upgrading a minor version, migrating the standby DB instance, and rebooting.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target DB instance or read replica name (click in front of a DB instance to locate its read replica).

Step 5 In the **Storage & Backup** area, toggle on the **Storage Autoscaling** switch. If this switch is not displayed, **submit a service ticket** to request required permissions.

Step 6 In the displayed dialog box, set the following parameters:

Figure 6-2 Configuring autoscaling

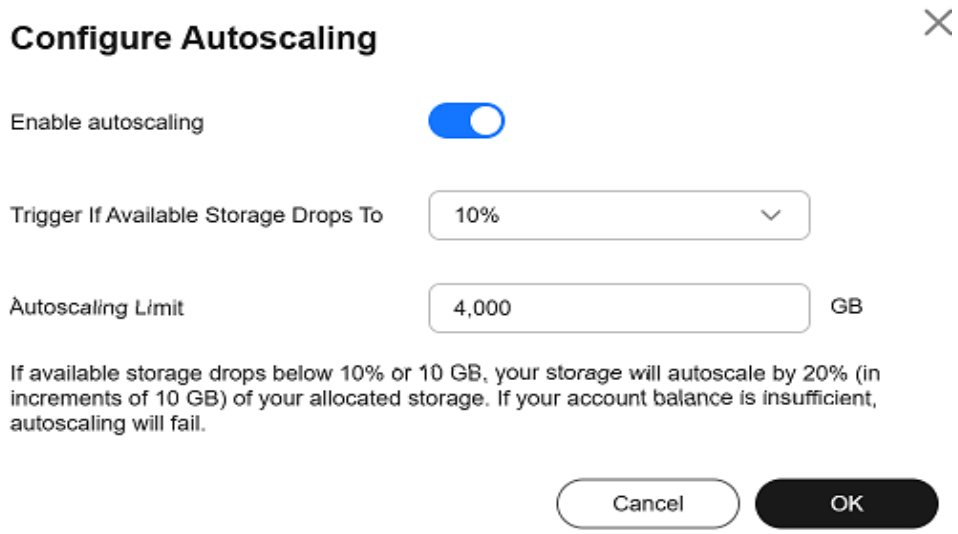

Table 6-1 Parameter description

Parameter	Description
Enable autoscaling	Select **Enable autoscaling**.
Trigger If Available Storage Drops To	If the available storage drops to a specified threshold or 10 GB, autoscaling is triggered.
Autoscaling Limit	The default value range is from 40 to 4,000, in GB. The limit must be no less than the storage of the DB instance.

Step 7 Click **OK**.

----End

6.10 Manually Switching Between Primary and Standby DB Instances

Scenarios

If the primary node of a primary/standby instance is unavailable due to a VM, disk, or network fault, RDS automatically triggers a primary/standby failover. In addition to automatic failover, you can manually switch between the primary and standby instances for rack-level DR.

Introduction to Primary/Standby DB Instances

A primary/standby instance uses an HA architecture and can be deployed across AZs. The primary and standby instances share the same IP address.

- Data is synchronized from the primary instance to the standby instance in real time. You can only access the primary instance. The standby instance serves as a backup.
- When a primary instance is being created, a standby instance is provisioned along with it to provide data redundancy. The standby instance is invisible to you after being created.
- If the primary instance fails, a failover occurs, during which database connection is interrupted. If there is a replication delay between the primary and standby instances, the failover takes an extended period of time. The client needs to be able to reconnect to the instance.

Constraints

- A manual switchover does not change the connection information of the DB instance, including its VPC, subnet, security group, floating IP address, private domain name, and database port.
- A primary/standby switchover may cause a brief interruption of several seconds or minutes (depending on the replication delay). If transaction logs are generated at a speed higher than 30 MB/s, services will probably be interrupted for several minutes. To prevent traffic congestion, perform a switchover during off-peak hours.
- You can perform a switchover only when all of the following conditions are met:
  - The DB instance is running properly.
  - The primary/standby replication is normal.
  - The replication delay is less than 5 minutes, and the data on the primary and standby instances is consistent.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target primary/standby instance name to go to the **Overview** page.

Step 5 Click  to show all the details.

Step 6 Under **DB Instance Type**, click **Switch**.

Step 7 (**Optional**) If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see the ***Identity and Access Management User Guide***.

Step 8 In the primary/standby switchover dialog box, click **OK**.

Step 9 After the switchover is successful, check the status of the DB instance on the **Instances** page.

- During the switchover, the DB instance status is **Switchover in progress**.
- In the upper right corner of the DB instance list, click  to refresh the list. After the switchover is successful, the DB instance status will become **Available**.

----End

6.11 Changing the Maintenance Window

Scenarios

The maintenance window is 02:00–06:00 by default and you can change it as required.

Precautions

- During the maintenance window, the DB instance will be intermittently disconnected once or twice. Ensure that your applications support automatic reconnection.
- To prevent service interruption, you are advised to set the maintenance window to off-peak hours.
- Changing the maintenance window does not affect the execution time of the scheduled tasks in the original maintenance period.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target instance name to go to the **Overview** page. Under **Maintenance Window**, click **Configure**.

Step 5 In the displayed dialog box, select an interval, select a start time from the drop-down list, and click **OK**.

Figure 6-3 Changing the maintenance window

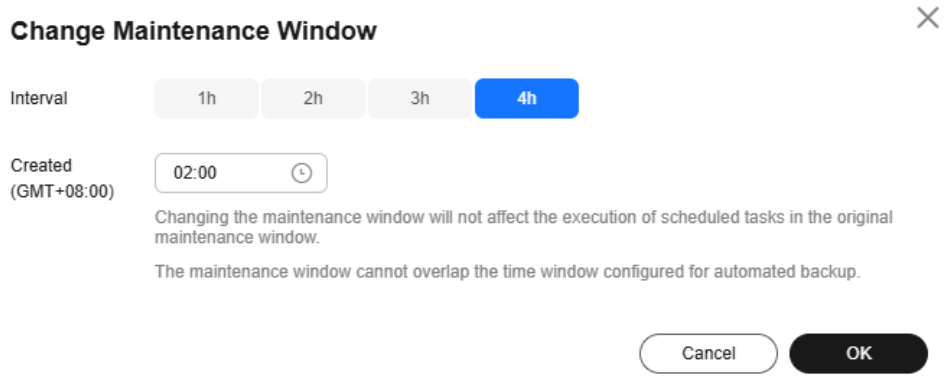

----End

7 Data Backups

7.1 Backup Solutions

RDS supports automated and manual backups. You can periodically back up databases. If a database is faulty or data is damaged, you can restore the database using backups to ensure data reliability.

RDS uses **sysbench** to import data models and a certain amount of data. After data is backed up, the compression ratio is about 80%. The more duplicate data there is, the higher the compression ratio is.

Compression ratio = Space occupied by backup files/Space occupied by data files x 100%

Backup Types

- Full backup: A full backup is to back up all data, even if no data has changed since the last backup.

  Full backups can be triggered automatically (by configuring a same-region backup policy) or manually.
- Incremental backup (binlog backup): RDS automatically backs up data modifications made after the most recent full or incremental backup every five minutes.

How RDS Backs Up Data

- Single-node instance

  A single-node architecture, which is more cost-effective than mainstream primary/standby DB instances. After a backup is triggered, data is backed up from the primary instance and stored as a package on OBS. The backup does not take up storage space of the instance.
- Primary/standby instance

  An HA architecture. In a primary/standby pair, each instance has the same instance class. After a backup is triggered, data is backed up from the standby instance and stored as a package on OBS. The backup does not take up storage space of the instance.

If a database or table in the primary instance is maliciously or mistakenly deleted, the database or table in the standby instance will also be deleted. In this case, you can only use backups to restore the deleted data.

Backup Solutions

Table 7-1 lists RDS backup solutions.

Table 7-1 Backup solutions

Task	Backup Type	Description
Backing up data	Automated backups	RDS automatically creates full backups for your instance during a backup window you specified and saves the backups based on the configured retention period. If necessary, you can restore data to any point in time within the backup retention period. Once the automated backup policy is enabled, a full physical backup is triggered immediately. After that, full backups will be created according to the specified time window and backup cycle. Incremental backups are automatically created every 5 minutes to ensure data reliability.
	Manual backups	Manual backups are user-initiated full backups of instances. The backup method is physical backup. Manual backups will not be deleted until you delete them manually.
	Incremental backups	Incremental backups are binary log (binlog) backups. Binary logging is enabled for RDS for MariaDB instances by default. You do not need to set an interval for incremental backups because RDS automatically backs up incremental data every 5 minutes. Incremental backups can be used to restore data to a specific point in time.
Downloading backups	Downloading a Full Backup File	You can use OBS Browser+, the browser, or the download URL to download a full backup.
	Downloading incremental backups	You can download a single binlog file or merged binlog file. To download a merged binlog file, use any of the following methods: OBS Browser+, the browser, or the download URL.

Billing

Backups are saved as packages in OBS buckets. Backups occupy backup space in OBS. If the free space RDS provides is used up, the additional space required will be billed. For the billing details, see **How Is RDS Backup Data Billed?**

Deleting Backups

- Manual backups and automated backups can be deleted in different ways:
  - Manual backups can only be manually deleted.
  - Automated backups cannot be manually deleted. You can adjust their retention period by referring to **Configuring a Same-Region Backup Policy**, and backups that expire will be automatically deleted.
- Local binlogs

  If the retention period is set to **0**, the binlogs of your DB instance will be deleted once they are synchronized to the standby instance and read replicas and successfully backed up to OBS.

  If the retention period is set to a value greater than 0, for example, 1 day, the binlogs will be retained for one day after they are synchronized to the standby instance and read replicas from the primary instance and successfully backed up to OBS. After the retention period expires, the binlogs will be automatically deleted.

7.2 Performing Backups

7.2.1 Configuring a Same-Region Backup Policy

If a DB instance fails or its data is damaged, you can restore it from backups to ensure data reliability. You can customize a same-region backup policy as required and then RDS backs up data based on the backup policy you configured. This section describes how to configure a same-region backup policy.

Notes

- The backups generated using the same-region backup policy are full backups. Binlog backups are incremental backups automatically generated by RDS every 5 minutes.
- RDS backs up data at the DB instance level, rather than the database level.
- Backups are saved as packages in OBS buckets to ensure data confidentiality and durability.
- When you create an RDS DB instance, same-region backup is enabled by default. For security purposes, this function cannot be disabled after the instance is created.

Precautions

- Since backing up data affects database read and write performance, the backup time window should be set to off-peak hours.
- Same-region backups cannot be manually deleted. To delete them, you can adjust the retention period specified in your **same-region backup policy**.

Retained backup files will be automatically deleted at the end of the retention period.

Constraints

- Rebooting the instance is not allowed during full backup. Exercise caution when selecting a backup time window.
- The system verifies the connection to the DB instance when starting a full backup task. If either of the following conditions is met, the verification fails and a retry is automatically performed. If the retry fails, the backup will fail.
  - DDL operations are being performed on the DB instance.
  - The backup lock failed to be obtained from the DB instance.

Billing

Backups are saved as packages in OBS buckets.

Viewing or Modifying a Same-Region Backup Policy

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the DB instance name.

Step 5 In the navigation pane on the left, choose **Backups & Restorations**. On the displayed page, click **Configure Same-Region Backup Policy**.

Figure 7-1 Configuring a backup policy

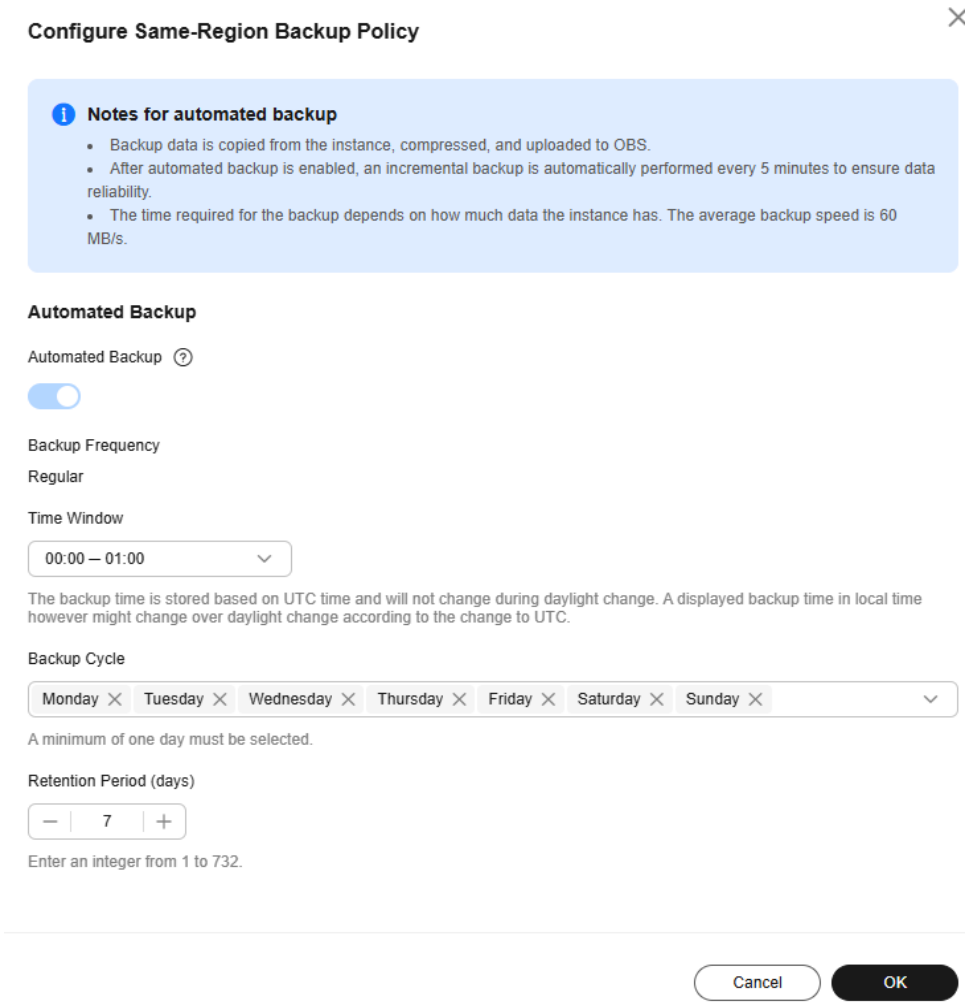

Step 6 View the configured backup policy. To modify the backup policy, adjust the values of the following parameters:

Table 7-2 Parameter description

Parameter	Description
Retention Period	How many days your automated full backups and binlog backups can be retained. The retention period is from 1 to 732 days and the default value is **7**. • Extending the retention period improves data reliability. • Reducing the retention period takes effect for all backups. Any backups that have expired will be automatically deleted.

Parameter	Description
Time Window	A one-hour period the backup will be scheduled within 24 hours, such as 01:00-02:00 or 12:00-13:00. The backup time window indicates when the backup starts. The backup duration depends on the data volume of your instance.
Backup Cycle	By default, each day of the week is selected. You can change the backup cycle and must select at least one day of the week.

Step 7 Click **OK**.

----End

7.2.2 Creating a Manual Backup

Scenarios

RDS allows you to create manual backups for a running primary DB instance. You can use these backups to restore data.

NOTE

When you delete a DB instance, its automated backups are also deleted but its manual backups are retained.

Constraints

- The number of tables in a DB instance affects the backup speed. The maximum number of tables is 500,000.
- The system verifies the connection to the DB instance when starting a full backup task. If either of the following conditions is met, the verification fails and a retry is automatically performed. If the retry fails, the backup will fail.
  - DDL operations are being performed on the DB instance.
  - The backup lock failed to be obtained from the DB instance.

Billing

Backups are saved as packages in OBS buckets.

Method 1

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and choose **More** > **Create Backup** in the **Operation** column.

Step 5 In the displayed dialog box, enter a backup name and description.

- The backup name must consist of 4 to 64 characters and start with a letter. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), and underscores (_).
- The description consists of a maximum of 256 characters and cannot contain carriage return characters or the following special characters: >!<"&'=
- The time required for creating a manual backup depends on the amount of data.

To check whether the backup has been created, you can click  in the upper right corner of the page to check the DB instance status. If the DB instance status becomes **Available** from **Backing up**, the backup has been created. You can manage the backup following the instructions provided in **Step 7**.

Figure 7-2 Creating a backup

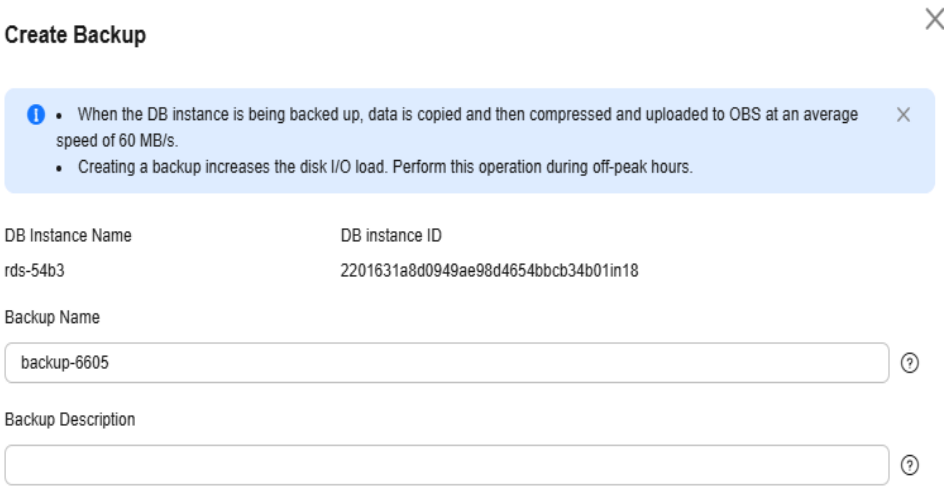

Step 6 Click **OK**.

Step 7 After a manual backup has been created, view and manage it on the **Backups** page.

Alternatively, click the target DB instance. On the **Backups & Restorations** page, you can view and manage the manual backup.

----End

Method 2

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target DB instance.

Step 5 On the **Backups & Restorations** page, click **Create Backup**. In the displayed dialog box, enter a backup name and description.

- The backup name can contain 4 to 64 characters and must start with a letter. Only letters (case-sensitive), digits, hyphens (-), and underscores (_) are allowed.
- The description can consist of a maximum of 256 characters and cannot contain carriage return characters or the following special characters: >!<"&'=
- The time required for creating a manual backup depends on the amount of data.

Step 6 Click **OK**.

Step 7 After a manual backup has been created, view and manage it on the **Backups** page.

Alternatively, click the target DB instance. On the **Backups & Restorations** page, you can view and manage the manual backup.

----End

7.2.3 Replicating a Backup

Scenarios

RDS supports replication of automated and manual backups.

Constraints

You can replicate backups and use them only within the same region.

Backup Retention Policy

- If a DB instance is deleted, the automated backups created for it are also deleted.
- If an **automated backup policy** is enabled, the automated backups will be deleted after the backup retention period expires.
- If you want to retain the automated backups for a long time, you can replicate them to generate manual backups, which will be always retained until you delete them.
- If the storage occupied by manual backups exceeds the provisioned free backup storage, additional storage costs may incur.
- Replicating a backup does not interrupt your services.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name. On the **Backups & Restorations** page, locate the backup to be replicated and click **Replicate** in the **Operation** column.

Alternatively, choose **Backups** in the navigation pane on the left. On the displayed page, locate the backup to be replicated and click **Replicate** in the **Operation** column.

Step 5 In the displayed dialog box, enter a new backup name and description and click **OK**.

- The backup name must consist of 4 to 64 characters and start with a letter. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), and underscores (_).
- The description consists of a maximum of 256 characters and cannot contain carriage return characters or the following special characters: >!<"&'=

Step 6 After the new backup has been created, you can view and manage it on the **Backups** page.

----End

7.3 Managing Backups

7.3.1 Downloading a Full Backup File

Scenarios

This section describes how to download a manual or an automated backup file to a local device and restore data from the backup file.

RDS for MariaDB allows you to download full backup files in .qp format.

Constraints

- Full backup files of frozen DB instances cannot be downloaded.
- When you use OBS Browser+ to download backup data, there is no charge for the outbound traffic from OBS.
- If the size of the backup data is greater than 400 MB, you are advised to use OBS Browser+ to download the backup data.

Method 1: Using OBS Browser+

Step 1 **Log in to the management console**.

Step 2 Click ⌖ in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Backups** page, locate the backup to be downloaded and click **Download** in the **Operation** column.

Alternatively, click the target DB instance. In the navigation pane on the left, choose **Backups & Restorations**. On the **Full Backups** page, locate the backup to be downloaded and click **Download** in the **Operation** column.

Step 5 If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

Step 6 In the displayed dialog box, select **Use OBS Browser+** for **Download Method** and download the backup as prompted.

Figure 7-3 Using OBS Browser+

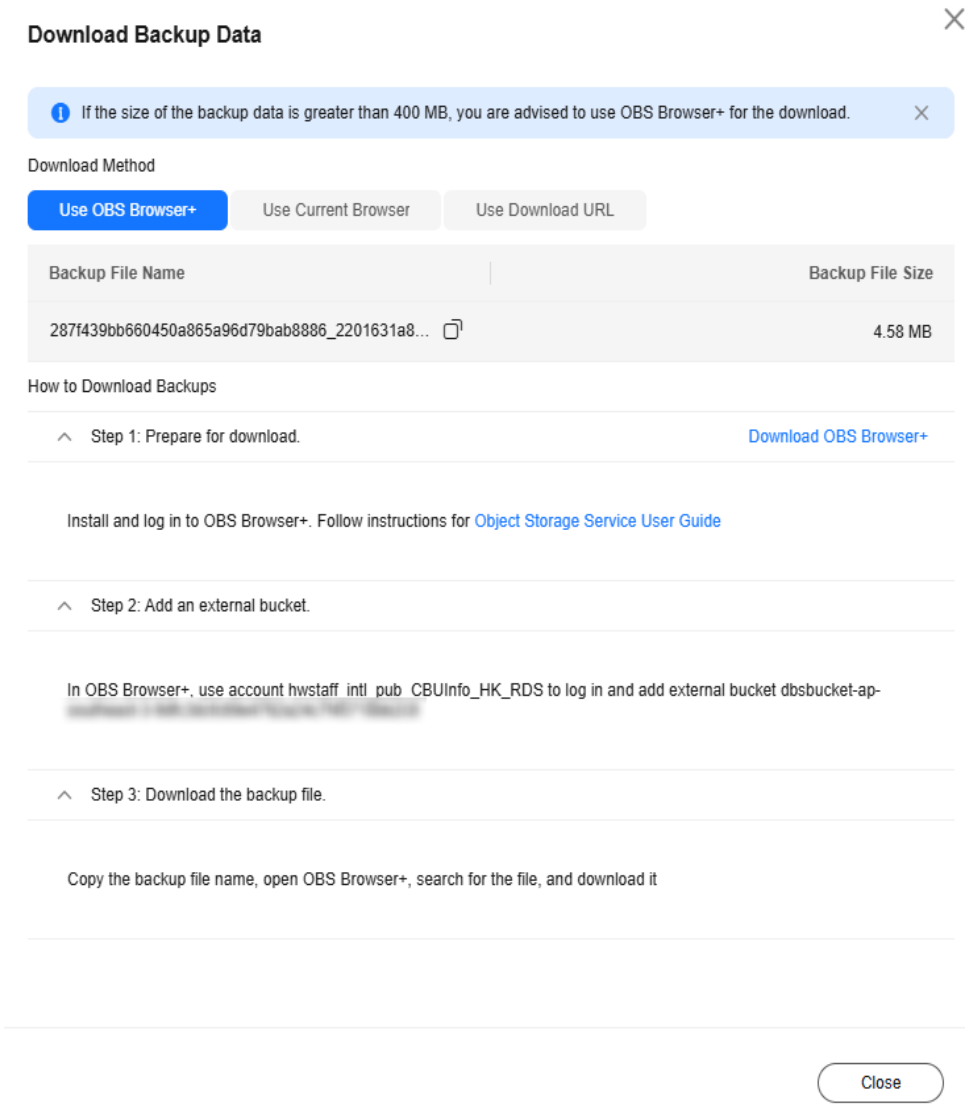

1. Download OBS Browser+ following step 1 provided on the download guide page.
2. Decompress and install OBS Browser+.
3. Log in to OBS Browser+ using the username provided in step 2 on the download guide page.

