

Huawei Cloud Flexus RDS

User Guide

Issue	01
Date	2025-09-26



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2025. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Cloud Computing Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI 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.

Contents

1 Working with FlexusRDS for MySQL.....	1
1.1 Using IAM to Grant Access to FlexusRDS.....	1
1.1.1 Creating a User and Granting Permissions for FlexusRDS.....	1
1.1.2 FlexusRDS Custom Policies.....	2
1.2 Buying a FlexusRDS for MySQL Instance.....	3
1.3 Connecting to a FlexusRDS for MySQL Instance.....	6
1.3.1 Using DAS to Connect to a FlexusRDS for MySQL Instance (Recommended).....	6
1.3.2 Using CLI to Connect to a FlexusRDS for MySQL Instance.....	7
1.4 Usage Rules.....	10
1.4.1 Instance Usage Rules.....	10
1.4.2 Database Usage Rules.....	11
1.5 Data Migration.....	15
1.5.1 Migrating Databases to FlexusRDS for MySQL Using DRS.....	15
1.5.2 Migrating Data to FlexusRDS for MySQL Using mysqldump.....	16
1.5.3 Migrating Data to FlexusRDS for MySQL Using the Export and Import Functions of DAS.....	19
1.6 Instance Management.....	23
1.6.1 Upgrading the Minor Version of a FlexusRDS for MySQL Instance.....	23
1.6.2 Upgrading a FlexusRDS for MySQL Instance to an RDS DB Instance.....	25
1.6.3 Changing the Name of a FlexusRDS for MySQL Instance.....	27
1.6.4 Rebooting FlexusRDS for MySQL Instances.....	28
1.6.5 Resetting the Administrator Password of a FlexusRDS for MySQL Instance.....	29
1.6.6 Enabling Storage Autoscaling for a FlexusRDS for MySQL Instance.....	30
1.6.7 Binding an EIP to a FlexusRDS for MySQL Instance or Unbinding an EIP from a FlexusRDS for MySQL Instance.....	32
1.6.8 Changing the VPC and Subnet of a FlexusRDS for MySQL Instance.....	33
1.6.9 Renewing FlexusRDS for MySQL Instances.....	35
1.6.10 Unsubscribing from a FlexusRDS for MySQL Instance.....	36
1.7 Backups and Restorations.....	36
1.7.1 Creating a Manual Backup for a FlexusRDS for MySQL Instance.....	36
1.7.2 Deleting a Manual Backup of a FlexusRDS for MySQL Instance.....	38
1.7.3 Downloading a Full Backup of a FlexusRDS for MySQL Instance.....	38
1.7.4 Checking and Exporting Backup Information of a FlexusRDS for MySQL Instance.....	42
1.7.5 Restoring a FlexusRDS for MySQL Instance.....	43

1.7.5.1 Restoring a FlexusRDS for MySQL Instance from Backups.....	43
1.7.5.2 Restoring a FlexusRDS for MySQL Instance to a Point in Time.....	45
1.8 Parameters.....	48
1.8.1 Suggestions on Parameter Tuning for a FlexusRDS for MySQL Instance.....	48
1.8.2 Modifying Parameters of a FlexusRDS for MySQL Instance.....	50
1.8.3 Exporting the Parameter List of a FlexusRDS for MySQL Instance.....	52
1.9 Monitoring.....	53
1.9.1 Viewing Monitoring Metrics of a FlexusRDS for MySQL Instance.....	53
1.10 Logs.....	55
1.10.1 Viewing Operation Logs of a FlexusRDS for MySQL Instance.....	55
1.10.2 Viewing and Downloading Error Logs of a FlexusRDS for MySQL Instance.....	56
1.10.3 Viewing and Downloading Slow Query Logs of a FlexusRDS for MySQL Instance.....	57
1.11 Interconnection with CTS.....	60
1.11.1 FlexusRDS Operations Supported by CTS.....	60
1.11.2 Viewing Traces.....	61
1.12 FlexusRDS for MySQL Tags.....	63
1.13 FlexusRDS for MySQL Quotas.....	64
2 Working with FlexusRDS for PostgreSQL.....	66
2.1 Permissions Management.....	66
2.1.1 Creating a User and Granting Permissions for FlexusRDS.....	66
2.1.2 FlexusRDS Custom Policies.....	67
2.2 Buying a FlexusRDS for PostgreSQL Instance.....	68
2.3 Connecting to a FlexusRDS for PostgreSQL Instance.....	71
2.3.1 Using DAS to Connect to a FlexusRDS for PostgreSQL Instance (Recommended).....	71
2.3.2 Using CLI to Connect to a FlexusRDS for PostgreSQL Instance.....	72
2.4 Usage Rules.....	77
2.4.1 Instance Usage Suggestions.....	77
2.4.2 Database Usage Rules.....	80
2.5 Data Migration.....	81
2.5.1 Migrating Databases to FlexusRDS for PostgreSQL Using DRS.....	81
2.5.2 Migrating Data to FlexusRDS for PostgreSQL Using psql.....	82
2.5.3 Migrating Data to FlexusRDS for PostgreSQL Using the Export and Import Functions of DAS.....	85
2.6 Instance Management.....	89
2.6.1 Changing the Name of a FlexusRDS for PostgreSQL Instance.....	89
2.6.2 Rebooting FlexusRDS for PostgreSQL Instances.....	90
2.6.3 Resetting the Administrator Password of a FlexusRDS for PostgreSQL Instance.....	91
2.6.4 Enabling Storage Autoscaling for a FlexusRDS for PostgreSQL Instance.....	92
2.6.5 Binding an EIP to a FlexusRDS for PostgreSQL Instance or Unbinding an EIP from a FlexusRDS for PostgreSQL Instance.....	94
2.6.6 Renewing FlexusRDS for PostgreSQL Instances.....	95
2.6.7 Unsubscribing from a FlexusRDS for PostgreSQL Instance.....	95
2.7 Backups and Restorations.....	96