Figure 7-4 Logging in to OBS Browser+

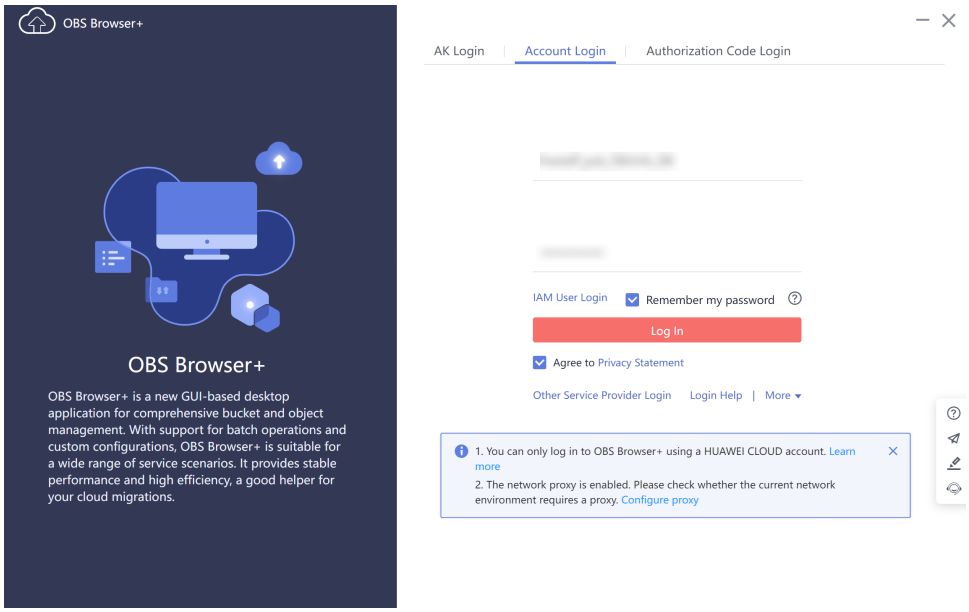

4. Add an external bucket using the bucket name provided in step 2 on the download guide page.

Figure 7-5 Adding an external bucket

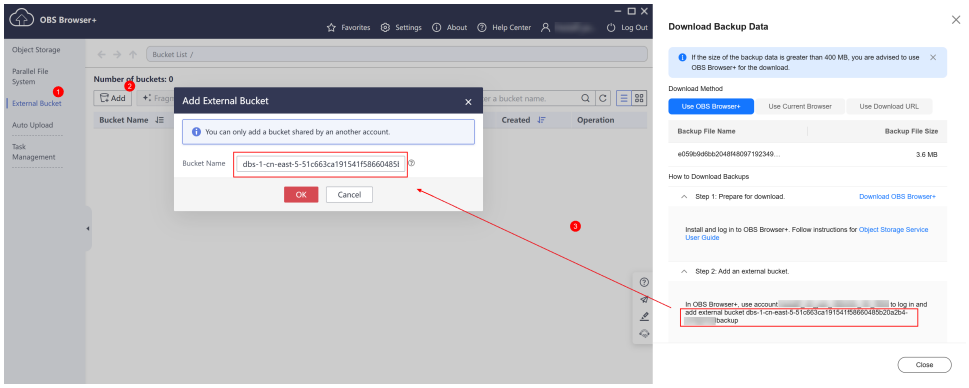

NOTE

If you want to access OBS external buckets across accounts, the access permission is required. For details, see **Granting IAM Users Under an Account the Access to a Bucket and Resources in the Bucket**.

5. Download the backup file.

   On the OBS Browser+ page, click the bucket that you added. In the search box on the right of the object list page, enter the backup file name provided in

step 3 on the download guide page. In the search result, locate the target backup and click in the **Operation** column.

Figure 7-6 Downloading a backup

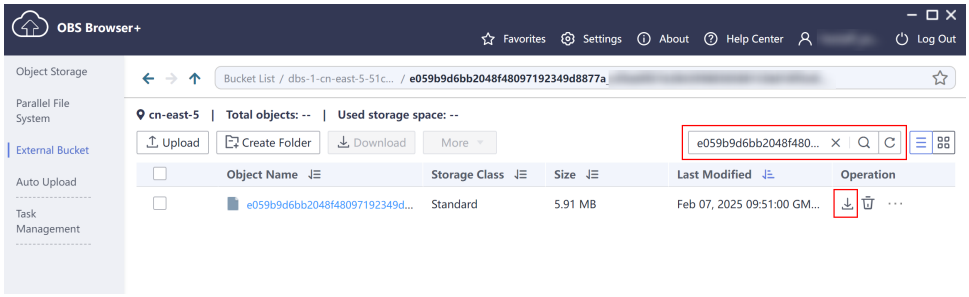

----End

Method 2: Using Current Browser

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Backups** page, locate the backup to be downloaded and click **Download** in the **Operation** column.

Alternatively, click the target DB instance. In the navigation pane on the left, choose **Backups & Restorations**. On the **Full Backups** page, locate the target backup to be downloaded and click **Download** in the **Operation** column.

Step 5 If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

Step 6 In the displayed dialog box, select **Use Current Browser** for **Download Method**.

Figure 7-7 Using the current browser

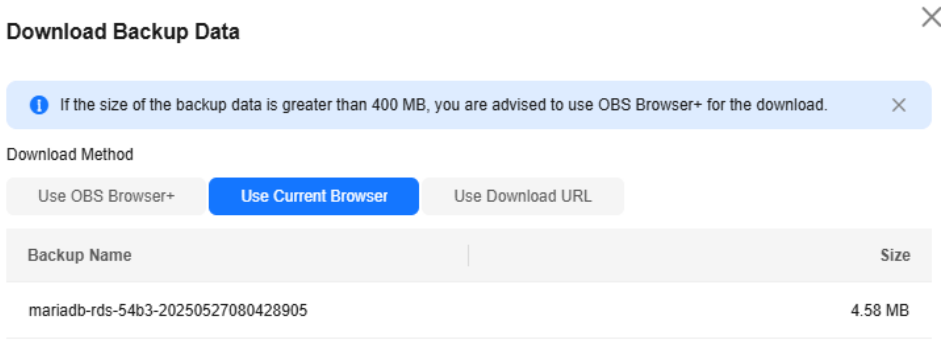

Step 7 Click **OK**.

----End

Method 3: Using Download URL

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Backups** page, locate the backup to be downloaded and click **Download** in the **Operation** column.

Alternatively, click the target DB instance. In the navigation pane on the left, choose **Backups & Restorations**. On the **Full Backups** page, locate the target backup to be downloaded and click **Download** in the **Operation** column.

Step 5 If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

Step 6 In the displayed dialog box, select **Use Download URL** for **Download Method**, click to copy the URL, and enter the URL in your browser.

Figure 7-8 Using the download URL

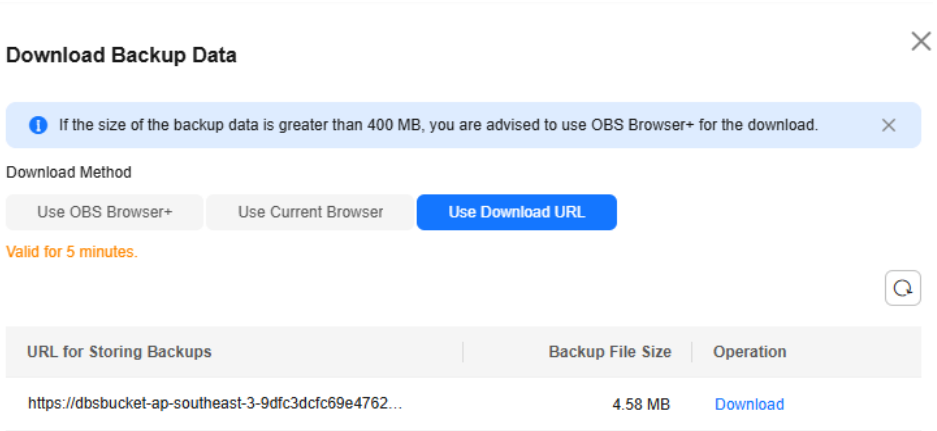

A valid URL for downloading the backup data is displayed. Download the backup file in either of the following ways:

- Using other download tools, such as your browser or Thunder, to download the backup file
- Running the **wget** command to download the backup file

  wget -O *FILE_NAME* **--no-check-certificate "***DOWNLOAD_URL***"**

Table 7-3 Parameter description

Parameter	Description
FILE_NAME	The new backup file name after the download is successful. The original backup file name may be too long and exceed the maximum characters allowed by the client file system. You are advised to use the **-O** argument with wget to rename the backup file.
DOWNLOAD_URL	The location of the backup file to be downloaded. If the location contains special characters, escape is required.

----End

7.3.2 Downloading a Binlog Backup File

Scenarios

RDS for MariaDB allows you to download binlog backup files to your client computer and use them to restore DB instances if necessary.

NOTE

The completion time displayed in the binlog backup file list indicates the time when the last transaction was committed.

Binlog backups on the management console are named in the format of "binlog name +timestamp" and use the row-based logging.

Binlog backup files of frozen DB instances cannot be downloaded.

Downloading a Binlog Backup File

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target instance name to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Backups & Restorations**. On the **Binlog Backups** page, locate the target backup to be downloaded and click **Download** in the **Operation** column.

You can also select the binlog backups to be downloaded and click **Download** above the list.

Step 6 After the download is complete, you can view the binlog backups on your computer.

----End

7.3.3 Checking and Exporting Backup Information

Scenarios

You can export backup information of RDS DB instances to an Excel file for further analysis. The exported information includes the DB instance name, backup start and end time, backup status, and backup size.

Procedure

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 In the navigation pane, choose **Backups**. On the displayed page, select the backups you want to export and click **Export** to export backup information.

- Only the backup information displayed on the current page can be exported. The backup information displayed on other pages cannot be exported.
- The backup information is exported to an Excel file for your further analysis.

Step 5 View the exported backup information.

----End

7.3.4 Deleting a Manual Backup

Scenarios

You can delete manual backups to free up backup storage.

Constraints

- Deleted manual backups cannot be recovered.
- Manual backups that are being created cannot be deleted.

Procedure

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 In the navigation pane on the left, choose **Backups**. On the displayed page, locate the manual backup to be deleted and choose **More** > **Delete** in the **Operation** column.

The following backups cannot be deleted:

- Automated backups
- Backups that are being restored
- Backups that are being replicated

Step 5 In the displayed dialog box, click **Yes**.

Step 6 If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

----End

7.4 Clearing Binlogs

7.4.1 Setting a Local Retention Period for RDS for MariaDB Binlogs

RDS for MariaDB deletes local binlogs after they are backed up to OBS. You can set the local retention period for binlogs as required.

NOTE

Binary logging is enabled for RDS by default and uses row-based logging.

Read replicas do not provide binlogs.

Binlogs can be retained from 0 to 168 (7 x 24) hours locally.

If the retention period is set to **0**, the binlogs of your DB instance will be deleted once they are synchronized to the standby instance and read replicas and successfully backed up to OBS. If the retention period is set to a value greater than 0, for example, 1 day, the binlogs will be retained for one day after they are synchronized to the standby instance and read replicas from the primary instance and successfully backed up to OBS. After the retention period expires, the binlogs will be automatically deleted. For details about how to view binlogs, see **Downloading a Binlog Backup File**.

Precautions

The binlog retention period is displayed in hour on the console. However, the value of **expire_logs_days** is displayed in day when you query the binlog retention period by running a command, which cannot be used as a reference. To check how long the binlogs can be retained, view the binlog retention period on the console.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target instance name to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Backups & Restorations**. On the **Binlog Backups** page, click **Set Binlog Retention Period**.

Step 6 In the displayed dialog box, set the local retention period and click **OK**.

Figure 7-9 Setting the binlog retention period

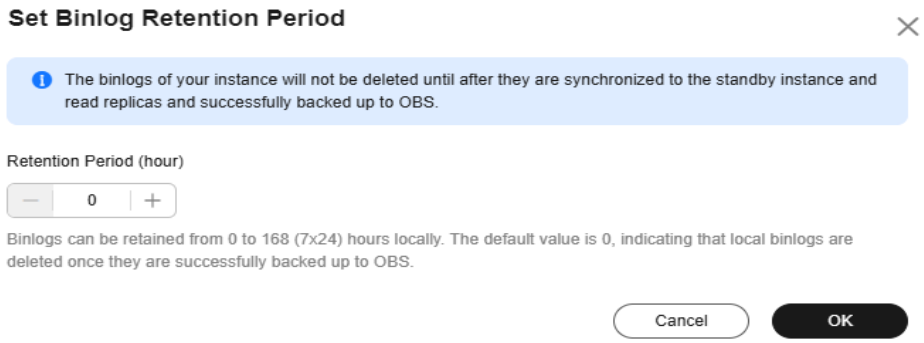

NOTE

When binlogs are deleted depends on the local retention period you configure on the console.

----End

8 Data Restorations

8.1 Restoration Solutions

If your database is damaged or mistakenly deleted, you can restore it from backups.

Table 8-1 Restoring a DB instance

When you...	Following Steps In...
Restore data to an RDS for MariaDB instance	**Restoring a DB Instance from a Backup**
	PITR: Restoring a DB Instance to a Point in Time

8.2 Restoring a DB Instance from a Backup

Scenarios

This section describes how to use an automated or manual backup to restore a DB instance to the status when the backup was created. The restoration is at the DB instance level.

When you restore a DB instance from a backup file, a full backup file is downloaded from OBS and then restored to the DB instance at an average speed of 40 MB/s.

Constraints

- You can restore to new DB instances from backups only when your account balance is greater than or equal to $0 USD. You will pay for the new instance specifications.
- If transparent page compression is enabled by specifying attributes in the CREATE TABLE statement for the original DB instance, the restoration may fail due to insufficient storage space.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Backups** page, select the backup to be restored and click **Restore** in the **Operation** column.

Alternatively, click the target DB instance on the **Instances** page. On the displayed page, choose **Backups & Restorations**. On the displayed page, select the backup to be restored and click **Restore** in the **Operation** column.

Step 5 Select a restoration method.

- Create New Instance

  Select **Create New Instance** for **Restoration Method** and click **OK**. The **Create New Instance** page is displayed.

  - The DB engine and version of the new DB instance are the same as those of the original DB instance and cannot be changed.
  - Storage space of the new DB instance is the same as that of the original DB instance by default and the new instance must be at least as large as the original DB instance.
  - Other settings are the same as those of the original DB instance by default and can be modified. For details, see Buying a DB Instance.

  Figure 8-1 Restoring to a new instance

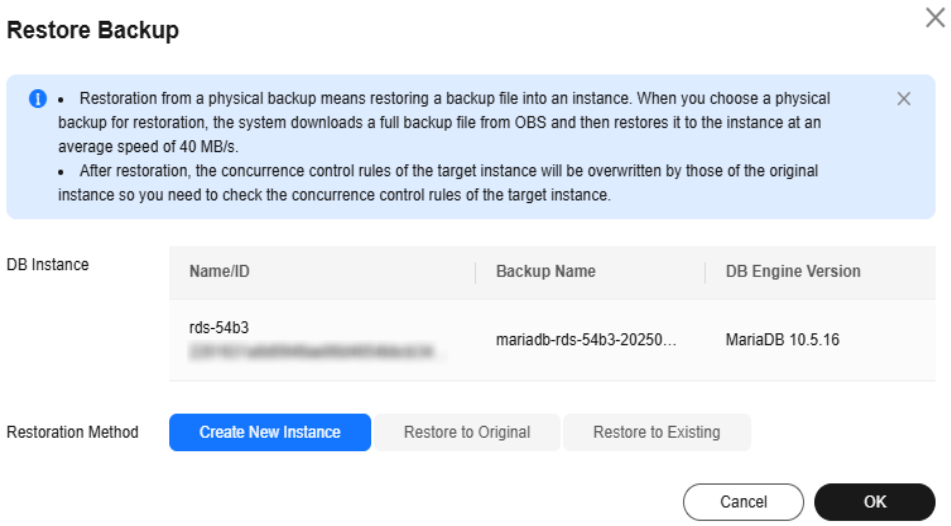

- Restore to Original

  a. Select "I acknowledge that after I select Restore to Original, data on the original databases will be overwritten and the original DB instance will be unavailable during the restoration." and click **Next**.

  b. Confirm the information and click **OK**.

NOTICE

- If the DB instance for which the backup is created has been deleted, data cannot be restored to the original DB instance.
- Restoring to the original DB instance will overwrite all existing data and the DB instance will be unavailable during the restoration process.

- Restore to Existing

  a. Select "I acknowledge that restoring to an existing DB instance will overwrite data on the instance and will cause the existing DB instance to be unavailable during the restoration. Only DB instances that can be used as target instances for the restoration are displayed here. Eligible instances must have the same DB engine type, version, and at least as much storage as the instance being restored."

  b. Select an existing instance and click **Next**.

  c. Confirm the information and click **OK**.

NOTICE

- If the target existing DB instance has been deleted, data cannot be restored to it.
- Restoring to an existing DB instance will overwrite data on it and cause the existing DB instance to be unavailable.
- To restore backup data to an existing DB instance, the selected DB instance must use the same DB engine and the same or a later version than the original DB instance.
- Ensure that the storage space of the selected existing DB instance is greater than or equal to the storage space of the original DB instance. Otherwise, data will not be restored.

Step 6 (**Optional**) If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

Step 7 View the restoration result. The result depends on which restoration method was selected:

NOTE

Restoring from backups does not affect the performance of original DB instances.

- Create New Instance

  A new DB instance is created using the backup data. The status of the DB instance changes from **Creating** to **Available**.

  The new DB instance is independent from the original one. If you need read replicas to offload read pressure, create one or more for the new DB instance.

After the new instance is created, a full backup will be automatically triggered.

- Restore to Original

  On the **Instances** page, the status of the original DB instance changes from **Restoring** to **Available**. If the original DB instance contains read replicas, the read replica status is the same as the original DB instance status.

  After the restoration is complete, a full backup will be automatically triggered.

- Restore to Existing

  On the **Instances** page, the status of the target existing DB instance changes from **Restoring** to **Available**. If the target existing DB instance contains read replicas, the read replica status is the same as the target existing DB instance status.

  After the restoration is complete, a full backup will be automatically triggered.

  You can view the detailed progress and result of the task on the **Task Center** page. For details, see **Viewing a Task**.

----End

FAQs

How Can I Restore Data If No Backup Is Available?

8.3 PITR: Restoring a DB Instance to a Point in Time

Scenarios

You can restore from automated backups to a specified point in time.

You can restore one or multiple DB instances at a time.

When you enter the time point that you want to restore the DB instance to, RDS downloads the most recent full backup file from OBS to the DB instance. Then, incremental backups are also restored to the specified point in time on the DB instance. Data is restored at an average speed of 30 MB/s.

Constraints

- You can restore to new DB instances from backups only when your account balance is greater than or equal to $0 USD. You will be billed for new instance specifications.
- Do not run the **reset master** command on RDS for MariaDB instances within their lifecycle. Otherwise, an exception may occur when restoring an RDS for MariaDB instance to a specified point in time.
- When you restore data to a new DB instance, large transactions in the original DB instance backup may cause a restoration failure. If the restoration fails, submit a service ticket.

Restoring a DB Instance

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target DB instance.

Step 5 In the navigation pane on the left, choose **Backups & Restorations**. On the displayed page, click **Restore to Point in Time**.

Step 6 Select the restoration date and time range, enter a time point within the selected time range, and select a restoration method. Then, click **OK**.

- Create New Instance

  The **Create New Instance** page is displayed.

  - The DB engine and version of the new DB instance are the same as those of the original DB instance and cannot be changed.
  - Other settings are the same as those of the original DB instance by default and can be modified. For details, see **Buying a DB Instance**.

- Restore to Original

  a. Select "I acknowledge that after I select Restore to Original, data on the original databases will be overwritten and the original DB instance will be unavailable during the restoration." and click **Next**.

  b. Confirm the information and click **OK**.

  NOTICE

  Restoring to the original DB instance will overwrite all existing data and the DB instance will be unavailable during the restoration process.

- Restore to Existing

  a. Select "I acknowledge that restoring to an existing DB instance will overwrite data on the instance and will cause the existing DB instance to be unavailable during the restoration. Only DB instances that can be used as target instances for the restoration are displayed here. Eligible instances must have the same DB engine type, version, and at least as much storage as the instance being restored."

  b. Select an existing instance and click **Next**.

  c. Confirm the information and click **OK**.

NOTICE

- Restoring to an existing DB instance will overwrite data on it and cause the existing DB instance to be unavailable.
- To restore backup data to an existing DB instance, the selected DB instance must use the same DB engine and the same or a later version than the original DB instance.
- Ensure that the storage space of the selected DB instance is greater than or equal to the storage space of the original DB instance. Otherwise, data will not be restored.

Step 7 (**Optional**) If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

Step 8 View the restoration result. The result depends on which restoration method was selected:

- Create New Instance

  A new DB instance is created using the backup data. The status of the DB instance changes from **Creating** to **Available**.

  The new DB instance is independent from the original one. If you need read replicas to offload read pressure, create one or more for the new DB instance.

  After the new DB instance is created, a full backup will be automatically triggered.

- Restore to Original

  On the **Instances** page, the status of the DB instance changes from **Restoring** to **Available**.

  A new restoration time range is available. There will be a difference between the new and original time ranges. This difference reflects the duration of the restoration.

  After the restoration is complete, a full backup will be automatically triggered.

- Restore to Existing

  On the **Instances** page, the status of the DB instance changes from **Restoring** to **Available**.

  You can view the detailed progress and result of the task on the **Task Center** page. For details, see **Viewing a Task**.

  After the restoration is complete, a full backup will be automatically triggered.

----End

FAQs

How Can I Restore Data If No Backup Is Available?

9 Read Replicas

9.1 Introduction to Read Replicas

Introduction

RDS for MariaDB supports read replicas.

In read-intensive scenarios, a single DB instance may be unable to handle the read pressure and workloads may be affected. To offload read pressure on the primary DB instance, you can create one or more read replicas in the same region as the primary instance. These read replicas can process a large number of read requests and increase application throughput.

A read replica uses a single-node architecture (without a standby node). Changes to the primary DB instance are also automatically synchronized to all associated read replicas through the native MariaDB replication function. The synchronization is not affected by network latency. Read replicas and the primary DB instance must be in the same region but can be in different AZs.

Applicable Scenario

In read-intensive scenarios, read replicas help offload read pressure from the primary instance.

Data is replicated from the primary instance to read replicas asynchronously. Although there is a replication delay, the data on read replicas will eventually be consistent with that on the primary instance. You can use read replicas if you do not mind such a replication delay.

Billing Standards

Read replicas are billed on a pay-per-use basis.

Functions

- Read replica specifications can be different from primary DB instance specifications. It is recommended that the read replica specifications be

greater than or equal to the primary DB instance specifications to prevent long delay and high load.

- Pay-per-use billing is supported. You only pay for what you use.
- Read replicas support system performance monitoring.

  RDS provides up to 20 monitoring metrics, including storage space, IOPS, database connections, CPU usage, and network traffic. You can view these metrics to learn about the load on read replicas.

Constraints

- A maximum of five read replicas can be created for a DB instance. To create more read replicas, **submit a service ticket** to request permissions. You can create up to 10 read replicas for each DB instance.
- You can purchase read replicas only for your created DB instance.
- All databases and tables in the primary instance are synchronized to read replicas. Data of the primary instance, standby instance, and read replicas is consistent.
- Read replicas do not support automated backups or manual backups.
- Read replicas do not support restoration from backups to new, existing, or original read replicas.
- Data cannot be migrated to read replicas.
- Read replicas do not support database creation and deletion.
- Read replicas do not support database account creation. Create database accounts on the primary DB instance. For details, see **Creating a Database Account**.
- Read replicas cannot be recycled after they are deleted.

Creating and Managing a Read Replica

- **Creating a Read Replica**
- **Creating Read Replicas in Batches**
- **Managing a Read Replica**

9.2 Creating a Read Replica

Scenarios

Read replicas enhance the read capabilities and reduce the load on your DB instances.

After an RDS instance is created, you can create read replicas for it as required.

Constraints

A maximum of five read replicas can be created for a DB instance. To create more read replicas, **submit a service ticket** to request permissions. You can create up to 10 read replicas for each DB instance.

For details about how to create read replicas in batches, see **Creating Read Replicas in Batches**.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and choose **More** > **Create Read Replica** in the **Operation** column.

Step 5 On the displayed page, configure required parameters and click **Next**.

Table 9-1 Basic information

Parameter	Description
Billing Mode	Yearly/monthly billing and pay-per-use billing are supported.
Region	By default, read replicas are in the same region as your DB instance.
DB Instance Name	Must start with a letter and consist of 4 to 64 characters. Only letters (case-sensitive), digits, hyphens (-), and underscores (_) are allowed.
DB Engine	Same as the DB engine of your DB instance by default and cannot be changed.
DB Engine Version	Same as the DB engine version of your DB instance by default and cannot be changed.
Storage Type	Determines the DB instance read/write speed. The higher the maximum throughput is, the higher the DB instance read/write speed can be. **Cloud SSD**: cloud disks used to decouple storage from compute. The maximum throughput is 350 MB/s.
AZ	RDS allows you to deploy your DB instance and read replicas in a single AZ or across AZs to improve reliability.

Table 9-2 Instance specifications

Parameter	Description
Instance Class	Refers to the vCPU and memory of a DB instance. Different instance classes support different numbers of database connections and maximum IOPS. After a DB instance is created, you can change its instance class. For details, see **Changing a DB Instance Class**.

Parameter	Description
Storage Space	Contains the system overhead required for inodes, reserved blocks, and database operation. By default, storage space of a read replica is the same as that of the primary DB instance.

Table 9-3 Network

Parameter	Description
VPC	Same as the primary DB instance's VPC.
Subnet	Same as the primary DB instance's subnet. A floating IP address is automatically assigned when you create a read replica. You can also enter an unused floating IP address in the subnet CIDR block. After the read replica is created, you can change the floating IP address.
Security Group	Same as the primary DB instance's security group.

Table 9-4 Enterprise project and tags

Parameter	Description
Enterprise Project	If your account has been associated with an enterprise project, select the target project from the **Enterprise Project** drop-down list. For more information about enterprise projects, see ***Enterprise Management User Guide***.
Tag	Optional. Tags help you easily identify and manage your read replicas. A maximum of 20 tags can be added for each read replica. After a read replica is created, you can view its tag details on the **Tags** page. For details, see **Managing Tags**.

Table 9-5 Yearly/monthly read replicas

Parameter	Description
Required Duration	The system will automatically calculate the configuration fee based on the selected required duration. The longer the required duration is, the larger discount you will enjoy.

Parameter	Description
Auto-renew	• By default, this option is not selected. • If you select this option, the auto-renewal cycle is determined by the selected required duration.

Step 6 Confirm specifications.

- If you need to modify your settings, click **Previous**.
- If you do not need to modify your settings, click **Submit** for pay-per-use read replicas.
- For yearly/monthly read replicas, click **Pay Now**.

Step 7 After a read replica is created, you can view and manage it.

You can view the detailed progress and result of the task on the **Task Center** page. For details, see **Task Center**.

----End

Follow-up Operations

Managing a Read Replica

9.3 Creating Read Replicas in Batches

Scenarios

Read replicas are used to enhance read capabilities and reduce the load on primary DB instances. On the **Instances** page, you can select one or more DB instances and create read replicas for them in batches.

NOTE

- To create read replicas in batches, **submit a service ticket** to apply for required permissions.
- A maximum of five read replicas can be created for a DB instance.
- You can create read replicas for a maximum of 50 DB instances at a time.
- Read replicas can be created in batches only for RDS for MariaDB instances running the same database version.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, select one or more DB instances and choose **More** > **Create Read Replica** above the instance list.

Step 5 On the displayed page, configure required information and click **Next**.

- By default, read replicas are named with "read" and two digits appended to the primary DB instance name. For example, if the primary instance name is instance-0001, the first read replica will be named instance-0001-read-01.
- The network and storage configurations are the same as those of the primary DB instance.
- In a batch creation, the number of read replicas you can create is limited by whichever DB instance already has the most replicas.

  For example, in a batch creation where most of the DB instances only have a single read replica, if any DB instance in the batch has more than one, for example, 3, you would only be able to add 2 more replicas for each DB instance in that particular batch operation.

Step 6 Confirm specifications.

- If you need to modify your settings, click **Previous**.
- If you do not need to modify your settings, click **Submit** for pay-per-use read replicas.
- For yearly/monthly read replicas, click **Pay Now**.

Step 7 After read replicas are created, you can view and manage them.

You can view the detailed progress and result of the task on the **Task Center** page. For details, see **Task Center**.

----End

Follow-up Operations

Managing a Read Replica

9.4 Managing a Read Replica

Entering the Management Interface Through a Read Replica

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click  in front of the DB instance and click the target read replica to go to the **Overview** page.

----End

Entering the Management Interface Through DAS

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and click in front of it. In the expanded panel, locate the read replica you want to manage and click **Log In** in the **Operation** column.

Step 5 On the displayed page, enter the username and password and click **Log In**.

----End

Deleting a Read Replica

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click in front of a DB instance, locate the read replica to be deleted, and choose **More** > **Delete** in the **Operation** column.

----End

10 Problem Diagnosis and SQL Analysis

10.1 Function Overview

DBA Assistant provides visualized database O&M and intelligent diagnosis for developers and database administrators (DBAs), making database O&M easy and efficient. By analyzing alarms, resources, health data, performance metrics, and storage usage, it helps users quickly locate faults and keep track of instance status.

Scenarios

- Setting a slow session threshold can help you quickly identify abnormal sessions and kill the sessions when an exception occurs in your instance, so that your instance can recover quickly and ensure database availability.
- If your DB instance is unstable due to a large number of concurrent SQL requests from new services, you can set SQL throttling rules for SQL statements to limit concurrent SQL statements and ensure instance stability.
- If your instance storage is full, you can learn about the storage usage and disk space distribution on the **Storage Analysis** page. Autoscaling is available for you to enable. After this function is enabled, the storage space is automatically scaled up when the storage space is too small.
- You can configure auto throttling to limit active connections in high burst traffic or abnormal read/write scenarios to ensure the availability of core workloads.

Functions

Table 10-1 lists the functions supported by DBA Assistant.

Table 10-1 Function description

Function	Description	Reference
Dashboard	Shows the status of your instance, including alarms, resource usages, and key performance metrics. DBA Assistant diagnoses instance health using operational data analytics and intelligent algorithms, and provides you with solutions and suggestions for handling detected exceptions.	**Viewing the Overall Status of a DB Instance**
Sessions	Displays a list of sessions and allows you to filter sessions. You can set a slow session threshold to identify abnormal sessions for urgent instance recovery, ensuring database availability.	• **Viewing Session Statistics** • **Setting a Slow Session Threshold**
Performance	Displays key metrics of your instance and provides metric comparison between different days. You can keep track of metric changes and detect exceptions in a timely manner.	**Viewing Performance Metrics**
Storage Analysis	Storage occupied by data and logs and historical changes of storage usage are important for database performance. The **Storage Analysis** page displays storage overview and disk space distribution of your instance. In addition, DBA Assistant can estimate the available days of your storage based on historical data and intelligent algorithms, so that you can scale up storage in a timely manner.	• **Viewing Storage Usage** • **Viewing Table Diagnosis Results** • **Setting a Diagnosis Threshold** • **Viewing Top Databases and Tables by Physical File Size**
Slow Query Logs	Displays slow queries within a specified time period. You can view top 5 slow query logs by user or client IP address, sort statistics, and identify sources of slow SQL statements.	**Viewing Slow Query Logs**
SQL Throttling	Restricts the execution of concurrent SQL statements based on specified rules when there are SQL statements that cannot be optimized timely or a resource (for example, vCPU) bottleneck occurs.	**Creating a SQL Throttling Rule**

Function	Description	Reference
Auto Throttling	Automatically detects database exceptions such as high vCPU usage and excessive active sessions, and limits traffic based on specified priorities. You can control traffic by database or user as required. Limiting traffic of non-core databases or from non-core users can ensure that core workloads remain stable.	**Configuring Auto Throttling**
Daily Reports	Provides overall information about your instance status of the previous day, including slow SQL analysis and performance & storage analysis. You can download and subscribe to analysis reports. A daily diagnosis is recommended.	• **Viewing Diagnosis Reports** • **Subscribing to Diagnosis Reports**

10.2 Performance Monitoring

10.2.1 Viewing the Overall Status of a DB Instance

On the **Overview** page, you can get knowledge of the overall status of your RDS for MariaDB instance, including alarms, intelligent anomaly diagnosis, and key performance metrics. DBA Assistant diagnoses instance health using operational data analytics and intelligent algorithms, and provides you with solutions and suggestions for handling detected exceptions.

Functions

Table 10-2 lists the functions provided on the **Overview** page.

Table 10-2 Function description

Function	Description
Alarms	To view alarm details, click the number next to an alarm severity.
Intelligent Anomaly Diagnosis	Shows the health status of your instance based on operational data analytics and intelligent algorithms.
Performance Monitoring	Shows key performance metrics of the instance, including the CPU usage, memory usage, slow SQL queries, and connections.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the DB instance name.

Step 5 On the **Overview** page, view the status of your instance.

- In the **Alarms** area, view alarm information of your instance.

  To view the list of all alarms, click **All Alarms**. To view alarm details, click the number next to an alarm severity.

  Figure 10-1 Alarms

  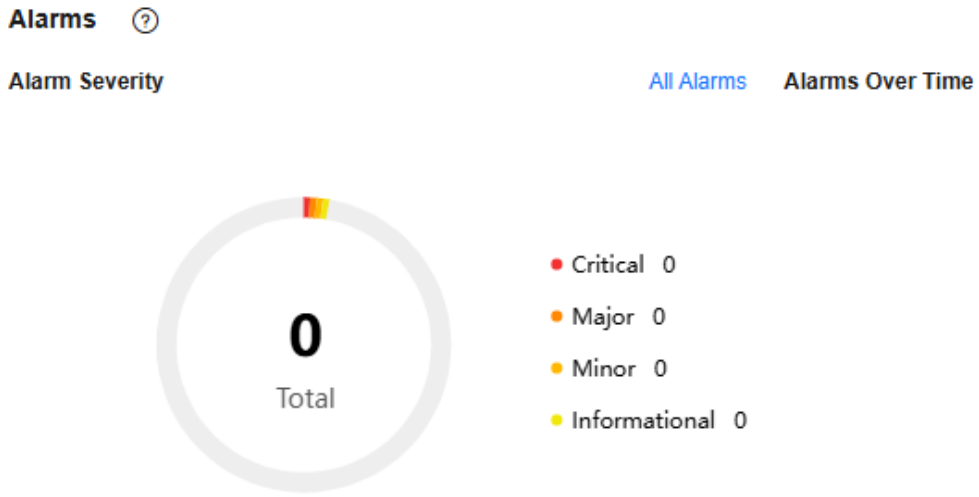

- In the **Intelligent Anomaly Diagnosis** area, view the health diagnosis results of your instance.

  Figure 10-2 Intelligent Anomaly Diagnosis

  

- In the **Performance Monitoring** area, view key performance metrics of your instance.

  Figure 10-3 Performance Monitoring

  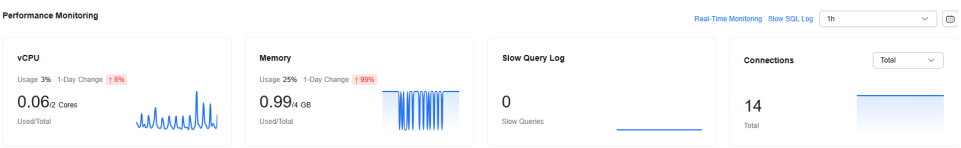

----End

10.2.2 Viewing Performance Metrics

DBA Assistant allows you to view the performance metrics of your DB instance. Historical trends of performance metrics within a specified time period help you learn about the status and resource usage of your DB instance. If any alarm is reported, you can take actions timely.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the DB instance name.

Step 5 In the navigation pane, choose **DBA Assistant** > **Real-Time Diagnosis**.

Step 6 Click the **Performance** tab to view the performance metrics of your instance.

Figure 10-4 Performance metrics

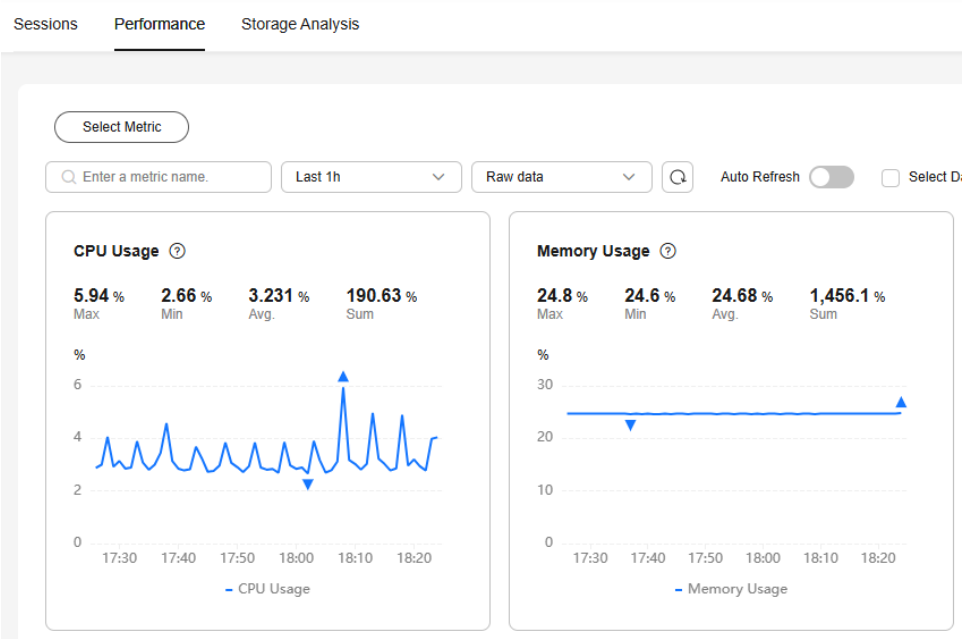

----End

10.3 Problem Diagnosis

10.3.1 Managing Real-Time Sessions

10.3.1.1 Viewing Session Statistics

DBA Assistant allows you to view session statistics of your instance, including slow sessions, active sessions, and total sessions, helping you learn about the distribution of sessions in different dimensions.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **DBA Assistant** > **Real-Time Diagnosis**.

Step 6 Click the **Sessions** tab. You can view slow sessions, active sessions, and total sessions by the following three dimensions:

- User
- Access host
- Database

Figure 10-5 Session statistics

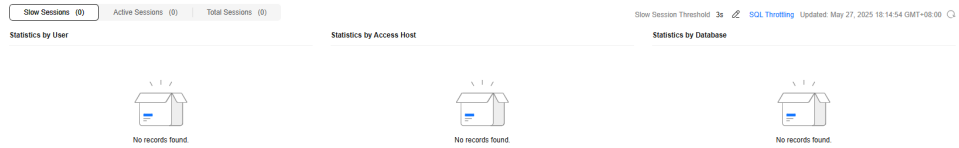

----End

10.3.1.2 Setting a Slow Session Threshold

You can set a slow session threshold to identify sessions whose execution time is longer than the threshold. This allows you to identify abnormal sessions and kill the sessions to restore the database.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **DBA Assistant** > **Real-Time Diagnosis**.

Step 6 Click the **Sessions** tab.

Step 7 Click next to the **Slow Session Threshold** field. In the displayed dialog box, set a slow session threshold and click **OK**. Sessions whose execution durations are greater than this threshold are automatically displayed.

Max. Execution Time for a Query (s) indicates how long a session has been executed before it can be considered as a slow session. The default value is **3**. The value ranges from 1 to 86,400, in seconds.

Figure 10-6 Setting a slow session threshold

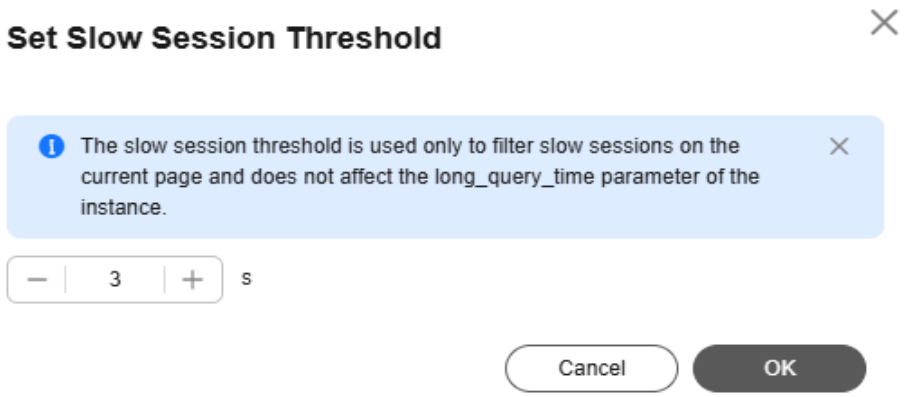

Step 8 In the **Sessions** area, select the target session IDs based on the instance status and service requirements, and click **Kill Session**.

NOTE

The **Kill Session** operation ends only the selected sessions. It does not affect other sessions or services.

Step 9 Click **OK**.

----End

10.3.2 Viewing Storage Usage

DBA Assistant allows you to view the storage usage of your DB instance in real time to prevent full storage space.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **DBA Assistant** > **Real-Time Diagnosis**.

Step 6 Click the **Storage Analysis** tab.

- In the **Overview** area, view the storage usage, including the available storage space and total storage space.

  If the storage usage reaches 87% or higher, you can click **Scale** to scale up the storage. For details about constraints and billing, see **Scaling Up Storage Space**.

  Figure 10-7 Storage usage

  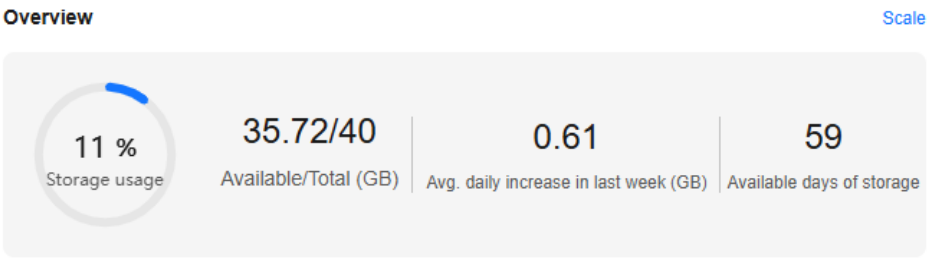

  NOTE

  If the average daily increase in last week is 0 GB, the estimated available days of storage are unlimited and are not displayed.

- In the **Disk Space Distribution** area, view the space distribution of your instance. For details, see **Table 10-3**.