2.7.1 Creating a Manual Backup for a FlexusRDS for PostgreSQL Instance.....	96
2.7.2 Deleting a Manual Backup of a FlexusRDS for PostgreSQL Instance.....	97
2.7.3 Downloading a Full Backup of a FlexusRDS for PostgreSQL Instance.....	98
2.7.4 Checking and Exporting Backup Information of a FlexusRDS for PostgreSQL Instance.....	102
2.7.5 Restoring a FlexusRDS for PostgreSQL Instance.....	103
2.7.5.1 Restoring a FlexusRDS for PostgreSQL Instance from Backups.....	103
2.7.5.2 Restoring a FlexusRDS for PostgreSQL Instance to a Point in Time.....	104
2.8 Parameters.....	105
2.8.1 Suggestions on Parameter Tuning for a FlexusRDS for PostgreSQL Instance.....	105
2.8.2 Modifying Parameters of a FlexusRDS for PostgreSQL Instance.....	107
2.8.3 Exporting the Parameter List of a FlexusRDS for PostgreSQL Instance.....	108
2.9 Monitoring.....	109
2.9.1 Viewing Monitoring Metrics of a FlexusRDS for PostgreSQL Instance.....	109
2.10 Logs.....	118
2.10.1 Viewing and Downloading Error Logs of a FlexusRDS for PostgreSQL Instance.....	118
2.10.2 Viewing and Downloading Slow Query Logs of a FlexusRDS for PostgreSQL Instance.....	119
2.11 Interconnection with CTS.....	122
2.11.1 FlexusRDS Operations Supported by CTS.....	122
2.11.2 Viewing Traces.....	123
2.12 FlexusRDS for PostgreSQL Tags.....	125
2.13 FlexusRDS for PostgreSQL Quotas.....	125

1 Working with FlexusRDS for MySQL

1.1 Using IAM to Grant Access to FlexusRDS

1.1.1 Creating a User and Granting Permissions for FlexusRDS

This section describes how to use [Identity and Access Management \(IAM\)](#) for fine-grained permissions management for your FlexusRDS 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 FlexusRDS 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 FlexusRDS resources.

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


This section describes the procedure for granting permissions (see [Figure 1-1](#)).

Prerequisites

Learn about the permissions (see [Permissions](#)) supported by FlexusRDS 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 FlexusRDS 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**.

1. **Create a 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.
2. **Log in** and verify permissions.
Log in to the console by using the created user, and verify the user permissions.
 - Go to the FlexusRDS console and click **Buy FlexusRDS Instance** in the upper right corner. 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. If a message appears indicating that you have insufficient permissions to access the service, the **RDS ReadOnlyAccess** policy has already taken effect.

1.1.2 FlexusRDS Custom Policies

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

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.

For details, see [Creating a Custom Policy](#). The following contains examples of common FlexusRDS custom policies.

Example Custom Policies

Example: Allowing users to create manual backups

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

1.2 Buying a FlexusRDS for MySQL Instance

Scenarios

This section describes how to purchase a FlexusRDS instance on the management console.

FlexusRDS for MySQL only supports the yearly/monthly billing mode. It allows you to tailor your compute resources and storage space to your business needs.

Prerequisites

- You have [created a Huawei ID and enabled Huawei Cloud services](#).
- Your account balance is greater than or equal to \$0 USD.

Procedure

Step 1 Go to the [FlexusRDS console](#).

Step 2 If this is your first time to create a DB instance, click **Buy**.

Step 3 Configure the instance information and click **Buy**.

Figure 1-2 Selecting an instance class

Region

CN-Hong KongAP-BangkokAP-SingaporeME-RiyadhTR-IstanbulLA-Sao Paulo1

Database Plan

DB Engine

MySQLPostgreSQL

MySQL 8.0

Instance Class

Standard 2U4G

For micro/small enterprise an...

DB Instance Type

Single

Storage

120 GB

Standard 2U8G

For micro/small enterprise an...

DB Instance Type

Single

Storage

240 GB

Standard 4U8G

For micro/small enterprise an...

DB Instance Type

Single

Storage

240 GB

High-availability 2U4G

For micro/small enterprise w...

DB Instance Type

Primary/Standby

Storage

120 GB

High-availability 2U8G

For micro/small enterprise w...

DB Instance Type

Primary/Standby

Storage

240 GB

Figure 1-3 Selecting the required duration

Network

vpc-default-smb

subnet-default-smb(192.168.0.0/20)

DB Instance Name

flexusrds-71db

If you buy multiple DB instances at a time, they will be named with four digits appended in the format "DB instance name-SN". For example, if the DB instance name is "instance", the first instance will be named "instance-0001", the second "instance-0002", and so on.

Required Duration and Quantity

Required Duration

1 month2 months3 months4 months5 months6 months7 months8 months9 months1 year2 years3 years

☐ Auto-renew Fee deduction and Renewal duration

Quantity

1

Price:

Buy

You will be charged based on the bill.

Table 1-1 Basic information

Parameter	Description
Region	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.
DB Engine	Select MySQL .

Issue 01 (2025-09-26)

Copyright © Huawei Cloud Computing Technologies Co., Ltd.

4

Parameter	Description
Instance Class	An instance class includes vCPUs, memory, storage, and DB instance type. • Storage: The purchased storage space. After a DB instance is purchased, you can configure storage autoscaling. The maximum allowed storage is 4,000 GB. For details, see Enabling Storage Autoscaling for a FlexusRDS for MySQL Instance. • DB Instance Type – Primary/Standby: An HA architecture. 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: A single-node architecture. It is less expensive than a primary/standby DB pair.
Network	A Virtual Private Cloud (VPC) is a virtual network in which your instance is located. A VPC can isolate networks for different workloads. A subnet provides dedicated network resources that are logically isolated from other networks for security purposes. For details about how to create a VPC and subnet, see Creating a VPC and Subnet. After a DB instance is purchased, you can change its VPC and subnet.
DB Instance Name	Must start with a letter and consist of 4 to 64 characters. Only letters (case-sensitive), digits, hyphens (-), underscores (_), and periods (.) are allowed. If you buy 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 <i>instance-0001</i>, <i>instance-0002</i>, and so on. If existing instances' suffixes have already reached up to <i>0010</i>, the new instance names will start from <i>instance-0011</i>.
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.
Auto-renew	• This option is not selected by default. • If you select this option, the auto-renew cycle is determined by the selected required duration.
Quantity	You can buy a maximum of 50 DB instances at a time. If you intend to create primary/standby DB instances and set Quantity to 1, a primary instance and a synchronous standby instance will be created.

Step 4 Confirm the order.

Figure 1-4 Order confirmation

Resource	Configuration	Required D...	Quantity	Price
FlexusRDS	Region: CN-Hong Kong DB Instance Name: flexusdb-7188 DB Engine: MySQL DB Engine Version: 5.5 DB Instance Type: Single Instance Specifications: 2U4G Storage: 120 GB	1 year	1	

Agreement ☒ I have read and agree to the terms and conditions of the [Service Level Agreement \(SLA\)](#).

Price:

[Previous](#) [Pay Now](#)

- If you need to modify your settings, click **Previous**.
- If you do not need to modify your settings, click **Pay Now**.

Step 5 Select a payment method and complete the payment.

Step 6 To view and manage your instance, go to the instance list page.