  Figure 10-8 Disk space distribution

  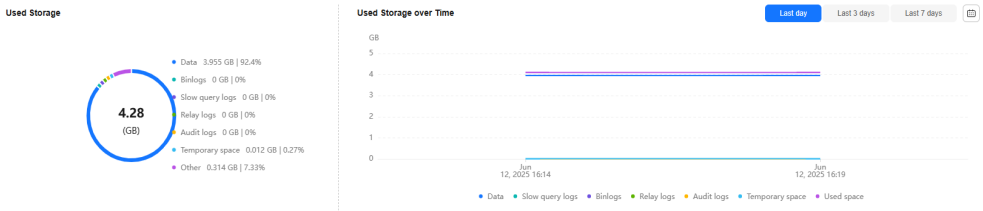

  NOTE

  If the total number of files in your disk space (including data space, binlog space, slow query log space, relay log space, audit log space, temporary space, and other space) exceeds 10,000, RDS will not collect information about the files or display disk space distribution and usages over time on the console. This prevents performance slowdowns caused by collecting statistics on too many files. If this happens, submit a service ticket.

Table 10-3 Parameter description

Parameter	Description
Data	Disk space for storing user data

Parameter	Description
Binlogs	Disk space for storing binlogs
Slow query logs	Disk space for storing slow logs
Relay logs	Disk space for storing relay logs
Audit logs	Disk space for storing audit logs
Temporary space	Disk space for storing temporary files
Other	Disk space for storing files such as **ib_buffer_pool**, **ib_doublewrite** and **error.log** generated by the instance.

----End

FAQ

Q: What can I do if the storage space of my DB instance is full?

A: Reduce the storage usage to below 87% so that the DB instance becomes available and data can be written to the instance. You can use either of the following methods to reduce the storage usage:

- Scale up the storage space: **Services are not interrupted during storage scale-up**. You can also enable autoscaling. When the available storage of a DB instance drops to the threshold, autoscaling is triggered.
- Reduce data: Delete useless historical data.
  a. If your instance becomes read-only, you need to submit a service ticket to cancel the read-only status first. If your instance is not in the read-only state, you can delete data directly.
  b. Check the top 50 databases and tables with large physical files and identify the historical table data that can be deleted. For details, see **Viewing Top Databases and Tables by Physical File Size**.
  c. To clear up space, you can optimize tables with a high fragmentation rate during off-peak hours.

     To delete data of an entire table, run **DROP** or **TRUNCATE**. To delete part of table data, run **DELETE** and **OPTIMIZE TABLE**.
- If temporary files generated by sorting queries occupy too much storage space, optimize your SQL query statements.

  You can query **slow query logs**, and analyze and optimize the problematic SQL statements.

10.3.3 Viewing Table Diagnosis Results

Intelligent table diagnosis can diagnose tables with abnormal tablespace growth, tables without primary keys, and tables without indexes, helping you quickly locate abnormal tables.

To use this function, you need to subscribe to Intelligent O&M first. For details, see **Subscribing to Intelligent O&M**.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **DBA Assistant** > **Real-Time Diagnosis**.

Step 6 Click the **Storage Analysis** tab to view the table diagnosis results of your instance.

Figure 10-9 Table diagnosis results

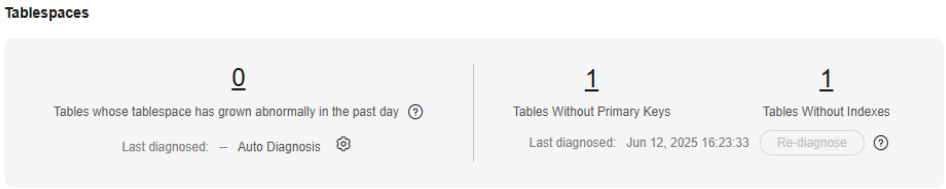

----End

10.3.4 Viewing Top Databases and Tables by Physical File Size

In combination with disk space distribution, top 50 databases and tables by physical file size help you identify the databases and tables with high disk space usage.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **DBA Assistant** > **Real-Time Diagnosis**.

Step 6 Click the **Storage Analysis** tab and view the top 50 databases and tables in the **Top Databases and Tables** area.

NOTE

- Physical file sizes are precisely recorded, but other fields' values are estimated. If there is a large gap between a file size and another field, run ANALYZE TABLE on the table.
- A database or table whose name contains reserved special characters, including slashes (/) and #p#p, is not counted. #p#p identifies the table as a partitioned table.
- If there are more than 50,000 tables in your instance, to prevent data collection from affecting the instance performance, top databases and tables will not be counted.

Step 7 Locate a database and click **View Chart** in the **Operation** column. You can view data volume changes in the last 7 days, last 30 days, or a custom time period (spanning no more than 30 days).

----End

10.3.5 Setting a Diagnosis Threshold

You can set a diagnosis threshold to identify abnormal tables whose tablespace is above the threshold.

To use this function, you need to subscribe to Intelligent O&M first. For details, see **Subscribing to Intelligent O&M**.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **DBA Assistant** > **Real-Time Diagnosis**.

Step 6 Click the **Storage Analysis** tab to view the table diagnosis results of your instance.

Step 7 Click next to **Auto Diagnosis**. In the displayed dialog box, configure a daily tablespace increase limit. The value ranges from 1 to 100,000,000, in MB.

Figure 10-10 Configuring a daily tablespace increase limit

Step 8 Click **OK**. When the size of a single table is greater than the threshold, the system automatically filters the table and collects statistics on the number of such tables on the **Tablespaces** page. The system filters tables once a day.

Step 9 Click a number and view the details about abnormal tables on the **Diagnosis Details** page.

----End

FAQ

- Q: What can I do if there are tables whose tablespace has grown abnormally in the past day?

  A: Check tablespace fragmentation and reclaim fragmented space. Do not use **DELETE** to clear data. If you have any other questions, submit a service ticket.
- Q: What is the impact of tables without primary keys on my DB instance?

  A: Tables without primary keys can cause slow SQL statements, affecting instance stability. You are advised to add primary keys to such tables to reduce the primary/standby replication delay.
- Q: What is the impact of tables without indexes on my DB instance?

  A: Tables without indexes can cause slow SQL statements, affecting instance stability. You are advised to add indexes to table fields for more efficient query.

10.3.6 Managing Diagnosis Reports

10.3.6.1 Viewing Diagnosis Reports

You can start a health diagnosis on your DB instance and view the current and historical diagnosis reports.

Procedure

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **DBA Assistant** > **Historical Diagnosis**.

Step 6 Click the **Daily Reports** tab.

Step 7 Click **Start Diagnosis** and select a time range for the diagnosis.

Step 8 Click **OK**. You can also view historical diagnosis reports or download a report to your local PC.

- To view historical diagnosis reports, click **View History** in the upper right corner of the page.

- To download a report to your local PC, click **Download** in the upper right corner of the page.

Figure 10-11 Daily reports

Step 9 In the **Diagnosis Dimensions** area, click **Slow SQL Analysis** or **Performance & Storage** to view details.

Figure 10-12 Viewing report analysis details

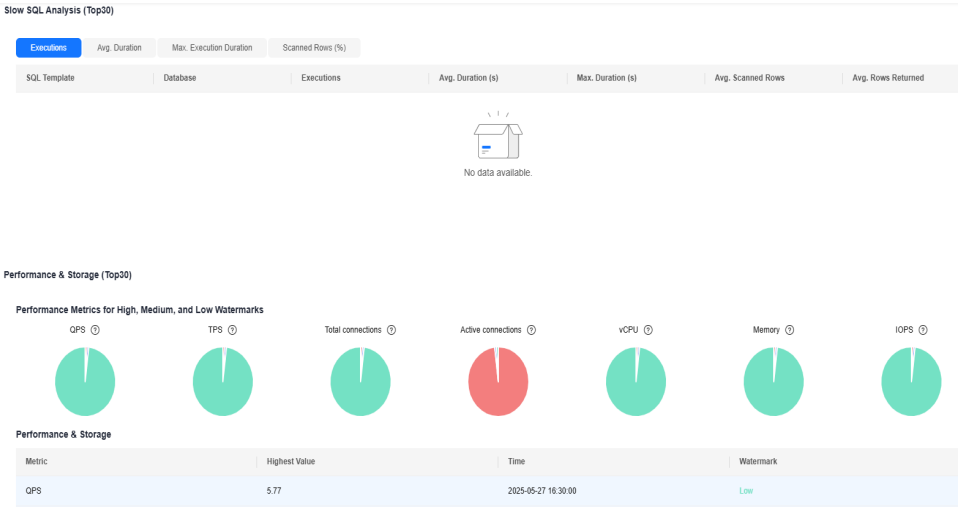

----End

10.3.6.2 Subscribing to Diagnosis Reports

After you subscribe to diagnosis reports, Simple Message Notification (SMN) will send diagnosis exception reports to the preset email address so that you can learn about the overall health status of your DB instance in real time.

Billing

When you use SMN, only pay for what you use. There are no minimum fees. For details, see **Billing**.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **DBA Assistant** > **Historical Diagnosis**.

Step 6 Click the **Daily Reports** tab.

Step 7 In the upper right corner of the page, click **Subscribe** and set subscription parameters. For details about the parameters, see **Table 10-4**.

Table 10-4 Subscription parameters

Parameter	Description
Subscription	Select **By topic** or **By email**.
Topics	A topic is used to publish messages and subscribe to notifications. It serves as a message transmission channel between publishers and subscribers. If there are no topics you want to select, **create one**. After a topic is created, click **Add Subscription** in the **Operation** column of the topic. In the displayed dialog box, specify a protocol (only **Email** is supported) and an endpoint.
Email Addresses	If you select **By email** for **Subscription**, you need to specify **Email Addresses**.

Step 8 Click **OK**.

----End

Related Operations

If you want to unsubscribe from diagnosis reports, click **Unsubscribe** in the upper right corner of the page. In the displayed dialog box, confirm the information and click **OK**.

10.3.7 Subscribing to Intelligent O&M

To use the **Tablespaces**, **Slow Query Logs**, and **Auto Throttling** functions, you need to subscribe to Intelligent O&M first. This section describes how to subscribe to Intelligent O&M.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click ⌖ in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **DBA Assistant** > **Real-Time Diagnosis**.

Step 6 Click the **Storage Analysis** tab.

Figure 10-13 Tablespaces

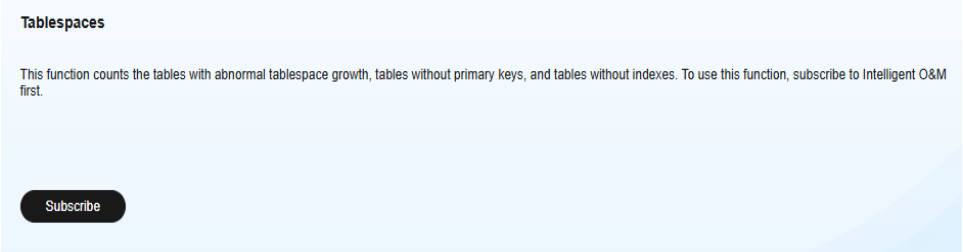

Step 7 Click **Subscribe**. In the displayed dialog box, you can learn about Intelligent O&M functions and pricing.

If you select **Collect Slow Query Logs**, slow SQL statements will be collected and analyzed. For details, see **Viewing Slow Query Logs**.

Figure 10-14 Subscribing to Intelligent O&M

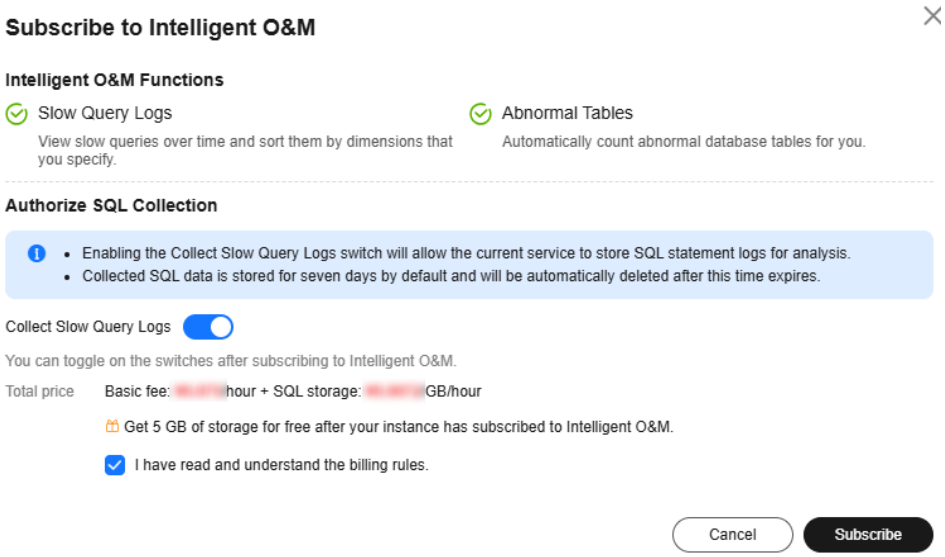

Step 8 Select "I have read and understand the billing rules." and click **Subscribe**.

----End

10.4 SQL Analysis

10.4.1 Viewing Slow Query Logs

Slow query logs help you locate SQL statements that process a large amount of data, scan a large number of rows, or run for a long time, so that you can optimize them to improve database performance.

To use this function, you need to subscribe to Intelligent O&M first. For details, see **Subscribing to Intelligent O&M**.

If you did not subscribe to Intelligent O&M, you can view only the data of the last hour. The data will be automatically deleted when it expires.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **DBA Assistant** > **Historical Diagnosis**.

Step 6 Click the **Slow Query Logs** tab.

Step 7 In the **Slow Queries over Time** area, you can view the slow query log and vCPU usage trends of your DB instance.

You can view the **Slow Queries over Time** chart in the last 1 hour, last 3 hours, last 12 hours, or a custom time period (spanning no more than one day).

Figure 10-15 Slow Queries over Time

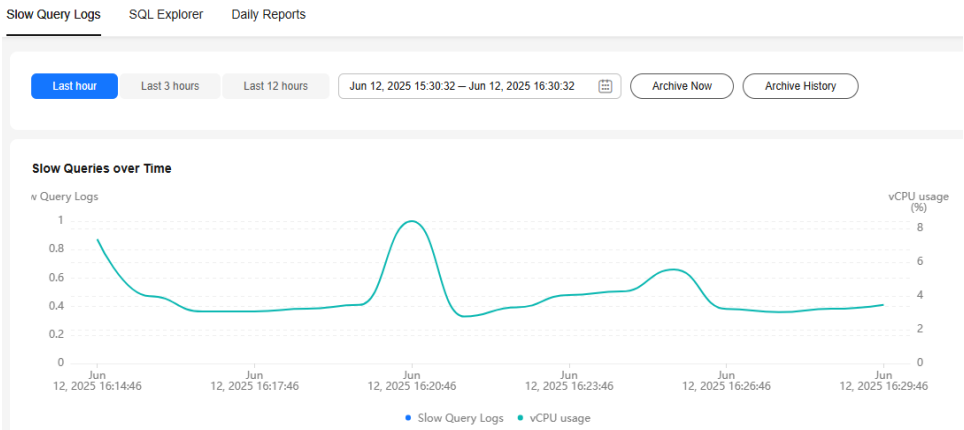

Step 8 Click **Archive History**. In the archived slow query logs list, view slow query log details.

The displayed details include the log start time, end time, and log file size.

Figure 10-16 Archived Slow Query Logs

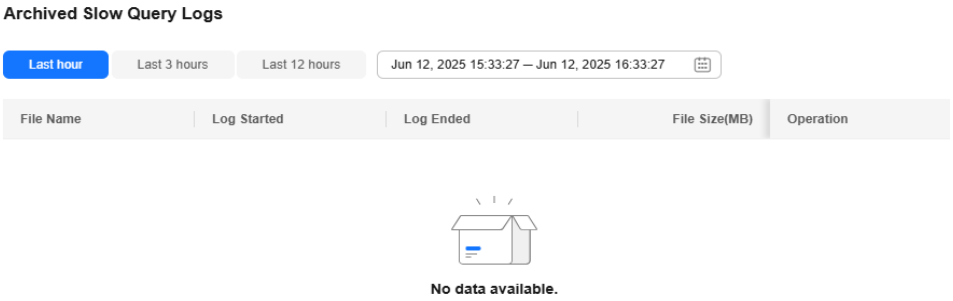

NOTE

Slow query logs are automatically archived every 3 minutes. To view the latest slow query logs, click **Archive Now**.

Step 9 In **Top 5 Slow Query Logs** area, view the top 5 slow SQL statements sorted by user or client IP address.

Figure 10-17 Top 5 Slow Query Logs

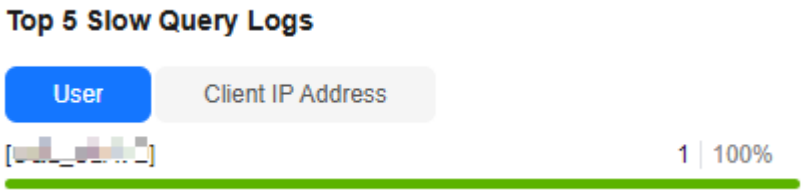

Step 10 In the details list, view details about slow query logs.

- You can filter slow query logs by SQL statement, database, client IP address, user, execution duration, and scanned rows.
- To export the slow query log list, click **Export**.
- To view log export history, click **View Export List**.

----End

10.4.2 Creating a SQL Throttling Rule

You can create rules to control concurrent execution of SQL statements by specifying SQL type, keywords, and maximum concurrency. To maintain better performance at high concurrency, SQL statements that meet the specified SQL type and keyword and exceed the maximum concurrency will not be executed.

Constraints

- The rule you are creating will be applied only to the current instance.
- If a SQL statement matches multiple SQL throttling rules, only the most recently added rule is applied.
- SQL statements that have been executed before a SQL throttling rule is added are not counted.

- If the replication delay is too long, adding or deleting a SQL throttling rule for a read replica does not take effect immediately.
- This function is only available for MariaDB 10.5. You are advised to upgrade the minor version of your instance to the latest.
- Too many SQL throttling rules affect the database performance. Delete unnecessary rules after using them.
- This function controls how many statements can run at the same time. However, it does not limit concurrency for:
  - System tables
  - Queries where no database data is involved, such as **SELECT sleep(*xxx*);**
  - Account **root**
  - SQL statements in stored procedures, triggers, and functions

Procedure

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **DBA Assistant** > **Historical Diagnosis**.

Step 6 Choose **SQL Explorer** > **SQL Throttling**.

Step 7 Click **Add Rule**. Configure the parameters listed in **Table 10-5**.

Figure 10-18 Adding a rule

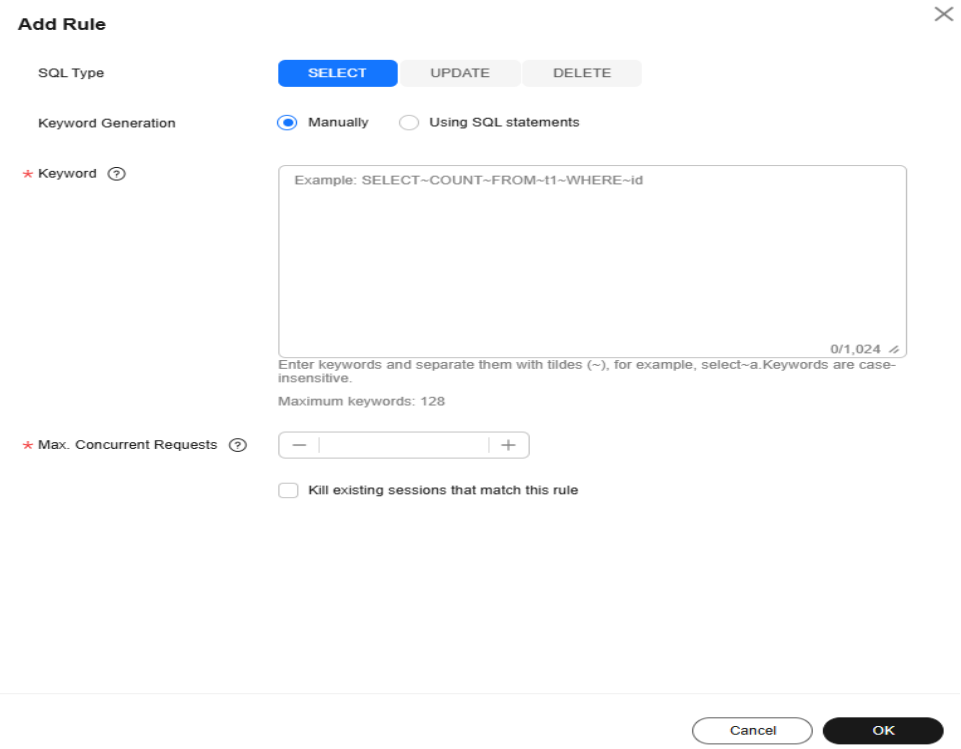

Table 10-5 Parameter description

Parameter	Description
SQL Type	There are three options: **SELECT**, **UPDATE**, and **DELETE**.
Keyword Generation	• **Manually**: Take **select~a** as an example. **select** and **a** are two keywords contained in a SQL throttling rule. The keywords are separated by a tilde (~). In this example, the rule restricts the execution of only the SQL statements containing keywords **select** and **a**. • **Generate keywords from a SQL statement**: You can enter a SQL statement and then click **Generate Keyword**. The generated keywords are for reference only. Exercise caution when using them.
Keyword	A maximum of 128 keywords (case-insensitive) are supported. SQL statements match the keywords from first to last. For example, if one rule contains the keyword **a~and~b**, the statement **xxx a>1 and b>2** can match the keyword, but **xxx b>2 and a>1** cannot.
Max. Concurrent Requests	If the number of concurrent SQL statements matching the keyword exceeds this limit, the SQL statements will not be executed. The value ranges from 0 to 1,000,000,000.

Parameter	Description
Kill existing sessions that match this rule	Selecting this option will not kill the connection sessions of user **root**.

Step 8 Confirm the settings and click **OK**.

Step 9 Toggle on the SQL throttling switch .

NOTE

SQL throttling rules take effect only after SQL throttling is enabled.

----End

Related Operations

To delete a SQL throttling rule, locate it in the rule list and click **Delete** in the **Operation** column. In the displayed dialog box, click **OK**.

Figure 10-19 Deleting a rule

10.4.3 Configuring Auto Throttling

Auto throttling helps ensure availability of your workloads. It restricts the execution of SQL statements during peak hours or when there are read/write exceptions by limiting how many SQL statements can be executed at the same time.

After it is enabled, the system performs flow control on sessions when the criteria you specify for your instance are met (for example, the number of active connections to your instance exceeds the **Max. Active Connections** parameter value).

To use this function, **subscribe to Intelligent O&M** first.

Constraints

To use this function, **submit a service ticket** to request permissions.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **DBA Assistant** > **Historical Diagnosis**.

Step 6 Choose **SQL Explorer** > **Auto Throttling**.

Step 7 Click **Auto Throttling**.

Step 8 Toggle on and configure required parameters. For details about the parameters, see **Table 10-6**.

Example for setting auto throttling parameters

Set **Time Window** to **15:00-18:00**, **Max. Duration** to **Last 5 minutes**, **vCPU usage ≥** to **90%**, **Active sessions ≥** to **20**, and **and Duration (min)** to **5**. When all the criteria are met, auto throttling is triggered. If your vCPU usage or number of active sessions falls below the threshold during the time window, auto throttling ends.

Figure 10-20 Example

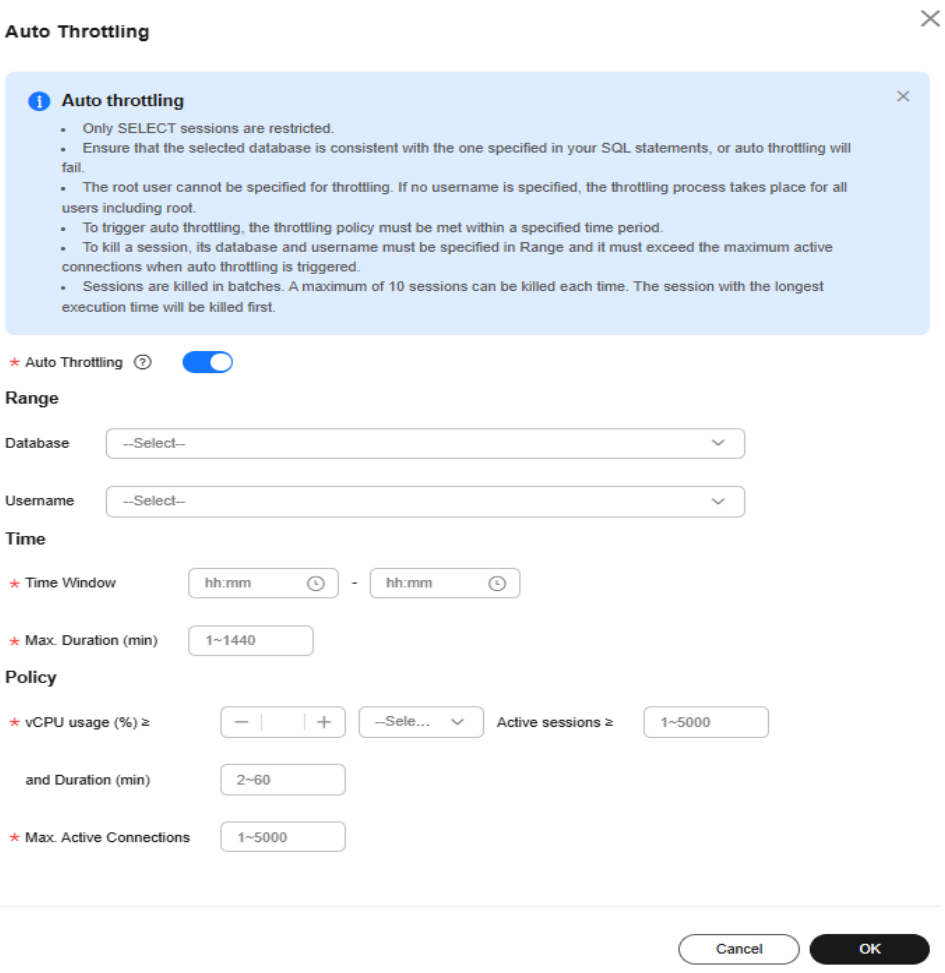

Table 10-6 Parameter description

Parameter	Description
Database	The name of the database for which auto throttling needs to be enabled. Ensure that the database you select is consistent with the one specified in the **use *<database>*** statement, or auto throttling will fail.
Username	The name of the user that auto throttling is applied. If no username is specified, auto throttling will be applied to user **root**.
Time Window	The time when auto throttling is applied. Auto throttling can be triggered only once within the time window.
Max. Duration	Maximum length of time that SQL statements matching the auto throttling policy can be throttled within the time window.
vCPU usage	vCPU usage threshold for the instance. You also need to specify the relationship between the vCPU usage and active sessions. Their relationship can be **and** or **or**.
Active sessions	Threshold for active sessions. Value range: 1 to 5000
Duration (min)	How long the vCPU usage and active sessions exceed the specified values. For example, if you set **vCPU usage ≥** to **90%**, **Active sessions ≥** to **1000**, and **Duration (min)** to **30**, auto throttling will be triggered only when the vCPU usage and active sessions exceed 90% and 1,000 for 30 minutes.
Max. Active Connections	Maximum number of active connections allowed. Value range: 1 to 5000 For example, if you set **Max. Active Connections** to **500**, the system will automatically end some active connections when necessary to keep the number of active sessions within 500.

Step 9 Click **OK**. You will see a record generated on the page every time auto throttling is triggered. You can see historical details, too.

----End

10.5 Common Performance Problems

10.5.1 How Do I Improve the Query Speed of My RDS Database?

The following are some suggestions provided for you to improve the database query speed:

- View the slow query logs to check if there are any slow queries, and view their performance characteristics to locate the cause. For details about how to view RDS for MariaDB logs, see **Viewing and Downloading Slow Query Logs**.
- View the CPU usage of your DB instance to facilitate troubleshooting. For details, see **Viewing Monitoring Metrics**.
- Create read replicas to offload read pressure on the primary DB instance.
- Increase the CPU or memory specifications for DB instances with high loads. For details, see **Changing a DB Instance Class**.
- Add indexes for associated fields in multi-table association queries.
- Specify a field or add a WHERE clause, which will prevent full table scanning triggered by the SELECT statement.

10.5.2 Identifying Why CPU Usage of RDS for MariaDB Instances Is High and Providing Solutions

If the CPU usage of your RDS for MariaDB instance is high or close to 100%, database performance deteriorates. For example, data read/write becomes slow, connecting to the instance takes a longer time, or errors are reported when you are trying to delete data.

Solution

Analyze slow SQL logs and CPU usage to locate slow queries and then optimize them.

1. View the slow SQL logs to check for slowly executed SQL queries and view their performance characteristics (if any) to locate the cause.

   For details about how to view RDS for MariaDB logs, see **Viewing and Downloading Slow Query Logs**.
2. View the CPU usage of your DB instance.

   For details, see **Viewing Monitoring Metrics**.
3. Create read replicas to reduce read pressure from primary DB instances.
4. Add indexes for associated fields in multi-table association queries.
5. Do not use the SELECT statement to scan all tables. You can specify fields or add the WHERE condition.

10.5.3 RDS for MariaDB Memory Usage Too High

For a DB instance storing mission-critical application data

Scale up the instance class.

For a DB instance not storing mission-critical application data

Check the memory usage of the local computer. If the memory usage curve is stable, no action is required.

For a DB instance storing mission-critical application data and configured with a large instance class

1. During off-peak hours, change the value of **performance_schema** to **OFF**. You need to reboot the instance for the change to take effect.
2. View the memory usage of your instance using DBA Assistant. For details, see **Viewing Performance Metrics**.

   If the space usage remains high, perform either of the following operations:

   - Scale up the instance class.
   - Change the **innodb_buffer_pool_size** value:
     - If the instance memory is 2 GB, change **innodb_buffer_pool_size** to **268,435,456** in byte (256 MB).
     - If the instance memory is 4 GB, change **innodb_buffer_pool_size** to **1,073,741,824** in byte (1 GB).
     - If the instance memory is 8 GB, change **innodb_buffer_pool_size** to **3,221,225,472** in byte (3 GB).
     - If the instance memory is greater than 8 GB, you do not need to adjust the **innodb_buffer_pool_size** value.

NOTICE

- Change the value of **innodb_buffer_pool_size** as needed.
- MariaDB has a dynamic memory balancing mechanism. If the memory usage is less than 90%, no action is required.
- RDS for MariaDB memory is allocated to the engine layer and server layer.
  - The memory allocated to the engine layer includes the InnoDB buffer pool, log buffer, and full text index cache. The InnoDB buffer pool is resident memory and accounts for a large proportion.

    The InnoDB buffer pool is a memory area that holds cached InnoDB data for tables, indexes, and other auxiliary buffers. You can use the **innodb_buffer_pool_size** parameter to define the buffer pool size.
  - The memory allocated to the server layer is occupied by the thread cache, binlog cache, sort buffer, read buffer, and join buffer. These caches and buffers are usually released when connections are closed.

  Such memory allocation keeps memory usage of a running RDS for MariaDB instance at about 80%.

10.5.4 What Should I Do If an RDS DB Instance Is Abnormal Due to Full Storage Space?

You can scale up storage space if it is no longer sufficient for your requirements. If the DB instance status is **Storage full** and no more data can be written to the database, the DB instance will be abnormal.

Solution

1. As your application data grows, the original storage space may be insufficient. You are advised to scale up storage space by referring to **Scaling Up Storage Space**.

   You can view the memory usage of your instance using DBA Assistant. For details, see **Viewing Storage Usage**.

   If the storage capacity has reached the upper limit of your DB instance class, change the instance class first.

   For details, see **Changing a DB Instance Class**.
2. Delete expired data files in a timely manner.
3. View performance metrics of your DB instance on the console, such as CPU, memory, storage, and connections. You can also set alarm rules for metric thresholds to identify risks in advance.

   For details, see **Viewing the Overall Status of a DB Instance**.

10.5.5 Troubleshooting Slow SQL Issues for RDS for MariaDB Instances

This section describes how to troubleshoot slow SQL statements on RDS for MariaDB instances. For any given service scenario, query efficiency depends on the architecture and on the database table and index design. Poorly designed architecture and indexes will cause many slow SQL statements.

Slow SQL Statements Caused by SQL Exceptions

- Causes and symptoms

  There are many causes for SQL exceptions, for example, unsuitable database table structure design, missing indexes, or too many rows that need to be scanned.

  On the **Slow Query Logs** page, you can download logs to identify the slow SQL statements and see how long they took to execute. For details, see **Viewing and Downloading Slow Query Logs**.
- Solution

  Optimize the SQL statements that you need to execute.

Slow SQL Statements Caused by DB Instance Limits

- Causes and symptoms

  DB instance performance can be limited because:

  – Your workloads have been increasing but the storage has not been scaled up accordingly.

  – The performance of your DB instance has been deteriorating as the physical server of the instance ages.

  – The amount of data has been increasing, and the data structure has been changing.

  You can view the resource usage of the DB instance on the console. If the values of all resource usage metrics are close to 100%, your DB instance may

reach its maximum performance. For details, see **Viewing the Overall Status of a DB Instance**.

- Solution

  Upgrade the instance class. For details, see **Changing a DB Instance Class**.

Slow SQL Statements Caused by Inappropriate Parameter Settings

- Causes and symptoms

  Inappropriate settings of some parameters (such as **innodb_spin_wait_delay**) can impact performance.

  You can view parameter modifications on the console. For details, see **Viewing Parameter Change History**.

- Solution

  Modify related parameters based on your specific service scenario.

Slow SQL Statements Caused by Batch Operations

- Causes and symptoms

  A large number of operations are performed to import, delete, and query data.

  You can view **Total Storage Space**, **Storage Space Usage**, and **IOPS** on the console. For details, see **Viewing the Overall Status of a DB Instance**.

- Solution

  Perform batch operations during off-peak hours, or split them.

Slow SQL Statements Caused by Scheduled Tasks

- Causes and symptoms

  If the load of your DB instance changes regularly over time, there may be scheduled tasks causing this.

  You can view **DELETE Statements per Second**, **INSERT Statements per Second**, **INSERT_SELECT Statements per Second**, **REPLACE Statements per Second**, **REPLACE_SELECTION Statements per Second**, **SELECT Statements per Second**, and **UPDATE Statements per Second** on the console to determine whether the load has been changing regularly. For details, see **Viewing Monitoring Metrics**.

- Solution

  Adjust the time when scheduled tasks are run. You are advised to run scheduled tasks during off-peak hours.

11 Security and Encryption

11.1 Database Account Security

Setting the Account Password Complexity

For information about the database password strength requirements on the console, see the database configuration table in **Buying a DB Instance**.

RDS for MariaDB has a password security policy for user-created database accounts. Passwords must:

- Contain at least eight characters.
- Contain at least one uppercase letter, one lowercase letter, one digit, and one special character.

When you are creating a DB instance, the password strength is checked. You can modify the password strength as user **root**. For security reasons, the new password strength must be at least as strong as the initial setting.

Account Description

To provide O&M services, the system automatically creates system accounts when you create RDS for MariaDB instances. These system accounts are unavailable to you.

NOTICE

Attempting to delete, rename, and change passwords or permissions for these accounts will result in an error. Exercise caution when performing these operations.

- rdsAdmin: a management account, used to query and modify instance information, rectify faults, migrate data, and restore data.
- rdsRepl: the replication account, which is used to synchronize data from primary DB instances to standby DB instances or read replicas.

- rdsBackup: the backup account, which is used for backend backup.
- rdsMetric: the metric monitoring account, which is used by watchdog to collect database status data.

Setting Password Complexity

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the DB instance to navigate to the **Overview** page.

Passwords must:

- Contain at least eight characters.
- Contain at least one uppercase letter, one lowercase letter, one digit, and one special character.
- Must be different from the username.

Step 5 In the navigation pane on the left, choose **Parameters**. On the displayed page, modify the required parameters.

RDS for MariaDB allows you to modify the following parameters:

- **simple_password_check_minimal_length**: Set this parameter to **8**.
- **simple_password_check_letters_same_case**: Set this parameter to **1**.
- **simple_password_check_digits**: Set this parameter to **1**.
- **simple_password_check_other_characters**: Set this parameter to **1**.

NOTICE

Check the value in the **Effective upon Reboot** column.

- If the value is **Yes** and the DB instance status on the **Instances** page is **Parameter change. Pending reboot**, a reboot is required for the modifications to take effect.
  - If you have modified parameters of a primary DB instance, you need to reboot the primary DB instance for the modifications to take effect. (For primary/standby DB instances, the parameter modifications are also applied to the standby DB instance.)
  - If you have modified parameters of a read replica, you need to reboot the read replica for the modifications to take effect.
- If the value is **No**, the modifications take effect immediately.

Step 6 Perform the following operations as needed:

- To save the modifications, click **Save**.
- To cancel the modifications, click **Cancel**.

- To preview the modifications, click **Preview**.

After the parameters are modified, you can click **Change History** to view parameter modification details.

----End

11.2 Resetting the Administrator Password to Restore Root Access

Scenarios

You can reset the administrator password only through a primary instance.

If you forget the password of the administrator account **root**, you can reset the password.

Precautions

- If you have changed the administrator password of the primary DB instance, the administrator passwords of the standby DB instance and read replicas (if any) will also be changed.
- The time required for the new password to take effect depends on the amount of service data currently being processed by the primary DB instance.
- To protect against brute force hacking attempts and ensure system security, change your password periodically.

Method 1

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and choose **More** > **Reset Password** in the **Operation** column.

Step 5 (**Optional**) If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

Step 6 Enter and confirm the new password.

NOTICE

Keep this password secure. The system cannot retrieve it.

The new password must consist of 8 to 32 characters and contain at least three types of the following characters: uppercase letters, lowercase letters, digits, and special characters (~!@$#%^*-_=+?,()&). Enter a strong password and periodically change it for security reasons.

Step 7 Click **OK**.

----End

Method 2

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the DB instance name.

Step 5 On the **Overview** page, find **Administrator** and click **Reset Password** under it.

Step 6 (**Optional**) If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

Step 7 Enter and confirm the new password.

NOTICE

Keep this password secure. The system cannot retrieve it.

The new password must consist of 8 to 32 characters and contain at least three types of the following characters: uppercase letters, lowercase letters, digits, and special characters (~!@$#%^*-_=+?,()&). Enter a strong password and periodically change it for security reasons.

Step 8 Click **OK**.

----End

11.3 Configuring an SSL Connection

Secure Sockets Layer (SSL) is an encryption-based Internet security protocol for establishing secure links between a server and a client. It provides authenticated Internet connections to ensure the privacy and integrity of online communications. SSL:

- Authenticates users and servers, ensuring that data is sent to the correct clients and servers.

- Encrypts data to prevent data theft.
- Ensures data integrity during transmission.

By default, SSL is disabled for new RDS for MariaDB instances. If your client has no SSL compatibility issues, you can enable SSL by referring to **Enabling SSL**. Enabling SSL will increase the network connection response time and CPU resource consumption. Before enabling it, evaluate any potential impacts on service performance.

You can connect to a DB instance through a client using an SSL or non-SSL connection.

- If SSL is disabled (default), use a non-SSL connection.
- If SSL is enabled, use an SSL connection. SSL encrypts connections to the instance, making in-transit data more secure.

Precautions

Enabling or disabling SSL will cause DB instances to reboot and interrupt connections. Exercise caution when performing this operation.

Enabling SSL

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the DB instance name.

Step 5 On the **Overview** page, find **SSL** and click **Enable**.

Step 6 In the displayed dialog box, click **OK**. Wait for some seconds and check that SSL has been enabled on the **Overview** page.

----End

Disabling SSL

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the DB instance name.

Step 5 On the **Overview** page, find **SSL** and click **Disable**.

Step 6 In the displayed dialog box, click **OK**. Wait for some seconds and check that SSL has been disabled on the **Overview** page.

----End

11.4 Configuring a Password Expiration Policy

Using the same password too long makes it easier for hackers to crack or guess your password. Requiring password changes after a certain amount of time can improve security. This section describes how to configure a password expiration policy.

Precautions

- Once your password expires, you cannot log in to the database.
- After the password expiration policy is configured, you need to periodically check whether your password is about to expire.

Modifying the Database Parameter

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **Parameters**.

Step 6 On the displayed page, change the value of **default_password_lifetime**.

The value of this parameter indicates how many days until a password expires. The default value is **0**, indicating that the created user password will never expire.

Step 7 Click **Save**. In the displayed dialog box, click **Yes**.

----End

Configuring the Password Expiration Policy Through DAS

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the DB instance and click **Log In** in the **Operation** column.

Alternatively, click the target DB instance on the **Instances** page. On the displayed **Basic Information** page, click **Log In** in the upper right corner of the page.

Step 5 Enter the username and password and click **Log In**.

Step 6 Choose **SQL Operations** > **SQL Query**.

Step 7 In the editing area, compile the statement shown below. The unit of **password_life_time** is day. You are advised to set it to **180**.

ALTER USER *username* PASSWORD EXPIRE INTERVAL *password_life_time* DAY;

Step 8 Click **Execute SQL**. Then, view SQL execution status on the **Executed SQL Statements**, **Messages**, and **Result** tab pages.

----End

11.5 Unbinding an EIP

Scenarios

The Elastic IP (EIP) service enables your RDS instances to communicate with the Internet using static public IP addresses and scalable bandwidths. But this increases the risk of network-wide attacks on your instances. Using an EIP leaves you open to DoS or DDoS attacks.

As an internal component, the database can be accessed using an internal IP address. Therefore, you are advised to unbind the EIP from the database.

Prerequisites

An EIP has been bound to your DB instance. For details, see **Binding an EIP**.

Unbinding an EIP

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the DB instance that has an EIP bound to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Connectivity & Security**. In the **Connection Information** area, click **Unbind** next to the **EIP** field. In the displayed dialog box, click **Yes**.

Alternatively, in the **Connection Topology** area, click **Public Connection** and then **Unbind** in the connection topology. In the displayed dialog box, click **Yes**.

Step 6 (**Optional**) If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account and cloud product. For details about how to enable operation protection, see ***Identity and Access Management User Guide***.

Step 7 On the **Connectivity & Security** page, view the results.

You can also view the progress and result of unbinding an EIP from a DB instance on the **Task Center** page.

To bind an EIP to the DB instance again, see **Binding an EIP**.

----End

11.6 Using DBSS (Recommended)

Database Security Service (DBSS) is an intelligent database security service. Based on the machine learning mechanism and big data analytics technologies, it can audit your databases, detect SQL injection attacks, and identify high-risk operations.

You are advised to use DBSS to provide extended data security capabilities. For details, see **Database Security Service**.

Advantages

- DBSS can help you meet security compliance requirements.
  - DBSS can help you comply with DJCP (graded protection) standards for database audit.
  - DBSS can help you comply with security laws and regulations, and provide compliance reports that meet data security standards (such as Sarbanes-Oxley).
- DBSS can back up and restore database audit logs and meet the audit data retention requirements.
- DBSS can monitor risks, sessions, session distribution, and SQL distribution in real time.
- DBSS can report alarms for risky behavior and attacks and respond to database attacks in real time.
- DBSS can locate internal violations and improper operations and keep data assets secure.

Deployed in bypass pattern, database audit can perform flexible audits on the database without affecting user services.

- Database audit monitors database logins, operation types (data definition, operation, and control), and operation objects based on risky operations to effectively audit the database.
- Database audit analyzes risks and sessions, and detects SQL injection attempts so you can stay apprised of your database status.
- Database audit provides a report template library to generate daily, weekly, or monthly audit reports according to your configurations. It sends real-time alarm notifications to help you obtain audit reports in a timely manner.

12 Parameters

12.1 Modifying Parameters of an RDS for MariaDB Instance

You can modify parameters in a custom parameter template to optimize database performance. This section describes how to modify parameters of an RDS for MariaDB instance.

Precautions

- You can only change parameter values in custom parameter templates. You cannot change the parameter values in default parameter templates.
- Pay attention to the following points when configuring parameters in a parameter template:
  - When you change a parameter value in a parameter template that has been applied to a DB instance and save the change, the change takes effect only to the DB instance and does not affect other DB instances.
  - When you modify dynamic parameters on the **Parameters** page of a DB instance and save the modifications, the modifications take effect immediately regardless of the **Effective upon Reboot** setting. However, when you modify static parameters on the **Parameters** page of a DB instance and save the modifications, the modifications do not take effect until you manually reboot the DB instance.
  - Modifying parameter template parameters: When you modify parameters in a custom parameter template on the **Parameter Templates** page and save the modifications, the modifications do not take effect until you have applied the template to your DB instances. When you modify static parameters in a custom parameter template on the **Parameter Templates** page and save the modifications, the modifications do not take effect until you have applied the template to your DB instances and manually rebooted those DB instances. For details, see **Applying a Parameter Template**.
  - Improper parameter settings may have unintended consequences, including reduced performance and system instability. Exercise caution when modifying database parameters and you need to back up data

before modifying parameters in a parameter template. Before applying parameter template changes to a production DB instance, you should try out these changes on a test DB instance.

- Global parameters must be modified on the console. Session-level parameters can be modified using SQL statements. When you modify a parameter, the time when the modification takes effect depends on the type of the parameter.

  The service console displays the statuses of DB instances that the parameter template applies to. For example, if the DB instance has not yet used the latest modifications made to its parameter template, its status is **Parameter change. Pending reboot**. Manually reboot the DB instance for the latest modifications to take effect for that DB instance.