- When your instance is being created, the status is **Creating**. The status changes to **Available** after the instance is created.
- Automated backup is enabled by default during instance creation. An automated full backup is immediately triggered once your DB instance is created.
- The default administrator account of your DB instance is **root**.
- During instance creation, the system randomly sets a password for the administrator account. You need to **reset the password** before you can connect to the instance.
- The default database port is **3306** and cannot be changed.

----End

1.3 Connecting to a FlexusRDS for MySQL Instance

1.3.1 Using DAS to Connect to a FlexusRDS for MySQL Instance (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 remotely 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** In the instance list, locate the target DB instance and click **Log In** in the **Operation** column.

Figure 1-5 Logging in to an instance



Alternatively, click the instance name in the instance list. On the displayed page, click **Log In** in the upper right corner.

- Step 2** On the displayed login page, enter the username and password and click **Log In**.
----End

1.3.2 Using CLI to Connect to a FlexusRDS for MySQL Instance

Scenarios

You can connect to your DB instance using the MySQL command-line interface (CLI) from a FlexusX instance with a MySQL client installed.

Procedure

- Step 1** **Log in to the FlexusX instance in the same region as your FlexusRDS DB instance.**
- Step 2** Download the MySQL client installation package for Linux to the FlexusX instance. The package [mysql-community-client-5.7.38-1.el6.x86_64.rpm](#) is used as an example.

```
wget https://dev.mysql.com/get/mysql-community-client-5.7.38-1.el6.x86_64.rpm
```

NOTE

A MySQL client running a version later than that of the FlexusRDS for MySQL DB instance is recommended.

- Step 3** Install the MySQL client.

```
rpm -ivh --nodeps mysql-community-client-5.7.38-1.el6.x86_64.rpm
```

NOTE

- If any conflicts occur during the installation, add the **replacefiles** parameter to the command and install the client again.

```
rpm -ivh --replacefiles mysql-community-client-5.7.38-1.el6.x86_64.rpm
```
- If a message is displayed prompting you to install a dependent package during the installation, add the **nodeps** parameter to the command and install the client again.

```
rpm -ivh --nodeps mysql-community-client-5.7.38-1.el6.x86_64.rpm
```

Step 4 Run the following command on the FlexusX instance to connect to the FlexusRDS for MySQL DB instance:


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

Example:

```
mysql -h 192.168.0.1 -P 3306 -u root -p
```

Table 1-2 Parameter description

Parameter	Description
<host>	- Private network connection (recommended): Click the DB instance name and obtain the floating IP address on the Overview page. - Public network connection: Click the DB instance name and obtain the EIP on the Overview page. For details about how to bind an EIP to a DB instance, see Binding an EIP to a FlexusRDS for MySQL Instance or Unbinding an EIP from a FlexusRDS for MySQL Instance. NOTE If the FlexusX instance and the FlexusRDS for MySQL DB instance are in the same region and VPC, you are advised to connect to the DB instance through a private network for higher security. If a DB instance cannot be accessed through a private network, connect to it through a public network. Figure 1-6 Network information Network Information Floating IP Address 192.168.0.85 Private Domain Name .rds.myhuaweicloud.com EIP No EIP bound Bind Database Port 3306 Security Group 1 security group ----- VPC vpc-default-smb Change Subnet subnet-default-smb(192.168.0.0/20) Change

Parameter	Description																																																								
	To connect to your DB instance through an EIP, add the EIP and port 3306 to an inbound rule of security group sg-default-smb. For details, see Adding a Security Group Rule. Figure 1-7 Adding an inbound rule  <table><tr><th>Priority</th><th>Action</th><th>Type</th><th>Protocol & Port</th><th>Source</th><th>Description</th><th>Last Modified</th><th>Operation</th></tr><tr><td>1</td><td>Deny</td><td>IPv4</td><td>All</td><td>--</td><td>Deny All port request Traffic</td><td>Mar 14, 2024 19:28:57 GMT+8</td><td>Modify Replace Delete</td></tr><tr><td>1</td><td>Allow</td><td>IPv4</td><td>TCP: 3306</td><td>0.0.0.0</td><td></td><td>Apr 15, 2024 18:12:48 GMT+8</td><td>Modify Replace Delete</td></tr><tr><td>1</td><td>Allow</td><td>IPv4</td><td>All</td><td>sg-default-ent</td><td></td><td>Mar 14, 2024 19:28:57 GMT+8</td><td>Modify Replace Delete</td></tr><tr><td>1</td><td>Allow</td><td>IPv4</td><td>TCP: 3389</td><td>0.0.0.0</td><td>Permit default Windows remot.</td><td>Mar 14, 2024 19:28:57 GMT+8</td><td>Modify Replace Delete</td></tr><tr><td>1</td><td>Allow</td><td>IPv6</td><td>All</td><td>sg-default-ent</td><td></td><td>Mar 14, 2024 19:28:57 GMT+8</td><td>Modify Replace Delete</td></tr><tr><td>1</td><td>Allow</td><td>IPv4</td><td>TCP: 22</td><td>0.0.0.0</td><td>Permit default Linux SSH port</td><td>Mar 14, 2024 19:28:57 GMT+8</td><td>Modify Replace Delete</td></tr></table>	Priority	Action	Type	Protocol & Port	Source	Description	Last Modified	Operation	1	Deny	IPv4	All	--	Deny All port request Traffic	Mar 14, 2024 19:28:57 GMT+8	Modify Replace Delete	1	Allow	IPv4	TCP: 3306	0.0.0.0		Apr 15, 2024 18:12:48 GMT+8	Modify Replace Delete	1	Allow	IPv4	All	sg-default-ent		Mar 14, 2024 19:28:57 GMT+8	Modify Replace Delete	1	Allow	IPv4	TCP: 3389	0.0.0.0	Permit default Windows remot.	Mar 14, 2024 19:28:57 GMT+8	Modify Replace Delete	1	Allow	IPv6	All	sg-default-ent		Mar 14, 2024 19:28:57 GMT+8	Modify Replace Delete	1	Allow	IPv4	TCP: 22	0.0.0.0	Permit default Linux SSH port	Mar 14, 2024 19:28:57 GMT+8	Modify Replace Delete
Priority	Action	Type	Protocol & Port	Source	Description	Last Modified	Operation																																																		
1	Deny	IPv4	All	--	Deny All port request Traffic	Mar 14, 2024 19:28:57 GMT+8	Modify Replace Delete																																																		
1	Allow	IPv4	TCP: 3306	0.0.0.0		Apr 15, 2024 18:12:48 GMT+8	Modify Replace Delete																																																		
1	Allow	IPv4	All	sg-default-ent		Mar 14, 2024 19:28:57 GMT+8	Modify Replace Delete																																																		
1	Allow	IPv4	TCP: 3389	0.0.0.0	Permit default Windows remot.	Mar 14, 2024 19:28:57 GMT+8	Modify Replace Delete																																																		
1	Allow	IPv6	All	sg-default-ent		Mar 14, 2024 19:28:57 GMT+8	Modify Replace Delete																																																		
1	Allow	IPv4	TCP: 22	0.0.0.0	Permit default Linux SSH port	Mar 14, 2024 19:28:57 GMT+8	Modify Replace Delete																																																		
<port>	3306																																																								
<userName>	root																																																								

Step 5 When the following information is displayed, enter the password of user **root**:

Enter password:

----End

1.4 Usage Rules

1.4.1 Instance Usage Rules

DB Instances

- Primary/Standby
 - A primary/standby pair provides an HA architecture.
 - 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 used.
 - If a fault occurs on a single-node instance, the instance cannot recover in a timely manner.

Database Connection

- Set database parameters based on the complexity of 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.
- Select an instance class and storage space appropriate to your workloads.

Backup and Restoration

- To prevent backup failures, perform manual backups during off-peak hours.
- Both automated and manual backups are deleted after your instance is unsubscribed from.

Routine O&M

- Periodically check slow query logs and error logs to identify problems in advance.
- 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 instance from being accessed from the Internet. If you want to allow the access from the Internet, bind an EIP to your instance.

1.4.2 Database Usage Rules

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 FlexusRDS for MySQL.
 - Reserved words and keywords for MySQL 8.0: <https://dev.mysql.com/doc/refman/8.0/en/keywords.html>
 - Reserved words and keywords for MySQL 5.7: <https://dev.mysql.com/doc/refman/5.7/en/keywords.html>
- 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 recover 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.
- The length of a single row in the table cannot exceed 1024 bytes.
- To avoid cross-partition queries, FlexusRDS 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.
- 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.
- Full-text indexes are not recommended because there are many limitations on full-text indexes for MySQL Community Edition.

Field Design

- Ensure that each table contains no more than 50 fields.
- 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.
- Each field should have the NOT NULL attribute. The default value for the numeric type such as INT is recommended to be 0, and that for the character type such as VARCHAR is recommended to be an empty string.
- 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

- 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 values can increment continuously as the primary key. So, the auto-increment ID column is recommended.
- 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 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 associated for joining 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 FlexusRDS 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, more efficient than SQL statements.
- Avoid implicit conversions because they may cause index 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 five 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. Do not use NATURAL JOIN.
- 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)
- Use ANSI-standard SQL statements instead of MySQL extended SQL statements for DML operations. Common MySQL extended SQL statements include:
 - REPLACE INTO
 - INSERT ... ON DUPLICATE KEY UPDATE
- Stored procedures are not recommended because they are difficult to debug, extend, and transplant.
- To avoid logical dependency on the database, do not use triggers, event schedulers, or views for service logic.
- Large transactions are not recommended. If possible, a transaction should contain no more than five SQL statements because large transactions have problems such as long data lock time, too many caches, and connection consumption.
- TRUNCATE TABLE is faster than DELETE and uses fewer system and log resources. If the table to be deleted does not have a trigger and the entire table needs to be deleted, TRUNCATE TABLE is recommended.
- Do not run the **flush logs** command frequently to prevent automatic binlog deletion failures.

1.5 Data Migration

1.5.1 Migrating Databases to FlexusRDS for MySQL Using DRS

Data Replication Service (DRS) provides multiple data migration solutions to help you migrate data from RDS for MySQL, DDM, GaussDB, TaurusDB, self-managed MySQL databases, self-managed Oracle databases, or MySQL databases built on other clouds to FlexusRDS for MySQL.

You are advised to use DRS to migrate data because it is easy to use and can complete a migration task in minutes. It facilitates data transfer between databases, helping you reduce DBA labor costs and hardware costs.

DRS Migration Solutions

The method for migrating databases to FlexusRDS for MySQL using DRS is the same as that for migrating databases to RDS for MySQL. For details, see [Migration Solution Overview](#).

DRS Migration Billing

- Real-time migration supports only the pay-per-use billing mode.
Real-time migration tasks are free of configuration and data transmission fees in the first seven days, lowering your costs for migrating data to the cloud.
- Real-time synchronization and DR support pay-per-use and yearly/monthly billing modes.
Real-time migration and synchronization will provide long-term discounts, lowering your costs for data transfers.

For more information, see [Data Replication Service Billing](#).

1.5.2 Migrating Data to FlexusRDS for MySQL Using mysqldump

Preparing for Data Migration

You can access your FlexusRDS for MySQL DB instance through an EIP or from a FlexusX instance.

1. Prepare a FlexusX instance for accessing your FlexusRDS for MySQL DB instance or prepare a device for accessing your FlexusRDS for MySQL DB instance through an EIP.
To connect to a FlexusRDS for MySQL instance through an EIP, [bind an EIP](#) to the instance.
2. Install a MySQL client of the same version as your FlexusRDS for MySQL instance on the prepared FlexusX instance or device.

NOTE

A MySQL client will provide mysqldump and mysql.

MySQL system databases **mysql** and **sys** cannot be imported to FlexusRDS for MySQL instances.

Exporting Data

Before migrating a database to FlexusRDS for MySQL, its data needs to be exported.

NOTICE

- The export tool must match the DB engine version.
 - Database migration is performed offline. Before the migration, you have to stop all applications using the source database.
-

Step 1 Log in to the source database.

Step 2 Use the mysqldump tool to export the table structure to an SQL file.

NOTICE

The **mysql** database is required for FlexusRDS for MySQL management. When exporting the table structure, do not specify **--all-database**. Otherwise, a database fault will occur.

```
mysqldump--databases<DB_NAME>--single-transaction --order-by-primary --hex-blob --no-data --routines --events --set-gtid-purged=OFF-u <DB_USER>-p -h<DB_ADDRESS>-P <DB_PORT>|sed -e 's/DEFINER[ ]*=[ ]*[^\]*\*/' -e 's/DEFINER[ ]*.*FUNCTION/FUNCTION/' -e 's/DEFINER[ ]*.*PROCEDURE/PROCEDURE/' -e 's/DEFINER[ ]*.*TRIGGER/TRIGGER/' -e 's/DEFINER[ ]*.*EVENT/EVENT/' ><BACKUP_FILE>
```

- *DB_NAME* indicates the name of the database to be migrated.
- *DB_USER* indicates the database username.
- *DB_ADDRESS* indicates the database address.
- *DB_PORT* indicates the database port.
- *BACKUP_FILE* indicates the name of the file to which the data will be exported.