  NOTE

  RDS has default parameter templates whose parameter values cannot be changed. You can view these parameter values by clicking the default parameter templates. If a custom parameter template is set incorrectly, the database startup may fail. If this happens, you can re-configure the custom parameter template based on the settings of the default parameter template.

Modifying a Custom Parameter Template and Applying It to DB Instances

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 Choose **Parameter Templates** in the navigation pane on the left. On the **Custom Templates** page, click the target parameter template.

Step 6 On the **Parameters** page, modify parameters as required.

Figure 12-1 Modifying parameters in a parameter template

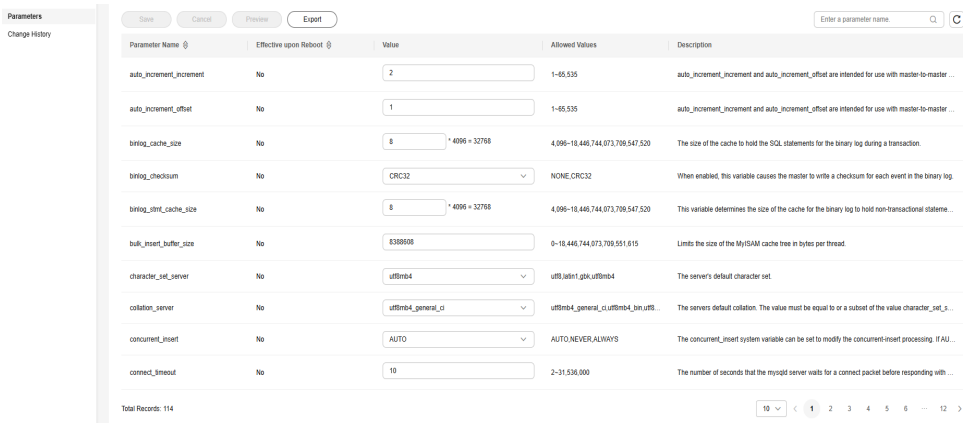

- To save the modifications, click **Save**.

- To cancel the modifications, click **Cancel**.
- To preview the modifications, click **Preview**.

Step 7 After the parameter values are modified, you can click **Change History** to view the details.

Step 8 The modifications do not take effect until you apply the parameter template to your DB instances. For details, see **Applying a Parameter Template**.

Step 9 View the status of the DB instance to which the parameter template was applied.

If the DB instance status is **Parameter change. Pending reboot**, a reboot is required for the modifications to take effect.

- The DB instance reboot caused by instance class changes will not make parameter modifications take effect.
- If you have modified parameters of a primary DB instance, you need to reboot the primary DB instance for the modifications to take effect. (For primary/standby DB instances, the parameter modifications are also applied to the standby DB instance.)
- If you have modified parameters of a read replica, you need to reboot the read replica for the modifications to take effect.

----End

Modifying Parameters of a DB Instance

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **Parameters**. On the displayed page, modify parameters as required.

NOTICE

Check the value in the **Effective upon Reboot** column.

- If the value is **Yes** and the DB instance status on the **Instances** page is **Parameter change. Pending reboot**, a reboot is required for the modifications to take effect.
  - If you have modified parameters of a primary DB instance, you need to reboot the primary DB instance for the modifications to take effect. (For primary/standby DB instances, the parameter modifications are also applied to the standby DB instance.)
  - If you have modified parameters of a read replica, you need to reboot the read replica for the modifications to take effect.
- If the value is **No**, the modifications take effect immediately.

- To save the modifications, click **Save**.
- To cancel the modifications, click **Cancel**.
- To preview the modifications, click **Preview**.

After parameters are modified, you can click **Change History** to view parameter modification details.

----End

12.2 Managing Parameter Templates

12.2.1 Creating a Parameter Template

You can use database parameter templates to manage the DB engine configuration. A database parameter template acts as a container for engine configuration values that can be applied to one or more DB instances. This section describes how to create a parameter template.

Scenarios

If you create a DB instance without specifying a custom DB parameter template, a default parameter template is used. This default template contains DB engine defaults and system defaults that are configured based on the engine, compute class, and allocated storage of the instance. Default parameter templates cannot be modified, but you can create your own parameter template to change parameter settings.

Precautions

- Not all of the DB engine parameters in a custom parameter template can be changed.
- If you want to use a custom parameter template, you simply create a parameter template and select it when you create a DB instance or apply it to an existing DB instance following the instructions provided in **Applying a Parameter Template**.
- When you have already created a parameter template and want to include most of the custom parameters and values from that template in a new parameter template, you can replicate that parameter template following the instructions provided in **Replicating a Parameter Template**.
- RDS does not share parameter template quotas with DDS. You can create a maximum of 100 parameter templates for RDS DB instances. All RDS DB engines share the parameter template quota.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, click **Create Parameter Template**.

Figure 12-2 Creating a parameter template

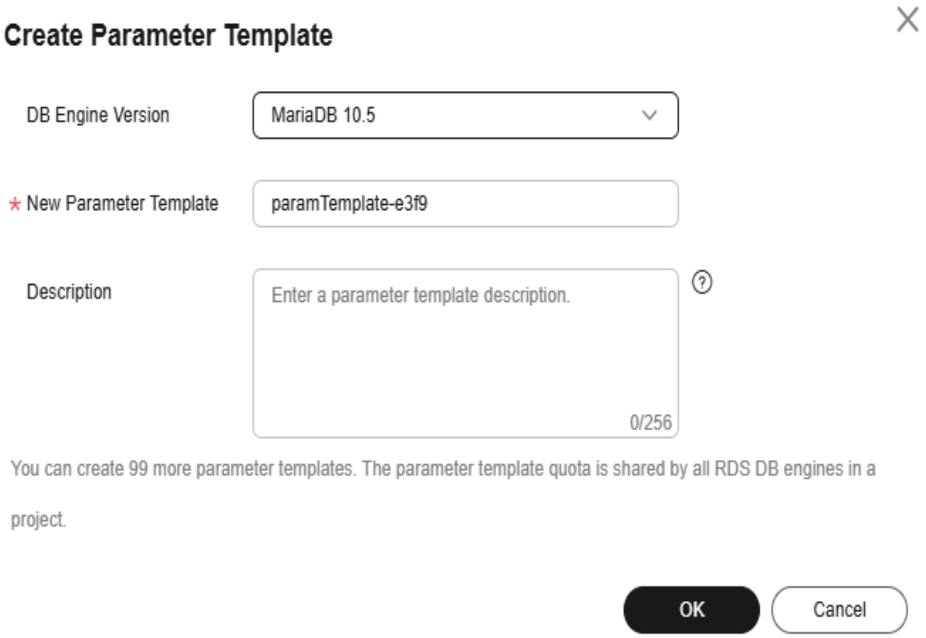

Step 6 In the displayed dialog box, configure required information.

- Select a DB engine for the parameter template.
- The template name must consist of 1 to 64 characters. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), underscores (_), and periods (.).
- The description consists of a maximum of 256 characters and cannot contain carriage return characters or the following special characters: >!<"&'=

Step 7 Click **OK** to create a parameter template.

----End

12.2.2 Applying a Parameter Template

Scenarios

You can apply parameter templates to DB instances as needed.

- The parameter **innodb_buffer_pool_size** is determined by the memory. DB instances of different specifications have different value ranges. If this parameter value is out of range of the DB instance that the parameter template is applied, the maximum value within the range is used.

- A parameter template can be applied only to DB instances of the same DB engine version.

Procedure

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, perform the following operations based on the type of the parameter template to be applied:

- If you intend to apply a default parameter template to DB instances, click **Default Templates**, locate the target parameter template, and click **Apply** in the **Operation** column.
- If you intend to apply a custom parameter template to DB instances, click **Custom Templates**, locate the target parameter template, and choose **More** > **Apply** in the **Operation** column.

A parameter template can be applied to one or more DB instances.

Step 6 In the displayed dialog box, select one or more DB instances to which the parameter template will be applied and click **OK**.

After the parameter template is successfully applied, you can view the application records by referring to **Viewing Application Records of a Parameter Template**.

----End

12.2.3 Replicating a Parameter Template

Scenarios

You can replicate a parameter template you have created. When you have already created a parameter template and want to include most of the custom parameters and values from that template in a new parameter template, you can replicate that parameter template. You can also export the parameter template to generate a new parameter template for future use.

After a parameter template is replicated, it takes about 5 minutes before the new template is displayed.

Default parameter templates cannot be replicated, but you can create parameter templates based on the default ones.

Procedure

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, click **Custom Templates**. Locate the target parameter template and click **Replicate** in the **Operation** column.

NOTE

To ensure that your parameter templates are applicable to all types of DB instances and databases can be started normally, the values of **innodb_flush_log_at_trx_commit** and **sync_binlog** exported from primary DB instances or read replicas are **1** by default.

Figure 12-3 Replicating a parameter template

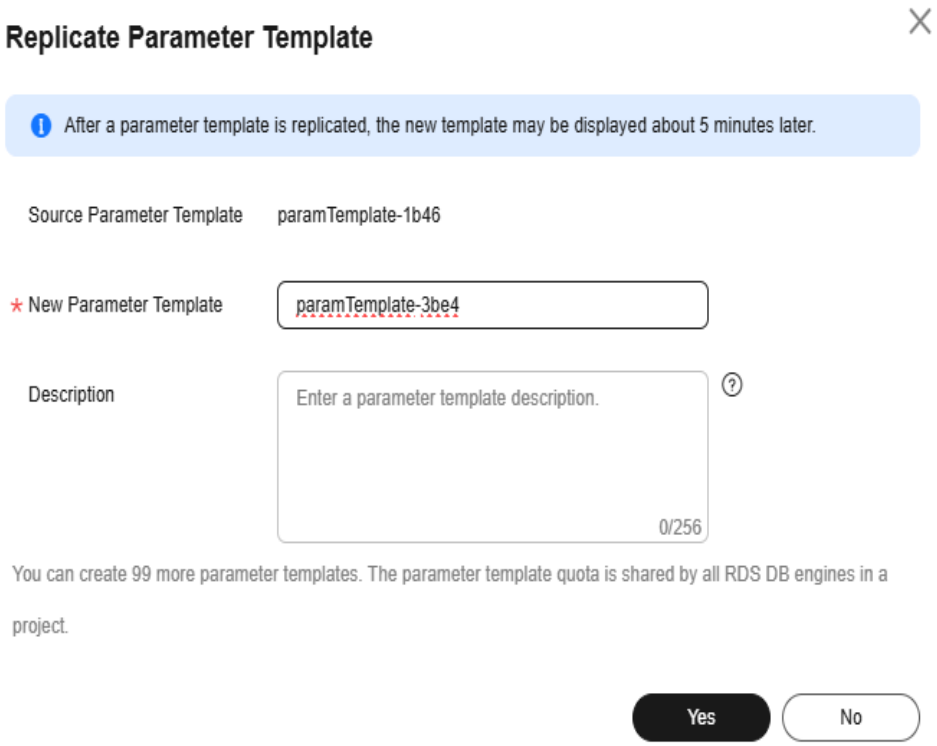

Step 6 In the displayed dialog box, configure required information and click **Yes**.

- The template name must consist of 1 to 64 characters. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), underscores (_), and periods (.).
- The description consists of a maximum of 256 characters and cannot contain carriage return characters or the following special characters: >!<"&'=

After the parameter template is replicated, a new template is generated in the list on the **Parameter Templates** page.

----End

12.2.4 Resetting a Parameter Template

Scenarios

You can reset all parameters in a custom parameter template to their default settings.

Procedure

Step 1 Log in to the management console.

Step 2 Click ☰ in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 3 In the navigation pane on the left, choose **Parameter Templates**.

Step 4 On the **Parameter Templates** page, click **Custom Templates**. Locate the target parameter template and choose **More** > **Reset** in the **Operation** column.

Step 5 Click **Yes**.

Figure 12-4 Confirming the reset

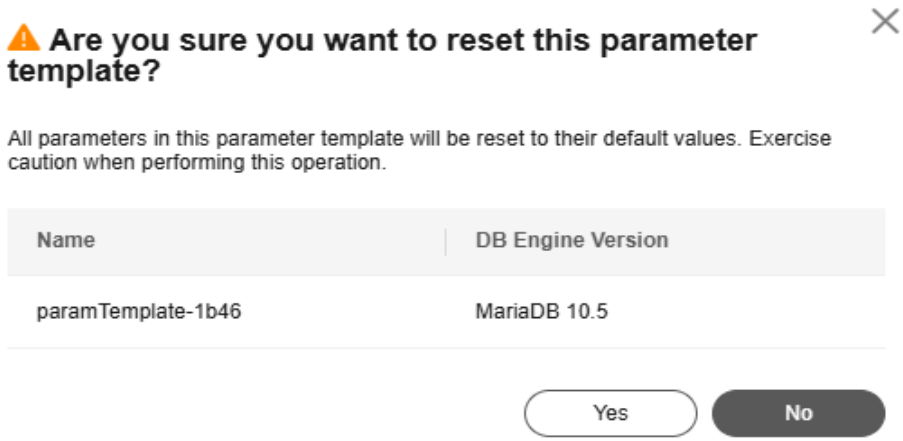

Step 6 The modifications take effect only after you apply the parameter template to DB instances. For details, see **Applying a Parameter Template**.

Step 7 View the status of the DB instance to which the parameter template is applied.

If the DB instance status is **Parameter change. Pending reboot**, a reboot is required for the modifications to take effect.

- The DB instance reboot caused by instance class changes will not make parameter modifications take effect.
- If you have modified parameters of a primary DB instance, you need to reboot the primary DB instance for the modifications to take effect. (For primary/standby DB instances, the parameter modifications are also applied to the standby DB instance.)

- If you have modified parameters of a read replica, you need to reboot the read replica for the modifications to take effect.

----End

12.2.5 Comparing Parameter Templates

Scenarios

You can compare DB instance parameters with a parameter template that uses the same DB engine to understand the differences of parameter settings.

You can also compare default parameter templates that use the same DB engine to understand the differences of parameter settings.

Comparing Instance Parameters with a Parameter Template

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Parameters**. On the displayed page, click **Compare** above the parameter list.

Figure 12-5 Comparing instance parameters with those in a specified parameter template

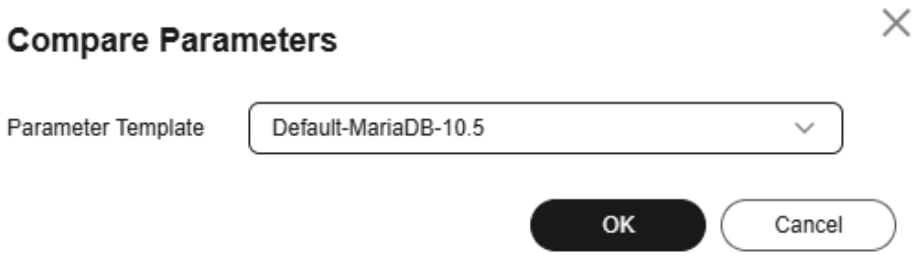

Step 6 In the displayed dialog box, select a parameter template to be compared and click **OK**.

- If their settings are different, the parameter names and values of both parameter templates are displayed.
- If their settings are the same, no data is displayed.

----End

Comparing Parameter Templates

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, click **Custom Templates**. Locate the target parameter template and click **Compare** in the **Operation** column.

Step 6 In the displayed dialog box, select a parameter template that uses the same DB engine as the target template and click **OK**.

Figure 12-6 Selecting a parameter template to be compared

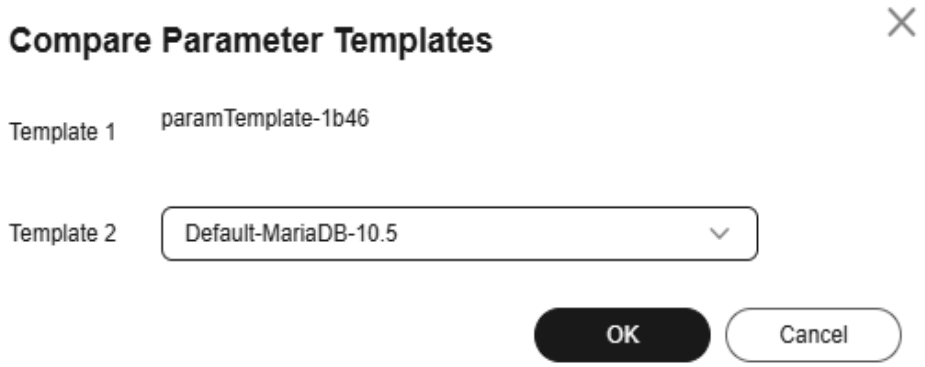

- If their settings are different, the parameter names and values of both parameter templates are displayed.
- If their settings are the same, no data is displayed.

----End

12.2.6 Exporting a Parameter Template

To view and use parameters of a DB instance, you can export the parameter template. This section describes how to export a parameter template.

Scenarios

- You can export a parameter template of a DB instance for future use. You can also apply the exported parameter template to DB instances by referring to **Applying a Parameter Template**.
- You can export the parameter template information (parameter names, values, and descriptions) of a DB instance to a CSV file for analysis.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Parameters**. On the displayed page, click **Export** above the parameter list.

Figure 12-7 Exporting parameters

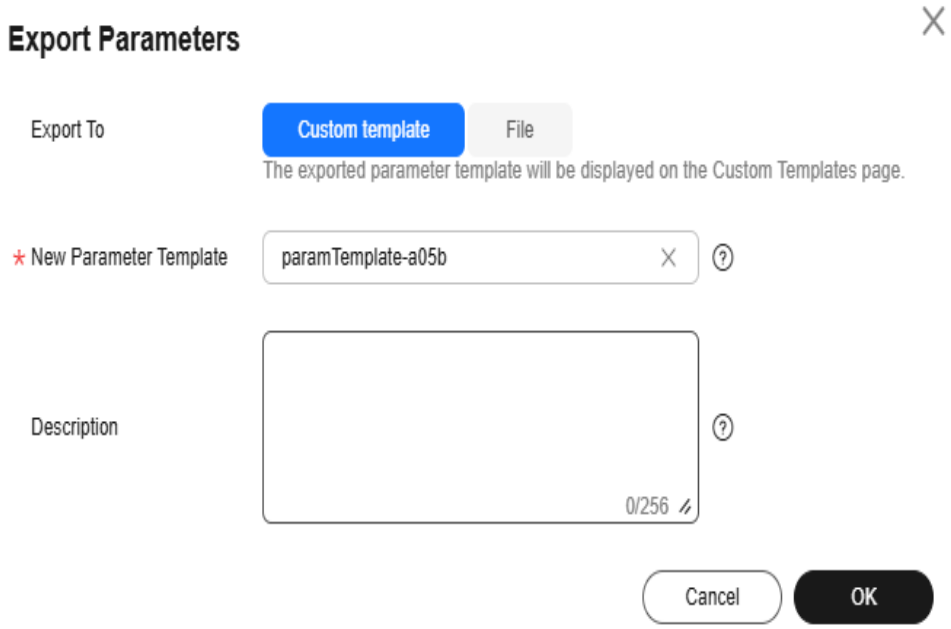

- Exporting to a custom template

  In the displayed dialog box, configure required information and click **OK**.

  - The template name must consist of 1 to 64 characters. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), underscores (_), and periods (.).
  - The description consists of a maximum of 256 characters and cannot contain carriage return characters or the following special characters: >! <"&'=

  After the parameter template is exported, a new template is generated in the list on the **Parameter Templates** page.

- Exporting to a file

  The parameter template information (parameter names, values, and descriptions) of the DB instance is exported to a CSV file. In the displayed dialog box, enter the file name and click **OK**.

  The file name must start with a letter and consist of 4 to 81 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.

----End

12.2.7 Importing a Parameter Template

RDS allows you to import new parameter templates for future use. To apply an imported parameter template to new DB instances, see **Applying a Parameter Template**.

Constraints

- Only parameter templates that were exported from the **Parameter Templates** page on the RDS console can be imported.
- If any modification to an exported parameter template causes a change in the file format, the template may not be able to be imported.
- The parameter template to be imported cannot contain parameters related to specifications.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Parameter Templates** page, click **Import Parameter Template**.

Step 5 In the displayed dialog box, click **Select File**, import the target parameter list (containing parameter names, values, and description), and click **OK**.

Only one file (CSV format) can be imported at a time. The file size cannot exceed 50 KB.

Figure 12-8 Importing a parameter template

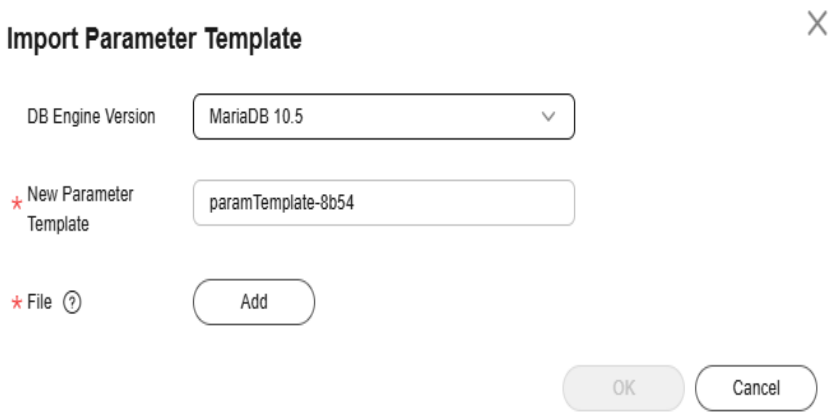

----End

12.2.8 Viewing Parameter Change History

Scenarios

You can view the change history of DB instance parameters or custom parameter templates.

NOTE

The change history for an exported or custom parameter template is initially blank.

Viewing Change History of a DB Instance

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Parameters**. On the displayed page, click **Change History**. The parameter change history of the last seven days is displayed.

You can view the parameter name, original parameter value, new parameter value, modification status, modification time, application status, and application time.

You can apply the parameter template to DB instances as required by referring to Applying a Parameter Template.

----End

Viewing Change History of a Parameter Template

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 Choose **Parameter Templates** in the navigation pane on the left. On the **Custom Templates** page, click the target parameter template.

Step 6 On the displayed page, choose **Change History** in the navigation pane on the left. The parameter change history of the last seven days is displayed.

You can view the parameter name, original parameter value, new parameter value, modification status, and modification time.

You can apply the parameter template to DB instances as required by referring to **Applying a Parameter Template**.

----End

Viewing Parameter Changes

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, click the **Parameter Changes** tab.