Enter the database password as prompted.

Example:

```
mysqldump --databases frdsdb --single-transaction --order-by-primary --hex-blob --no-data --routines --events --set-gtid-purged=OFF -u root -p -h 192.168.151.18 -P 3306 |sed -e 's/DEFINER[ ]*=[ ]*[^\]*\*/' -e 's/DEFINER[ ]*.*FUNCTION/FUNCTION/' -e 's/DEFINER[ ]*.*PROCEDURE/PROCEDURE/' -e 's/DEFINER[ ]*.*TRIGGER/TRIGGER/' -e 's/DEFINER[ ]*.*EVENT/EVENT/' > dump-defs.sql
```

Enter password:

After this command is executed, a **dump-defs.sql** file will be generated as follows:

```
[rds@localhost ~]$ ll dump-defs.sql
-rw-r-----. 1 rds rds 2714 Sep 21 08:23 dump-defs.sql
```

Step 3 Use the mysqldump tool to export data to an SQL file.

NOTICE

The **mysql** database is required for FlexusRDS for MySQL management. When exporting data, do not specify **--all-database**. Otherwise, a database fault will occur.

```
mysqldump --databases<DB_NAME>--single-transaction --hex-blob --set-gtid-purged=OFF --no-create-info --skip-triggers-u<DB_USER>-p-h<DB_ADDRESS>-P<DB_PORT>-r<BACKUP_FILE>
```

For details on the parameters in the preceding command, see [2](#).

Enter the database password as prompted.

Example:

```
mysqldump --databases frdsdb --single-transaction --hex-blob --set-gtid-purged=OFF --no-create-info --skip-triggers -u root -p -h 192.168.151.18 -P 3306 -r dump-data.sql
```

After this command is executed, a **dump-data.sql** file will be generated as follows:

```
[rds@localhost ~]$ ll dump-data.sql
-rw-r-----. 1 rds rds 2714 Sep 21 08:23 dump-data.sql
```

----End

Importing Data

You can connect your client to the FlexusRDS for MySQL instance and import exported SQL files into it.

NOTICE

If the source database calls triggers, stored procedures, functions, or events, you must set **log_bin_trust_function_creators** to **ON** on the destination database before importing data.

- Step 1** Log in to the FlexusX instance or device that can access the FlexusRDS for MySQL instance.
- Step 2** Connect to the FlexusRDS for MySQL instance through a client.
- Step 3** Import the table structure into the FlexusRDS for MySQL instance.

```
# mysql -f -h<DB_ADDRESS>-P<DB_PORT>-uroot-p < <BACKUP_DIR>/dump-defs.sql
```

- *DB_ADDRESS* indicates the IP address of the FlexusRDS for MySQL instance.
- *DB_PORT* indicates the database port of the DB instance.
- *BACKUP_DIR* indicates the directory where **dump-defs.sql** is stored.

Example:

```
# mysql -f -h 172.16.66.198 -P 3306 -u root -p < dump-defs.sql
```

Enter password:

NOTE

If you intend to import SQL statements of a table to FlexusRDS, specify a database (*mydb*) in the command. Otherwise, the error message "No database selected" may be displayed.
Example:

```
# mysql -f -h 172.16.66.198 -P 3306 -u root -p mydb < dump-defs.sql
```

Enter password:

- Step 4** Import data into the FlexusRDS for MySQL instance.

```
# mysql -f -h<DB_ADDRESS>-P<DB_PORT>-uroot-p< <BACKUP_DIR>/dump-  
data.sql
```

- *DB_ADDRESS* indicates the IP address of the FlexusRDS for MySQL instance.
- *DB_PORT* indicates the database port of the DB instance.
- *BACKUP_DIR* indicates the directory where **dump-data.sql** is stored.

Example:

```
# mysql -f -h 172.16.66.198 -P 3306 -u root -p < dump-data.sql
```

Enter password:

NOTE

If you intend to import SQL statements of a table to FlexusRDS, specify a database (*mydb*) in the command. Otherwise, the error message "No database selected" may be displayed.
Example:

```
# mysql -f -h 172.16.66.198 -P 3306 -u root -p mydb < dump-defs.sql
```

Enter password:

Step 5 View the import result.

```
mysql> show databases;
```

The following result indicates that database **frdsdb** has been imported.

```
mysql> show databases;  
+-----+  
| Database |  
+-----+  
| information_schema |  
| frdsdb |  
| mysql |  
| performance_schema |  
+-----+  
4 rows in set (0.00 sec)
```

----End

1.5.3 Migrating Data to FlexusRDS for MySQL Using the Export and Import Functions of DAS

Scenarios

Data Admin Service (DAS) is a one-stop management platform that allows you to manage Huawei Cloud databases on a web console. It offers database development, O&M, and intelligent diagnosis, making it easy to use and maintain databases.

To back up or migrate data, you can use DAS to export data from the source database first and then import the data to from your local PC or OBS bucket to the destination database.

For more information, see [Import and Export](#).

Constraints

- Only one file that is no larger than 1 GB can be imported at a time.
- Only data files in the CSV or SQL format can be imported.
- Binary fields such as BINARY, VARBINARY, TINYBLOB, BLOB, MEDIUMBLOB, and LONGBLOB are not supported.
- Data cannot be exported or imported using cross-region OBS buckets.

Exporting Data

- Step 1** In the instance list, locate the target DB instance and click **Log In** in the **Operation** column.
- Step 2** On the displayed login page, enter the username and password and click **Log In**.
- Step 3** On the top menu bar, choose **Import and Export** > **Export**.
- Step 4** On the displayed page, click **Create Task** and choose **Export Database** or **Export SQL Result** as required. The following takes database export as an example.

Alternatively, click **Quick Export** and select the target database. On the displayed page, select a storage path and click **OK**.

Figure 1-8 Quick export

Quick Export

Export Database

Storage ?

No OBS bucket? [Create OBS Bucket](#)

Creating an OBS bucket is free of charge, but storing files in it will incur fees.

OK Cancel

- Step 5** On the displayed page, set parameters as required in areas **Basic Information** and **Advanced Settings**. Then, select the tables to be exported on the right.

Figure 1-9 Creating an export task

Export Database

Basic Information

Database

create_new_db1

☐ Export all tables

Allowed Rows

10,000

File Type

SQL

CSV

Object to Export

Data

Structure

Data and structure

Charset

UTF8

GBK

Storage

abcsdfdf

No OBS bucket? [Create OBS Bucket](#)

Creating an OBS bucket is free of charge, but storing files in it will incur fees.

Options

☒ Combine INSERT statements. (Combine INSERT statements into files, with each file smaller than 5 MB.)

☐ Generate a file for each table. (Downloading table files in the details slows down the export.)

Remarks

Advanced Settings

Tables

Selected Tables: 0

Enter a table name

☐	Table Name	Column	WHERE Clause
☐	asd	Edit	Edit
☐	dddd	Edit	Edit
☐	new_db1_tb1	Edit	Edit
☐	rule_aml_main	Edit	Edit
☐	test1	Edit	Edit

10 / page Total Records: 5 < 1 >

OK

Cancel

NOTE

In a SQL result export task, the executed SQL statements cannot exceed 5 MB.

Export SQL Result

Basic Information

Database

create_new_db1

Allowed Rows

10,000

File Type

SQL-insert

CSV

Charset

UTF8

GBK

Storage

abcsdfdf

No OBS bucket? [Create Bucket](#)

Creating an OBS bucket is free of charge, but storing files in it will incur fees.

Options

☒ Combine INSERT statements. (Combine INSERT statements into files, with each file smaller than 5 MB.)

☐ Generate one file for each result.

SQL to Execute

SELECT * FROM "asd"

Remarks

Advanced Settings

OK

Cancel

NOTE

- Databases are classified into user databases and system databases. System databases cannot be exported. If system database data is required, deploy system database services in a created user database, so that you can export the system database data from the user database.
- DAS connects to your standby database to export data. This prevents the primary database from being affected by data export. However, if the standby database has a high replication delay, the exported data may not be the latest.

Issue 01 (2025-09-26)

Copyright © Huawei Cloud Computing Technologies Co., Ltd.


21

Step 6 After settings are complete, click **OK**.

Step 7 In the task list, view the task ID, type, status, and progress.

Step 8 Click **Details** in the **Operation** column to view task details.

Figure 1-10 Task list



Task ID	Task Type	Database	Started	Ended	File Size	File Type	Status	Elapsed Time	Exported Rows	Progress	Remarks	Operation
c468c5366a0b43276c5366a0b5201	Quick E...	db_01	2020-09-07 20:16:45	2020-09-07 20:16:55	4.53 MB	SQL	Successful	10 secs...	202415	100%		Details Download
c02843785a47e1a7a43785a47e1a2	Database	create_new_db1	2020-09-03 16:50:45	2020-09-03 16:52:14	16.36 MB	SQL	Successful	1 min...	10000	100%		Details Download
7ad59a2d58db460999a2d58db9f8ca	Database	create_new_db1	2020-09-03 16:47:05	2020-09-03 16:47:22	3.94 MB	SQL	Successful	17 secs...	2414	100%		Details Download


----End

Importing Data

Step 1 On the top menu bar, choose **Import and Export > Import**.

Step 2 Import a file from your local PC or an OBS bucket.

Figure 1-11 Creating an import task



Create Task

Import Type: **sql** | CSV

File Source: **Upload file** | Choose from OBS

Attachment Storage: 407154 | No OBS bucket? [Create OBS Bucket](#)

Creating an OBS bucket is free of charge, but storing files in it will incur fees.

Attachment:
Click here to upload a file, or drag one here. (.sql)
Upload only one attachment that is no larger than 1 GB.

Database: db_4eb3_0000

Charset: **Auto Detect** | UTF8 | GBK

Options:
☐ Ignore errors, that is, skip the step where the SQL statement fails to be executed.
☒ Delete the uploaded file upon an import success.

Remarks:

Create Cancel

- From your local PC
In the upper left corner, click **Create Task**. On the displayed page, select an import type, select **Upload file** for **File Source**, set the attachment storage, and upload the file. Then, set other parameters as required.

For security purposes, imported files are stored in OBS buckets.

 NOTE

- To keep your data secure, provide your own OBS bucket to store the attachments you upload. In this way, DAS automatically connects to your OBS bucket for in-memory reading. No data is stored on DAS.
- If you select **Delete the uploaded file upon an import success.**, the file you uploaded will be automatically deleted from the OBS bucket after being imported to the destination database.
- From an OBS bucket

In the upper left corner, click **Create Task**. On the displayed page, select an import type, select **Choose from OBS** for **File Source**, and select a file from the bucket. Then, set other parameters as required.

 NOTE

The file uploaded from an OBS bucket will not be deleted upon an import success.

Step 3 After setting import parameters, click **Create**. Confirm the information again before you click **OK** because original data may be overwritten after data import.

Step 4 View the import progress in the task list or check task details.

----End

1.6 Instance Management

1.6.1 Upgrading the Minor Version of a FlexusRDS for MySQL Instance

Scenarios

FlexusRDS for MySQL supports minor version upgrades to improve performance, add new functions, and fix bugs.

By default, a newly created DB instance uses the latest minor version. For details about minor kernel versions, see [RDS for MySQL Kernel Version Description](#).

Precautions

- 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 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.
- When a minor version is upgraded, the network is intermittently disconnected during the primary/standby switchover. In addition, semi-synchronous replication is performed between the primary and standby instances by default. During the upgrade, there can be two waits of up to 10s for a single

SQL statement to update or write data. To avoid this problem, change the replication mode to asynchronous.

- If primary and standby DB instances are deployed in the same AZ, a minor version upgrade will trigger a failover. If primary and standby DB instances are deployed in different AZs, a minor version upgrade will trigger two failovers.
- 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.
- A minor version can be upgraded in minutes.
- DDL operations on events, such as CREATE EVENT, DROP EVENT, and ALTER EVENT, are not allowed during a minor version upgrade.

During a minor version upgrade, if you are prompted that there are DDL operations being executed on the primary instance, do as follows:

- Change the status of the event whose **STATUS** is **SLAVESIDE_DISABLED** to **ENABLED** or **DISABLED**, and then perform the upgrade.
- Delete the events whose **STATUS** is **SLAVESIDE_DISABLED** and then perform the upgrade.

Constraints

- If the replication delay between primary and standby DB instances is longer than 300 seconds, the minor version cannot be upgraded.
- For primary/standby DB instances, the standby DB instance is upgraded first and then the primary DB instance is upgraded afterwards.
- Minor versions cannot be upgraded for DB instances with abnormal nodes.
- TLSv1.1 is not supported for version 8.0.28 or later versions. To modify the TLS version, change the value of the parameter **loose_tls_version**.

Procedure

Step 1 In the instance list, click the target instance name.

Step 2 On the **Overview** page, click **Upgrade** under the **DB Engine Version** field.

Step 3 In the displayed dialog box, confirm the upgrade information and click **OK**.

Upon submission: The system upgrades the minor version immediately after you have submitted your upgrade request.