Figure 12-9 Viewing parameter changes

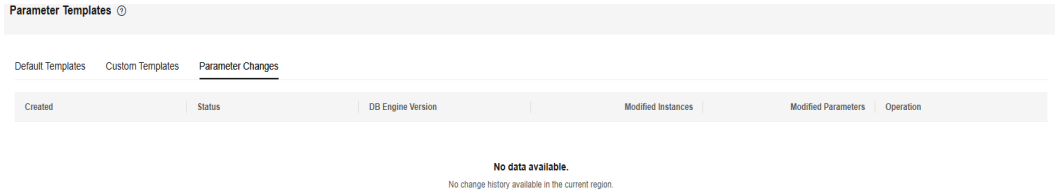

Step 6 Click **View Details** in the **Operation** column.

You can view detailed information about the modified parameters.

----End

12.2.9 Viewing Application Records of a Parameter Template

Scenarios

You can view the application records of a parameter template.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Default Templates** or **Custom Templates** page, locate the target parameter template and choose **More** > **View Application Record** in the **Operation** column.

You can view the name or ID of the DB instance that the parameter template is applied, as well as the application status, application time, and failure cause (if failed).

Figure 12-10 Viewing application records of a parameter template

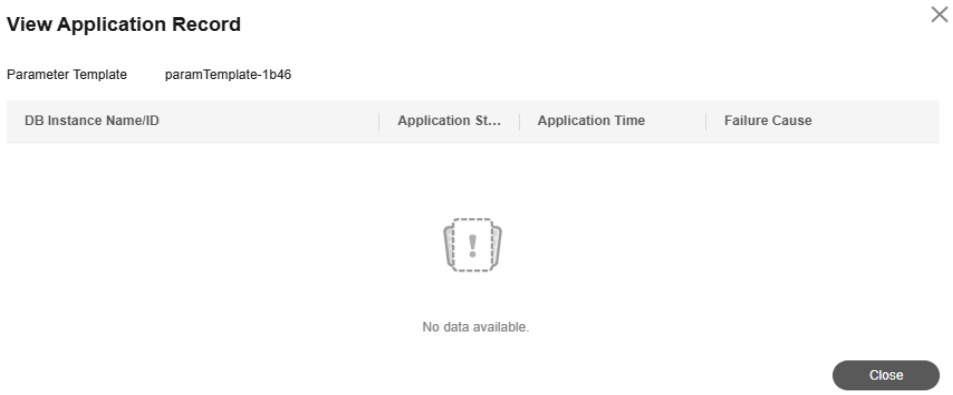

----End

12.2.10 Modifying a Parameter Template Description

Scenarios

You can modify the description of a parameter template you have created.

NOTE

You cannot modify the description of a default parameter template.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, click **Custom Templates**. Locate the target parameter template and click in the **Description** column.

Step 6 Enter a new description and click **OK** to submit the modification or click **Cancel** to cancel the modification.

- The description consists of a maximum of 256 characters and cannot contain carriage return characters or the following special characters: >!<"&'=

- After the modification is successful, you can view the new description in the **Description** column of the parameter template list.

----End

12.2.11 Deleting a Parameter Template

Scenarios

You can delete a custom parameter template that is no longer needed.

NOTICE

- Deleted parameter templates cannot be recovered. Exercise caution when performing this operation.
- Default parameter templates cannot be deleted.

Procedure

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, click **Custom Templates**. Locate the parameter template to be deleted and choose **More** > **Delete** in the **Operation** column.

Step 6 In the displayed dialog box, click **Yes**.

----End

13 Log Management

13.1 Viewing and Downloading Error Logs

RDS log management allows you to view database-level logs, including error logs and slow SQL query logs.

Error logs help you analyze problems with databases. You can download error logs for further analysis.

You can view error logs generated within the last month.

Viewing Log Details

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Logs**. On the **Error Logs** page, click **Log Details** to view details about error logs.

- You can select a log level in the upper right corner to view logs of the selected level.

  NOTE

  For RDS for MariaDB instances, the following levels of logs are displayed:
  - All log levels
  - ERROR
  - WARNING
  - NOTE

- Currently, a maximum of 2,000 error log records can be displayed.
- You can click in the upper right corner to view logs generated in different time segments.

- Only error logs generated within the last one month can be viewed.
- If the description of a log is truncated, locate the log and move your pointer over the description in the **Description** column to view details.

----End

Downloading an Error Log

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **Logs**. On the **Error Logs** page, click the **Downloads** tab.

Step 6 Locate a log file whose status is **Preparation completed** and click **Download** in the **Operation** column.

- The system automatically loads the downloading preparation tasks. The time required depends on the log file size and the network environment.
  - When the log is being prepared for download, the log status is **Preparing**.
  - When the log is ready for download, the log status is **Preparation completed**.
  - If the preparation for download fails, the log status is **Abnormal**.

  Logs in the **Preparing** or **Abnormal** status cannot be downloaded.
- If the size of a log to be downloaded is greater than 40 MB, you need to use OBS Browser+ to download it. For details, see **Method 1: Using OBS Browser +**.
- The download link is valid for 5 minutes. After the download link expires, a message is displayed indicating that the download link has expired. If you need to download the log, click **OK**.
- The downloaded logs contain only the logs of the primary node.

----End

13.2 Viewing and Downloading Slow Query Logs

Scenarios

Slow query logs record statements that exceed **long_query_time** (1 second by default). You can view log details to identify statements that are executing slowly and optimize the statements. You can also download slow query logs for service analysis.

Slow query logs generated within the last month can be viewed.

You can search slow query logs by SQL statement type and time range.

Parameter Description

Table 13-1 Parameters related to MariaDB slow queries

Parameter	Description
long_query_time	Specifies how many microseconds a SQL query has to take to be defined as a slow query log. The default value is 1s. When the execution time of an SQL statement exceeds the value of this parameter, the SQL statement is recorded in slow query logs. The recommended value is **1s**. Note: The lock wait time is not calculated into the query time.
log_queries_not_using_indexes	Specifies whether to record the slow queries without indexes. The default value is **OFF**.

Viewing Log Details

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane on the left, choose **Logs**. On the **Slow Query Logs** page, click **Log Details** to view details about slow query logs.

NOTE

- You can view the slow query log records of a specified execution statement type or a specific time period.
- Only SELECT statements return the number of result rows. The number of result rows for the INSERT, UPDATE, DELETE, and CREATE statements is 0 by default.
- You can view slow query logs of a specified database name (which cannot contain any special characters). The database name supports only exact search.
- Slow query logs only record executed statements whose execution duration exceeds the threshold.
- The **long_query_time** parameter determines when a slow query log is recorded. However, changes to this parameter do not affect already recorded logs. If **long_query_time** is changed from 1s to 0.1s, RDS starts recording statements that meet the new threshold and still displays the previously recorded logs that do not meet the new threshold. For example, a 1.5s SQL statement that was recorded when the threshold was 1s will not be deleted now that the new threshold is 2s.
- A maximum of 2,000 slow log records can be displayed. To view more slow log records, submit a service ticket.
- If the length of a single line of an SQL statement exceeds 10 KB or the total number of lines exceeds 200, the SQL statement will be truncated. When you view slow query log details, the SQL statement may be incomplete after special processing and is for reference only.

----End

Downloading a Slow Query Log

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **Logs**. On the **Slow Query Logs** page, click the **Downloads** tab.

Step 6 Locate a log file whose status is **Preparation completed** and click **Download** in the **Operation** column.

- The system automatically loads the downloading preparation tasks. The loading duration is determined by the log file size and network environment.
  - When the log is being prepared for download, the log status is **Preparing**.
  - When the log is ready for download, the log status is **Preparation completed**.
  - If the preparation for download fails, the log status is **Abnormal**.

  Logs in the **Preparing** or **Abnormal** status cannot be downloaded.
- Only logs no more than 40 MB can be downloaded directly from this page. The time range is calculated from the time you download the logs back to the time when the accumulated file size reaches 40 MB.

- It is impossible to generate a log file much larger than 40 MB, like 100 MB or 200 MB. If a log file that is a little larger than 40 MB is required, use OBS Browser+ to download it by referring to **Method 1: Using OBS Browser+**.
- The download link is valid for 5 minutes. After the download link expires, a message is displayed indicating that the download link has expired. If you need to download the log, click **OK**.
- The downloaded logs contain only the logs of the primary node.

----End

13.3 Enabling or Disabling SQL Audit

After you enable SQL audit, all SQL operations will be recorded in log files. You can **download** audit logs to view log details.

By default, SQL audit is disabled because enabling this function may affect database performance. This section describes how to enable, modify, or disable SQL audit.

Notes

- Both DB instances and read replicas support SQL audit logging.
- Audit logs use the Coordinated Universal Time (UTC) format and are not affected by the time zone configuration.
- After SQL audit is enabled, RDS records SQL operations in audit logs. The generated audit log files are temporarily stored in the instance and then uploaded to OBS and stored in the backup space. If there is not enough free backup space available for generated audit logs, the additional space required is billed.
- Audit logs are cleared every hour. After you change the retention period of audit logs, expired audit logs will be deleted 1 hour later.
- After SQL audit is enabled, a large number of audit logs may be generated during peak hours. As a result, there are many audit log files temporarily stored in the instance, and the storage may be full. You are advised to enable **storage autoscaling**.

Precautions

- Enabling SQL audit deteriorates instance performance by about 5%.
- After SQL audit is disabled, all audit logs will be deleted immediately and cannot be recovered. Exercise caution when performing this operation.

Enabling SQL Audit

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **Logs**. On the **SQL Audit Logs** tab page, click **Set SQL Audit**.

Figure 13-1 Setting SQL audit

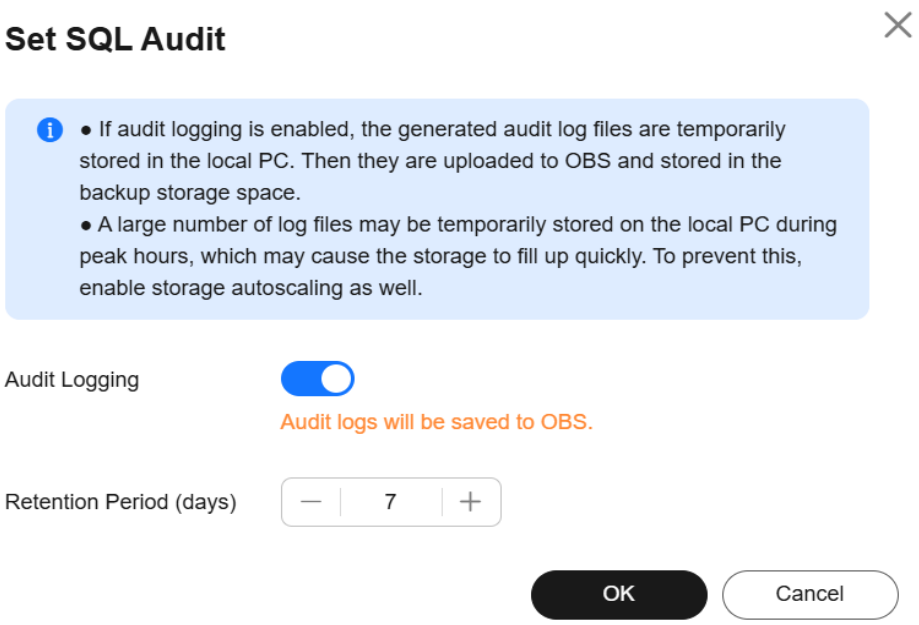

Step 6 In the displayed dialog box, toggle on the **Audit Logging** switch and set the log retention period.

- Set  to  to enable audit logging.
- Audit logs are retained for 7 days by default but can be retained from 1 to 732 days if needed.

Step 7 Click **OK**.

----End

Disabling SQL Audit

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **Logs**. On the **SQL Audit Logs** tab page, click **Set SQL Audit**.

Figure 13-2 Setting SQL audit

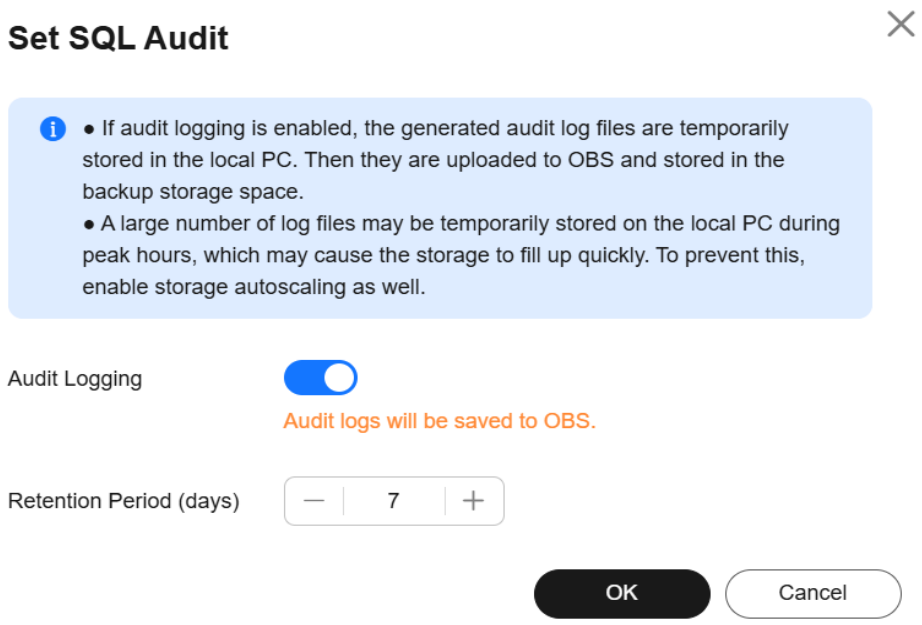

Step 6 In the displayed dialog box, toggle off the **Audit Logging** switch and select the check box "I acknowledge that after audit log is disabled, all audit logs are deleted."

NOTICE

Deleted audit logs cannot be recovered. Exercise caution when performing this operation.

Step 7 Click **OK**.

----End

13.4 Downloading SQL Audit Logs

If you enable SQL audit, all SQL operations will be logged, and you can download audit logs to view details. The minimum time unit of audit logs is second. By default, SQL audit is disabled. Enabling this function may affect database performance.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **Logs**.

Step 6 On the **SQL Audit Logs** page, select a time range in the upper right corner, select SQL audit logs to be downloaded in the list, and click **Download** in the upper left corner of the list to download SQL audit logs in batches.

Alternatively, select an audit log and click **Download** in the **Operation** column to download an individual SQL audit log.

Step 7 The following figure shows the SQL audit log content. For field descriptions, see **Table 13-2**.

Figure 13-3 RDS for MariaDB audit logs

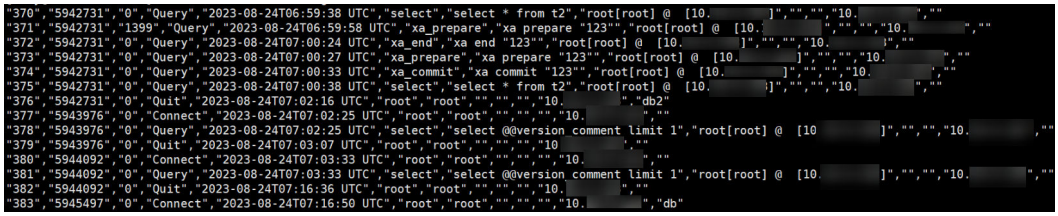

Table 13-2 Audit log field description

Parameter	Description
record_id	ID of a single record, which is the unique global ID of each SQL statement recorded in the audit log.
connection_id	ID of the session executed for the record, which is the same as the ID in the **show processlist** command output.
connection_status	Session status, which is usually the returned error code of a statement. If a statement is successfully executed, the value **0** is returned.
name	Recorded type name. Generally, DML and DDL operations are QUERY, connection and disconnection operations are CONNECT and QUIT, respectively.
timestamp	UTC time for the record.
command_class	SQL command type. The value is the parsed SQL type, for example, select or update. (This field does not exist if the connection is disconnected.)
sqltext	Executed SQL statement content. (This field does not exist if the connection is disconnected.)
user	Login account.

Parameter	Description
host	Login host. The value is **localhost** for local login and is empty for remote login.
external_user	External username.
ip	IP address of the remotely-connected client. For local connection, the field is empty.
default_db	Default database on which SQL statements are executed. **NOTE** Only when you have specified a database name using **-D** in the command for connecting to your DB instance, can the database name be queried in audit logs. If no database name has been specified, this parameter is left blank in audit logs. In the following example, the specified database name is **db**. **mysql -h** 10.10.0.233 **-P** 3306 **-u** root **-p -D** db

----End

14 Metrics and Alarms

14.1 Configuring Displayed Metrics

The Agent monitors RDS DB instances and collects monitoring metrics only.

This section describes the metrics supported by RDS. It covers description, namespaces, monitoring metrics, and dimensions.

Description

This section describes the RDS for MariaDB metrics that can be monitored by Cloud Eye as well as their namespaces and dimensions. You can use APIs provided by Cloud Eye to query the monitoring metrics and alarms generated for RDS for MariaDB.

The monitoring interval is 1 minute.

Namespace

SYS.RDS

RDS for MariaDB Metrics

Table 14-1 Monitored metrics

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds001_cpu_util	CPU Usage	CPU usage of the monitored object	0–100	%	N/A	mariadb_cluster_id	1 minute

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds002_mem_util	Memory Usage	Memory usage of the monitored object	0–100	%	N/A	mariadb_cluster_id	1 minute
rds003_iops	IOPS	Average number of I/O requests processed by the system in a specified period	≥ 0	counts/s	N/A	mariadb_cluster_id	1 minute
rds004_bytes_in	Network Input Throughput	Incoming traffic in bytes per second	≥ 0	bytes/s	N/A	mariadb_cluster_id	1 minute
rds005_bytes_out	Network Output Throughput	Outgoing traffic in bytes per second	≥ 0	bytes/s	N/A	mariadb_cluster_id	1 minute
rds006_conn_count	Total Connections	Total number of connections that attempt to connect to the MariaDB server	≥ 0	counts	N/A	mariadb_cluster_id	1 minute

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds007_conn_active_count	Current Active Connections	Number of current active connections	≥ 0	counts	N/A	mariadb_cluster_id	1 minute
rds008_qps	QPS	Query times of SQL statements (including stored procedures) per second	≥ 0	queries/s	N/A	mariadb_cluster_id	1 minute
rds009_tps	TPS	Execution times of submitted and rollback transactions per second	≥ 0	transactions/s	N/A	mariadb_cluster_id	1 minute
rds010_innodb_buf_usage	Buffer Pool Usage	Ratio of idle pages to the total number of buffer pool pages in the InnoDB buffer	0–1	ratio	N/A	mariadb_cluster_id	1 minute
rds011_innodb_buf_hit	Buffer Pool Hit Ratio	Ratio of read hits to read requests	0–1	ratio	N/A	mariadb_cluster_id	1 minute
rds012_innodb_buf_dirty	Buffer Pool Dirty Block Ratio	Ratio of dirty data to used pages in the InnoDB buffer	0–1	ratio	N/A	mariadb_cluster_id	1 minute

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds013_innodb_reads	InnoDB Read Throughput	Number of read bytes per second in the InnoDB buffer	≥ 0	bytes/s	N/A	mariadb_cluster_id	1 minute
rds014_innodb_writes	InnoDB Write Throughput	Number of write bytes per second in the InnoDB buffer	≥ 0	bytes/s	N/A	mariadb_cluster_id	1 minute
rds015_innodb_read_count	InnoDB File Read Frequency	Number of times that InnoDB reads data from files per second	≥ 0	bytes/s	N/A	mariadb_cluster_id	1 minute
rds016_innodb_write_count	InnoDB File Write Frequency	Number of times that InnoDB writes data to files per second	≥ 0	counts/s	N/A	mariadb_cluster_id	1 minute
rds017_innodb_log_write_req_count	InnoDB Log Write Requests per Second	Number of log write requests per second	≥ 0	counts/s	N/A	mariadb_cluster_id	1 minute

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds018_innodb_log_write_count	InnoDB Log Physical Write Frequency	Number of physical write times to log files per second	≥ 0	counts/s	N/A	mariadb_cluster_id	1 minute
rds019_innodb_log_fsync_count	InnoDB Log fsync() Write Frequency	Number of completed fsync() write times to log files per second	≥ 0	counts/s	N/A	mariadb_cluster_id	1 minute
rds020_temp_tbl_rate	Temporary Tables Created per Second	Number of temporary tables created on hard disks per second	≥ 0	counts/s	N/A	mariadb_cluster_id	1 minute
rds021_myisam_buf_usage	Key Buffer Usage	MyISAM key buffer usage	0–1	ratio	N/A	mariadb_cluster_id	1 minute
rds022_myisam_buf_write_hit	Key Buffer Write Hit Ratio	MyISAM key buffer write hit ratio	0–1	ratio	N/A	mariadb_cluster_id	1 minute

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds023_myisam_buf_read_hit	Key Buffer Read Hit Ratio	MyISAM key buffer read hit ratio	0–1	ratio	N/A	mariadb_cluster_id	1 minute
rds024_myisam_disk_write_count	MyISAM Disk Write Frequency	Number of times that indexes are written to disks per second	≥ 0	counts/s	N/A	mariadb_cluster_id	1 minute
rds025_myisam_disk_read_count	MyISAM Disk Read Frequency	Number of times that indexes are read from disks per second	≥ 0	counts/s	N/A	mariadb_cluster_id	1 minute
rds026_myisam_buf_write_count	MyISAM Buffer Pool Write Requests per Second	Number of requests for writing indexes into the buffer pool per second	≥ 0	counts/s	N/A	mariadb_cluster_id	1 minute

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds027_myisam_buf_read_count	MyISAM Buffer Pool Read Requests per Second	Number of requests for reading indexes from the buffer pool per second	≥ 0	counts/s	N/A	mariadb_cluster_id	1 minute
rds028_comdml_del_count	DELETE Statements per Second	Number of DELETE statements executed per second	≥ 0	queries/s	N/A	mariadb_cluster_id	1 minute
rds029_comdml_ins_count	INSERT Statements per Second	Number of INSERT statements executed per second	≥ 0	queries/s	N/A	mariadb_cluster_id	1 minute
rds030_comdml_ins_sel_count	INSERT_SELECT Statements per Second	Number of INSERT_SELECT statements executed per second	≥ 0	queries/s	N/A	mariadb_cluster_id	1 minute

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds031_comdml_rep_count	REPLACE Statements per Second	Number of REPLACE statements executed per second	≥ 0	queries/s	N/A	mariadb_cluster_id	1 minute
rds032_comdml_rep_sel_count	REPLACE_SELECTION Statements per Second	Number of REPLACE_SELECTION statements executed per second	≥ 0	queries/s	N/A	mariadb_cluster_id	1 minute
rds033_comdml_sel_count	SELECT Statements per Second	Number of SELECT statements executed per second	≥ 0	queries/s	N/A	mariadb_cluster_id	1 minute
rds034_comdml_upd_count	UPDATE Statements per Second	Number of UPDATE statements executed per second	≥ 0	queries/s	N/A	mariadb_cluster_id	1 minute
rds035_innodb_del_row_count	Row Delete Frequency	Number of rows deleted from the InnoDB table per second	≥ 0	rows/s	N/A	mariadb_cluster_id	1 minute