During the upgrade, the instance status is **Upgrading minor version**. After the upgrade is complete, the instance status changes to **Available**.

----End

1.6.2 Upgrading a FlexusRDS for MySQL Instance to an RDS DB Instance

Scenarios

You can change a FlexusRDS for MySQL instance to an RDS DB instance to experience more functions. For comparison between FlexusRDS for MySQL and RDS DB instances, see [FlexusRDS Overview](#).

Constraints

- Your account balance must be no less than \$0 USD.
- When a FlexusRDS for MySQL instance is being upgraded to an RDS instance, the following operations cannot be performed on the instance: deleting the instance, rebooting the instance, modifying the parameter template, or creating backups.
- If there is any large transaction being performed during the upgrade, the upgrade may fail.
- The upgrade takes 5 to 15 minutes (during off-peak hours). If it takes an extended period of time, contact customer service.
- The instance will be rebooted and services may be interrupted during the upgrade. The length of the interruption depends on the workloads and how much data there is. Upgrade your instance during off-peak hours.

Parameter Changes

vCPU-related parameters, such as **threadpool_size** and **slave_parallel_workers**, will be reset according to the following rules.

Table 1-3 Parameter value changes with vCPU changes

Scenario	Rule for a Parameter with No Changes Made to Its Value	Rule for a Parameter Changed to a Custom Value
vCPU increase	The parameter will be reset to the default value of the new instance class.	The larger one between the custom value and the default value of the new instance class will be used.

Scenario	Rule for a Parameter with No Changes Made to Its Value	Rule for a Parameter Changed to a Custom Value
vCPU decrease	The parameter will be reset to the default value of the new instance class.	The smaller one between the custom value and the default value of the new instance class will be used.

Memory-related parameters, such as **innodb_buffer_pool_size**, **innodb_log_buffer_size**, **innodb_log_files_in_group**, **max_connections**, **innodb_page_cleaners**, **innodb_buffer_pool_instances**, and **back_log**, will be reset according to the following rules.

Table 1-4 Parameter value changes with memory changes

Scenario	Rule for a Parameter with No Changes Made to Its Value	Rule for a Parameter Changed to a Custom Value
Memory increase	The parameter will be reset to the default value of the new instance class.	The larger one between the custom value and the default value of the new instance class will be used.
Memory decrease	The parameter will be reset to the default value of the new instance class.	The smaller one between the custom value and the default value of the new instance class will be used.

However, values of **innodb_io_capacity** and **innodb_io_capacity_max** will be reset to the default values of the new instance class if no custom values have been

specified for them or they will remain unchanged if you have specified custom values for them.

Procedure

Step 1 In the instance list, locate the target instance and click **Upgrade to RDS** in the **Operation** column.

Step 2 On the displayed page, select the target RDS instance class.

Figure 1-12 Upgrading to RDS

vCPUs Memory	Recommended Connections	TPS QPS	IPv6
☒ 2 vCPUs 4 GB	4,000	--	Supported
☐ 2 vCPUs 4 GB	4,000	--	Supported
☐ 2 vCPUs 8 GB	6,000	--	Supported
☐ 2 vCPUs 16 GB	8,000	--	Not supported
☐ 4 vCPUs 4 GB	4,000	--	Supported
☐ 4 vCPUs 8 GB	6,000	--	Supported

Certain parameter settings will be adjusted based on the specifications of the target instance. [View Details](#)

Current	New
Instance Class Lightweight 2U4G	Instance Class General-enhanced 2U4G
Price Before Upgrade	Price After Upgrade

Amount Due ⓘ Pay Now

Step 3 After confirming the configuration, click **Pay Now** and complete the payment as prompted.

During the upgrade, the instance status is **Upgrading to RDS....** After the upgrade is complete, the instance status changes to **Available**.

Step 4 To manage the RDS instance, click  in the upper left corner of the page and choose **Databases > Relational Database Service**. On the **Instances** page, search for the instance by its name.

NOTE

The RDS instance name and ID are the same as those of the FlexusRDS for MySQL instance before the upgrade.

----End

1.6.3 Changing the Name of a FlexusRDS for MySQL Instance

Scenarios

You can change the name of a DB instance as required.

Procedure

Step 1 In the instance list, locate the instance that you want to edit name for and click  next to the instance name. Then, change the name and click **OK**.

Alternatively, click the target instance name. On the displayed page, click  under the **DB Instance Name** field and change the instance name.

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

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

Step 2 Check the result in the instance list.

----End

1.6.4 Rebooting FlexusRDS for MySQL Instances

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 one or more DB instances at a time on the console.


Constraints

- If the DB instance status is **Abnormal**, the reboot may fail.
- Rebooting a DB instance will reboot the DB engine service, causing service interruptions. During this period, the 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.
- After a primary/standby DB instance is rebooted, it takes about one minute to establish the replication relationship. During this period, some operations, such as changing the instance class, cannot be performed.

Procedure

Step 1 In the instance list, select one or more DB instances (maximum: 50) to be rebooted and click **Reboot** above the instance list.

Step 2 In the displayed dialog box, enter **YES** and click **OK**.

Figure 1-13 Rebooting a DB instance

Step 3 View the instance status. If the status is **Available**, the instance has been rebooted successfully.

----End

1.6.5 Resetting the Administrator Password of a FlexusRDS for MySQL Instance

Scenarios

If you forget the password of the administrator account **root**, you can reset the password. The new password is applied immediately without rebooting the instance.

Precautions

- If the password you provide is regarded as a weak password by the system, you will be prompted to enter a stronger password.
- If you change the administrator password of a primary instance, the administrator password of the standby instance 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.

Procedure

Step 1 In the instance list, locate the target instance and click **Reset Password** in the **Operation** column.

Step 2 In the displayed dialog box, enter a new password and confirm the password.

Figure 1-14 Resetting a password

✕

Reset Password

DB instance ID
dcad90a5d32b4998a0f7f34dddafb954in01

DB Instance Name
flexusrds-71b8

New Password

.....
✕

Confirm Password

.....
✕

i After the password is reset, use the new password to access the DB instance.

Cancel
OK

NOTICE

Keep this password secure. The system cannot retrieve it.

The 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.

- To submit the new password, click **OK**.
- To cancel the reset operation, click **Cancel**.

----End

1.6.6 Enabling Storage Autoscaling for a FlexusRDS for MySQL Instance

Scenarios

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

Constraints

- If your account balance is insufficient, storage autoscaling will fail.

- The storage space can be autoscaled up only when your instance status is **Available** or **Storage full**.
- For a primary/standby DB instance, autoscaling the storage for the primary node will also autoscale the storage for the standby node.
- If a yearly/monthly DB instance has pending orders, it will not be autoscaled.