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds036_innodb_ins_row_count	Row Insert Frequency	Number of rows inserted into the InnoDB table per second	≥ 0	rows/s	N/A	mariadb_cluster_id	1 minute
rds037_innodb_read_row_count	Row Read Frequency	Number of rows read from the InnoDB table per second	≥ 0	rows/s	N/A	mariadb_cluster_id	1 minute
rds038_innodb_upd_row_count	Row Update Frequency	Number of rows updated into the InnoDB table per second	≥ 0	rows/s	N/A	mariadb_cluster_id	1 minute
rds039_disk_util	Storage Space Usage	Storage space usage of the monitored object	0–100	%	N/A	mariadb_cluster_id	1 minute
rds047_disk_total_size	Total Storage Space	Total storage space of the monitored object	40–4000	GB	1024	mariadb_cluster_id	1 minute
rds048_disk_used_size	Used Storage Space	Used storage space of the monitored object	0–4000	GB	1024	mariadb_cluster_id	1 minute

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds049_disk_read_throughput	Disk Read Throughput	Number of bytes read from the disk per second	≥ 0	bytes/s	N/A	mariadb_cluster_id	1 minute
rds050_disk_write_throughput	Disk Write Throughput	Number of bytes written into the disk per second	≥ 0	bytes/s	N/A	mariadb_cluster_id	1 minute
rds072_conn_usage	Connection Usage	Percent of used MariaDB connections to the total number of connections	0–100	%	N/A	mariadb_cluster_id	1 minute
rds073_replication_delay	Real-Time Replication Delay	Real-time replication delay between standby DB instances or read replicas and primary DB instances, corresponding to seconds_behind_master.	≥ 0	s	N/A	mariadb_cluster_id	1 minute 5 seconds
rds074_slow_queries	Slow Query Logs	Number of slow query logs generated per minute by MariaDB	≥ 0	counts/min	N/A	mariadb_cluster_id	1 minute

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds075_avg_disk_ms_per_read	Disk Read Time	Average time required for each disk read in a specified period	≥ 0	ms	N/A	mariadb_cluster_id	1 minute
rds076_avg_disk_ms_per_write	Disk Write Time	Average time required for each disk write in a specified period	≥ 0	ms	N/A	mariadb_cluster_id	1 minute
rds077_vma	VMA	Virtual memory area size of an RDS process	≥ 0	counts	N/A	mariadb_cluster_id	1 minute
rds078_threads	Threads	Number of threads in a process	≥ 0	counts	N/A	mariadb_cluster_id	1 minute
rds079_vm_hwm	Peak Resident Set Size	Peak physical memory usage of an RDS process	≥ 0	KB	1024	mariadb_cluster_id	1 minute
rds080_vm_peak	Peak Virtual Memory Size	Peak virtual memory usage of an RDS process	≥ 0	KB	1024	mariadb_cluster_id	1 minute
rds082_semi_sync_tx_avg_wait_time	Transaction Wait Time	Average wait time of transactions in semi-synchronous mode	≥ 0	μs	N/A	mariadb_cluster_id	1 minute

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds173_replication_delay_avg	Average Replication Delay	Average replication delay between standby DB instances or read replicas and primary DB instances, corresponding to seconds_behind_master	≥ 0	s	N/A	mariadb_cluster_id	1 minute
rds_buffer_pool_wait_free	Dirty Pages to Be Flushed to Disks	When InnoDB needs to read or create a page and no clean pages are available, InnoDB flushes some dirty pages first and waits for that operation	≥ 0	counts	N/A	mariadb_cluster_id	1 minute
rds_bytes_recv_rate	Received Bytes per Second	Number of bytes received by the database per second	≥ 0	bytes/s	N/A	mariadb_cluster_id	1 minute
rds_bytes_sent_rate	Sent Bytes per Second	Number of bytes sent from the database per second	≥ 0	bytes/s	N/A	mariadb_cluster_id	1 minute

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds_conn_active_usage	Active Connection Usage	Usage of active connections	0–100	%	N/A	mariadb_cluster_id	1 minute
rds_created_tmp_tables_rate	Temporary Tables Created per Second	Number of temporary tables created per second	≥ 0	counts/s	N/A	mariadb_cluster_id	1 minute
rds_innodb_buffer_pool_pages_flushed_rate	Innodb_buffer_pool Page Flushes per Second	Number of innodb_buffer_pool page flushes per second	≥ 0	counts/s	N/A	mariadb_cluster_id	1 minute
rds_innodb_buffer_pool_read_requests_rate	Innodb_buffer_pool Read Requests per Second	Number of innodb_buffer_pool read requests per second	≥ 0	counts/s	N/A	mariadb_cluster_id	1 minute

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds_innodb_buffer_pool_write_requests_rate	Innodb_buffer_pool Write Requests per Second	Number of innodb_buffer_pool write requests per second	≥ 0	counts/s	N/A	mariadb_cluster_id	1 minute
rds_innodb_lock_waits	Row Locks Waits Transactions	Number of InnoDB transactions waiting for row lock	≥ 0	counts	N/A	mariadb_cluster_id	1 minute
rds_innodb_log_waits_count	Log Buffer Status	Number of times that the log buffer was too small and a wait was required for it to be flushed before continuing	≥ 0	counts	N/A	mariadb_cluster_id	1 minute

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds_innodb_log_waits_rate	Flush Times to Disks Due to Insufficient Log Buffer	Times of transaction logs flushed to disks due to insufficient log buffer	≥ 0	counts/s	N/A	mariadb_cluster_id	1 minute
rds_innodb_os_log_written_rate	Redo Log Size Written per Second	Size of redo logs written per second	≥ 0	bytes/s	N/A	mariadb_cluster_id	1 minute
rds_innodb_pages_read_rate	Data Volume Read By InnoDB per Second	Data volume read by InnoDB per second	≥ 0	Pages/s	N/A	mariadb_cluster_id	1 minute

Metric ID	Name	Description	Value Range	Unit	Conversion Rule	Dimension	Monitoring Interval (Raw Data)
rds_innodb_pages_written_rate	Data Volume Written by InnoDB per Second	Data volume written by InnoDB per second	≥ 0	Pages/s	N/A	mariadb_cluster_id	1 minute
rds_innodb_row_lock_current_waits	Current Row Lock Waits	Number of current InnoDB row lock waits	≥ 0	counts	N/A	mariadb_cluster_id	1 minute
rds_innodb_row_lock_time_avg	Row Lock Wait Time	Average wait time of row locks	≥ 0	ms	N/A	mariadb_cluster_id	1 minute
rds_wait_thread_count	Waiting Threads	Number of waiting threads	≥ 0	counts	N/A	mariadb_cluster_id	1 minute

Dimension

Key	Value
mariadb_cluster_id	RDS for MariaDB instance ID You can obtain the value by referring to **Querying DB Instances**.

14.2 Viewing Monitoring Metrics

Scenarios

Cloud Eye monitors the statuses of RDS instances. You can view RDS metrics on the management console. For details, see **Procedure**.

Monitored data takes some time before it can be displayed. The instance status displayed on the Cloud Eye console is about 5 to 10 minutes delayed. When you create a new RDS instance, it takes 5 to 10 minutes before monitoring data is displayed on Cloud Eye.

Prerequisites

- RDS is running properly.

  Monitoring metrics of the RDS instances that are faulty or have been deleted are not displayed on the Cloud Eye console. You can view their monitoring metrics after they are rebooted or restored to normal.

  NOTE

  If an RDS instance has been faulty for 24 hours, Cloud Eye considers it to no longer exist and deletes it from the monitoring object list. You need to manually clear the alarm rules created for the DB instance.

- RDS has been running properly for about 10 minutes.

  For a newly created RDS instance, you need to wait a bit before you can view the metrics.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and click **View Metrics** in the **Operation** column to go to the Cloud Eye console.

Alternatively, click the target DB instance. On the displayed page, click **View Metrics** in the upper right corner of the page.

Step 5 On the displayed page, view the instance metrics.

- On the Cloud Eye console, click **Select Metric** in the upper right corner. In the displayed dialog box, you can select the metrics to be displayed and sort them by dragging them to desired locations.
- You can sort graphs by dragging them based on service requirements.
- You can view the performance metrics in the last 1 hour, 3 hours, 12 hours, 1 day, and 7 days.

----End

14.3 Setting Alarm Rules

Scenarios

You can set alarm rules to customize the monitored objects and notification policies and keep track of the instance status.

RDS alarm rules include alarm rule names, resource types, dimensions, monitored objects, metrics, alarm thresholds, monitoring period, and whether to send notifications.

Procedure

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Choose **Management & Governance** > **Cloud Eye** from the service list.

Step 4 In the navigation pane on the left, choose **Cloud Service Monitoring** > **Relational Database Service**.

Step 5 Locate the instance you want to add an alarm rule for and click **Create Alarm Rule** in the **Operation** column.

Step 6 On the displayed page, set required parameters.

- Basic information

  Table 14-2 Basic information

Parameter	Description
Name	Alarm rule name. The system generates a random name, which you can modify. Example value: alarm-wnat
Description	(Optional) Supplementary information about the alarm rule.

- Alarm parameters

Table 14-3 Alarm parameters

Parameter	Description
Method	You are advised to select **Use existing template**. The existing templates already contain three common alarm metrics: CPU usage, memory usage, and storage space usage. **NOTE** If you select **Associate template**, after the associated template is modified, the policies contained in this alarm rule to be created will be modified accordingly.
Template	Select the template to be used. You can select a default alarm template or create a custom template.
Alarm Policy	If you select **Configure manually** for **Method**, you need to configure alarm policies. An alarm is triggered when the metric configured for this alarm reaches the preset threshold in consecutive periods. For example, Cloud Eye triggers an alarm if the average CPU usage of the monitored object is 80% or more for three consecutive 5-minute periods. **NOTE** A maximum of 50 alarm policies can be added to an alarm rule. If any one of these alarm policies is met, an alarm is triggered.
Alarm Severity	If you select **Configure manually** for **Method**, you need to configure alarm severity. The alarm severity can be **Critical**, **Major**, **Minor**, or **Informational**.

- Alarm notification parameters

Table 14-4 Alarm notification parameters

Parameter	Description
Alarm Notification	Whether to notify users when alarms are triggered. Notifications can be sent by email, text message, or HTTP/HTTPS message.
Notification Recipient	You can select a notification group or topic subscription as required.
Notification Group	Notification group the alarm notification is to be sent to.

Parameter	Description
Notification Object	Object the alarm notification is to be sent to. You can select the account contact or a topic. – The account contact is the mobile phone number and email address of the registered account. – A topic is used to publish messages and subscribe to notifications. If the required topic is unavailable, create one first and add subscriptions to it. For details, see **Creating a Topic** and **Adding Subscriptions**.
Notification Window	Cloud Eye sends notifications only within the notification window specified in the alarm rule. If **Notification Window** is set to **08:00-20:00**, Cloud Eye sends notifications only within 08:00-20:00.
Trigger Condition	Condition for triggering an alarm notification. You can select **Generated alarm** (when an alarm is generated), **Cleared alarm** (when an alarm is cleared), or both.

- Advanced settings

  Table 14-5 Advanced settings

Parameter	Description
Enterprise Project	Enterprise project that the alarm rule belongs to. Only users with the enterprise project permissions can view and manage the alarm rule.
Tag	A tag is a key-value pair. Tags identify cloud resources so that you can easily categorize and search for your resources. – A key can contain a maximum of 128 characters, and a value can contain a maximum of 225 characters. – A maximum of 20 tags can be added.

Step 7 After the configuration is complete, click **Create**.

----End

15 Interconnection with CTS

15.1 Key Operations Supported by CTS

Cloud Trace Service (CTS) records operations related to RDS for further query, audit, and backtrack.

Table 15-1 RDS operations that can be recorded by CTS

Operation	Resource Type	Trace Name
Creating a DB instance or a read replica, or restoring data to a new DB instance	instance	createInstance
Scaling up storage space and changing instance class	instance	instanceAction
Rebooting a DB instance	instance	instanceRestart
Restoring data to the original DB instance	instance	instanceRestore
Renaming a DB instance	instance	instanceRename
Resetting a password	instance	resetPassword
Setting database version parameters	instance	setDBParameters
Resetting database version parameters	instance	resetDBParameters
Enabling, modifying, or disabling a backup policy	instance	setBackupPolicy
Changing a database port	instance	changeInstancePort
Binding or unbinding an EIP	instance	setOrResetPublicIP
Modifying a security group	instance	modifySecurityGroup

Operation	Resource Type	Trace Name
Adding a tag	instance	createTag
Deleting a tag	instance	deleteTag
Editing a tag	instance	modifyTag
Deleting a DB instance	instance	deleteInstance
Performing a primary/standby switchover	instance	instanceFailOver
Changing the replication mode	instance	instanceFailOver-Mode
Changing a failover priority	instance	instanceFailOver-Strategy
Creating a backup	backup	createManualSnap-shot
Replicating a backup	backup	copySnapshot
Downloading a backup (using OBS)	backup	downLoadSnapshot
Downloading a backup (using a browser)	backup	backupsDownLoad
Deleting a backup	backup	deleteManualSnap-shot
Downloading a merged backup	backup	packBackupsDown-Load
Creating a parameter template	parameterGroup	createParameterGrou p
Modifying parameters in a parameter template	parameterGroup	updateParameterGro up
Deleting a parameter template	parameterGroup	deleteParameterGrou p
Replicating a parameter template	parameterGroup	copyParameterGroup
Resetting a parameter template	parameterGroup	resetParameterGroup
Applying a parameter template	parameterGroup	applyParameterGrou p
Saving parameters in a parameter template	parameterGroup	saveParameterGroup
Deleting a frozen DB instance	all	rdsUnsubscribeIn-stance

Operation	Resource Type	Trace Name
Freezing a DB instance	all	rdsfreezeInstance

15.2 Viewing Traces

Scenarios

After CTS is enabled, operations on cloud resources are recorded. You can view the operation records of the last 7 days on the CTS console.

This section describes how to query the operation records of last 7 days on the CTS console.

NOTE

Before using CTS, you need to enable it. For details, see **Enabling CTS**.

Procedure

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 In the upper left corner of the page, click and choose **Management & Governance** > **Cloud Trace Service**.

Step 4 Choose **Trace List** in the navigation pane on the left.

Step 5 Filter conditions to query traces. The details are described as follows:

- **Trace Type**, **Trace Source**, **Resource Type**, and **Search By**: Select a filter from the drop-down list.

  When you select **Resource ID** for **Search By**, you also need to select or enter a resource ID.
- **Operator**: Select a specific operator from the drop-down list.
- **Trace Status**: Available options include **All trace statuses**, **Normal**, **Warning**, and **Incident**. You can only select one of them.
- In the upper right corner of the page, you can specify a time range for querying traces.

Step 6 Click **Query**.

Step 7 Click on the left of the required trace to expand its details.

Step 8 Click **View Trace** in the **Operation** column. On the displayed dialog box, the trace structure details are displayed.

Step 9 Click **Export** on the right. CTS exports traces collected in the past seven days to a CSV file. The CSV file contains all information related to traces on the management console.

For details about key fields in the trace structure, see sections "Trace Structure" and "Trace Examples" in the *Cloud Trace Service User Guide*.

----End

16 Task Center

16.1 Viewing a Task

You can view the progresses and results of scheduled and instant tasks on the **Task Center** page.

Task Details

RDS allows you to view and manage the following instant tasks:

- Creating DB instances
- Creating read replicas
- Scaling up storage space
- Switching primary/standby DB instances
- Rebooting DB instances
- Binding EIPs to DB instances
- Unbinding EIPs from DB instances
- Restoring data to new DB instances

RDS allows you to view and manage the following scheduled tasks:

- Changing MariaDB instance classes
- Rebooting MariaDB instances

Viewing an Instant Task

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 Choose **Task Center** in the navigation pane on the left. Locate the target task and view its details on the displayed **Instant Tasks** page.

- To identify the target task, you can use the task name, order ID, or DB instance name/ID, or simply enter the target task name in the search box in the upper right corner.
- You can view the progress and status of tasks in a specific period. The default period is seven days.

  The task list can only show up to 30 days of past tasks.
- You can view instant tasks in the following statuses:
  - Running
  - Completed
  - Failed
- You can view the task creation and completion time.

----End

Viewing a Scheduled Task

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 Choose **Task Center** in the navigation pane on the left. On the **Scheduled Tasks** page, view the task progress and results.

- To identify the target task, you can use the DB instance name/ID or enter the target DB instance ID in the search box in the upper right corner.
- You can view the scheduled tasks in the following statuses:
  - Running
  - Completed
  - Failed
  - Canceled
  - To be executed
  - To be authorized

----End

16.2 Deleting a Task Record

You can delete task records so that they are no longer displayed in the task list. This operation only deletes the task records, and does not delete the DB instances or terminate the tasks that are being executed.

Precautions

Deleted task records cannot be recovered. Exercise caution when performing this operation.

Deleting an Instant Task Record

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 Choose **Task Center** in the navigation pane on the left. On the displayed **Instant Tasks** page, locate the task record to be deleted and click **Delete** in the **Operation** column. In the displayed dialog box, click **OK**.

You can delete the records of instant tasks in any of the following statuses:

- Completed
- Failed

----End

Deleting a Scheduled Task Record

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 Choose **Task Center** in the navigation pane on the left. On the **Scheduled Tasks** page, locate the task record to be deleted and check whether the task status is **To be executed** or **To be authorized**.

- If yes, go to **Step 5**.
- If no, go to **Step 6**.

Step 5 Click **Cancel** in the **Operation** column. In the displayed dialog box, click **OK** to cancel the task. Then, click **Delete** in the **Operation** column. In the displayed dialog box, click **OK** to delete the task record.

Step 6 Click **Delete** in the **Operation** column. In the displayed dialog box, click **OK** to delete the task record.

You can delete the records of scheduled tasks in any of the following statuses:

- Completed
- Failed
- Canceled
- To be executed
- To be authorized

----End

17 RDS for MariaDB Tags

Tag Management Service (TMS) enables you to use tags on the management console to manage resources. TMS works with other cloud services to manage tags. TMS manages tags globally. Other cloud services manage only their own tags.

Constraints

- Log in to the management console. Click **Service List** and choose **Management & Governance** > **Tag Management Service**. Set predefined tags on the TMS console.
- A tag consists of a key and value. You can add only one value for each key.
- A maximum of 20 tags can be added for each DB instance.

Adding or Editing a Tag

Step 1 **Log in to the management console**.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the instance name to go to the **Overview** page.

Step 5 In the navigation pane, choose **Tags** and click **Edit Tag**.

Step 6 In the displayed dialog box, click **Add**, enter a tag key and value, and click **OK**.

Figure 17-1 Adding a tag

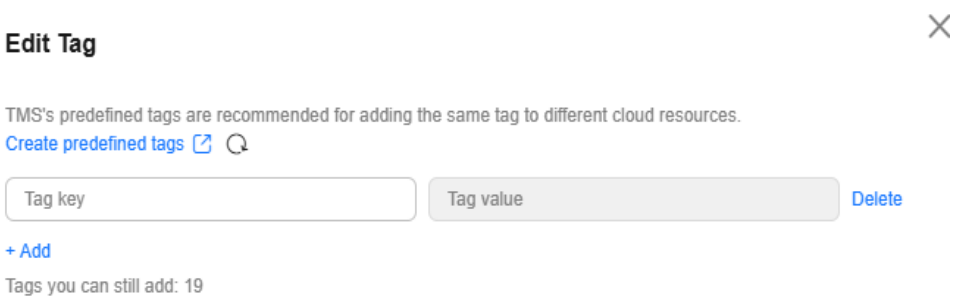

- When you enter a tag key and value, the system automatically displays all tags (including predefined tags and resource tags) associated with your DB instances (except the current instance).
- The tag key must be unique. It must consist of 1 to 128 characters and can include letters, digits, spaces, and the following characters: _ **. : = + - @()**. It cannot start or end with a space, or start with _**sys**_.
- The tag value can be an empty string or consist of up to 255 characters. It can include letters, digits, spaces, and the following characters: _ **. : / = + - @()**. It cannot start or end with a space.

Step 7 View and manage the tag on the **Tags** page.

----End

Deleting a Tag

Step 1 **Log in to the management console**.

Step 2 Click  in the upper left corner and select a region and a project.

Step 3 Click  in the upper left corner of the page and choose **Databases** > **Relational Database Service**.

Step 4 On the **Instances** page, click the target DB instance.

Step 5 In the navigation pane, choose **Tags** and click **Edit Tag**.

Step 6 In the displayed dialog box, locate the tag to be deleted, click **Delete**, and click **OK**.

Verify that the tag is no longer displayed on the **Tags** page.

----End

18 RDS for MariaDB Quotas

What Is a Quota?

A quota is a limit on the quantity or capacity of a certain type of service resources available to you. Examples of RDS quotas include the maximum number of DB instances that you can create. Quotas are put in place to prevent excessive resource usage.

If a quota cannot meet your needs, apply for a higher quota.

Viewing Quotas

Step 1 Click in the upper left corner and select a region.

Step 2 In the upper right corner, choose **Resources** > **My Quotas**.

Figure 18-1 My quotas

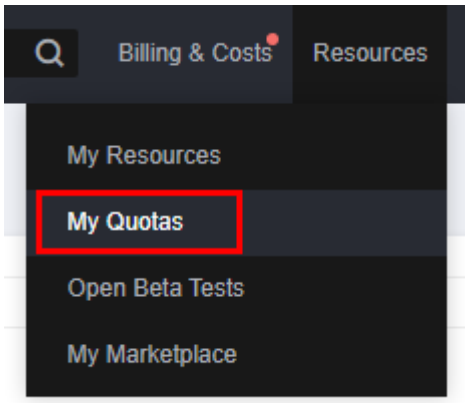

Step 3 On the **Quotas** page, view the used and total quotas of each type of resources.

----End

Increasing Quotas

Step 1 Click in the upper left corner and select a region.

Step 2 In the upper right corner, choose **Resources** > **My Quotas**.

Step 3 In the upper right corner of the page, click **Increase Quota**.

Figure 18-2 Increasing quotas

Step 4 On the **Create Service Ticket** page, configure parameters as required.

In the **Problem Description** area, fill in the content and reason for adjustment.

Step 5 After all required parameters are configured, select the agreement and click **Submit**.

----End