Procedure

Step 1 In the instance list, click the target instance name to go to the **Overview** page.

Step 2 On the **Overview** page, click **Configure** under the **Configure Autoscaling** field.

Step 3 In the displayed dialog box, click  and configure the required parameters.

Figure 1-15 Configuring autoscaling

×

Configure Autoscaling

Enable Autoscaling



Additional storage will be billed. [Learn more](#) 

Trigger If Available Storage Drops To

10% 

If available storage drops to or below this value, your storage will autoscale.

Increment (%)

20

Enter an integer.

Autoscaling Limit (GB)

4000

Storage can autoscale to no more than 4000 GB.

 If your account balance is insufficient, autoscaling will fail.

Cancel

OK

Table 1-5 Parameter description

Parameter	Description
Enable Autoscaling	If you select this option, autoscaling is enabled.
Trigger If Available Storage Drops To	If the available storage drops to a specified threshold (10%, 15%, or 20%), autoscaling is triggered.
Increment	Autoscaling increment, as a percentage. The default value range is from 5% to 50% .
Autoscaling Limit	The default value range is from 120 to 4000 , in GB. The limit must be no less than the storage of the DB instance.

Step 4 Click **OK**.

----End

1.6.7 Binding an EIP to a FlexusRDS for MySQL Instance or Unbinding an EIP from a FlexusRDS for MySQL Instance

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.

Precautions

- You can buy an EIP [on the network console](#) and bind it to a FlexusRDS for MySQL instance. One EIP can be bound to only one instance. For pricing details, see [Elastic IP pricing details](#).
- If a DB instance has already been bound with an EIP, you must unbind the EIP from the instance first before binding a new EIP to it.

Binding an EIP to a DB Instance

Step 1 In the instance list, click the target instance name to go to the **Overview** page.

Step 2 On the displayed **Overview** page, click **Bind** under the **EIP** field.

Step 3 In the displayed dialog box, all available EIPs are listed. Select the required EIP and click **Yes**.

Step 4 View the EIP that has been bound to the DB instance.

----End

Unbinding an EIP from a DB Instance

Step 1 In the instance list, click the target instance name to go to the **Overview** page.

Step 2 On the displayed **Overview** page, click **Unbind** under the **EIP** field.

Step 3 In the displayed dialog box, click **Yes**.

----End

1.6.8 Changing the VPC and Subnet of a FlexusRDS for MySQL Instance

Scenarios

A Virtual Private Cloud (VPC) is an isolated private virtual network environment provided for cloud databases using network virtualization technologies. A subnet is a unique CIDR block with a range of IP addresses in a VPC. You can change the VPC and subnet of a FlexusRDS for MySQL instance.

Constraints

- Subnets in a secondary CIDR block cannot be selected.
- Changing a VPC or subnet will briefly interrupt services. Perform this operation during off-peak hours.
- VPCs that conflict with the CIDR block preoccupied by the system cannot be selected.
- You can only change the VPC of your instance to one within the same enterprise project.

Procedure

Step 1 In the instance list, click the target instance name to go to the **Overview** page.

Step 2 In the **Network Information** area, click **Change** under **VPC** or **Subnet**.

Figure 1-16 Network Information

Network Information

Floating IP Address

192.168.0.85

Private Domain Name

[Redacted]

[Redacted].rds.myhuaweicloud.com 

EIP

No EIP bound [Bind](#)

Database Port

3306 

Security Group

1 security group

VPC

vpc-default-smb [Change](#)

Subnet

subnet-default-smb(192.168.0.0/20) [Change](#)

Step 3 In the displayed dialog box, select a VPC and subnet, and click **OK**.

Figure 1-17 Modify Network

Modify Network

- Subnets in a secondary CIDR block cannot be selected.
- Modifying a VPC or subnet will briefly interrupt services. Perform this operation during off-peak hours.

VPC

vpc-a33b

* Subnet

subnet-c42d(192.168.1.0/24)

Cancel

OK

-----End

1.6.9 Renewing FlexusRDS for MySQL Instances

Scenarios

You can renew one or multiple yearly/monthly DB instances at a time.

Procedure

- Step 1

In the instance list, select the target DB instance and click **Renew** above the instance list.

You can also click the target instance name to go to the **Overview** page to renew the instance.
- Step 2

In the displayed dialog box, confirm the instance to be renewed.

Figure 1-18 Renewing an instance

Renew

You are renewing instance flexusrds-71b8.

Name/ID	Status	DB Instance Type	Storage (GB)
flexusrds-71b8	Available	Single	120
dcad90a5d32b4998a0f7f34dd...		2 vCPUs 4 GB	

Cancel

OK

- Step 3

Click **OK** to go to the renewal page and renew the instance.

-----End

1.6.10 Unsubscribing from a FlexusRDS for MySQL Instance

Scenarios

To delete a DB instance billed on the yearly/monthly basis, you need to unsubscribe the order. For unsubscription fees, see [Unsubscription Rules](#).

Constraints

- A DB instance cannot be unsubscribed when any operations are being performed on it. It can be unsubscribed only after the operations are complete.
- If a backup of a DB instance is being restored, the instance cannot be unsubscribed.

Procedure

Step 1 In the instance list, select the target instance and click **Unsubscribe** above the instance list.

Step 2 In the displayed dialog box, enter **YES**.

Step 3 Click **OK**.

After you unsubscribe from an instance order, the instance will be deleted and it is no longer displayed in the instance list.

----End

1.7 Backups and Restorations

1.7.1 Creating a Manual Backup for a FlexusRDS for MySQL Instance

Scenarios

FlexusRDS for MySQL allows you to create manual backups for an available DB instance. You can use these backups to restore data.

Constraints

- You can create manual backups only when your account balance is no less than \$0 USD.
- Unsubscribing from a DB instance will delete its automated and manual backups.
- 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. For the billing details, see [How Is FlexusRDS Backup Data Billed?](#)

Procedure

- Step 1** In the instance list, click the target instance name.
- Step 2** Click **Backups & Restorations** and then click **Create Backup**.
- Step 3** In the displayed dialog box, enter a backup name and description, and click **OK**.

Figure 1-19 Creating a backup

Create Backup

1

When the DB instance is being backed up, data is copied and then compressed and uploaded to OBS at an average speed of 300 MB/s. Creating a backup increases the disk I/O load. Perform this operation during off-peak hours.

DB instance ID

dcad90a5d32b4998a0f7f34dddafb954in01

DB Instance Name

flexusrds-71b8

★ Backup Name

?

Description

0/256

?

Cancel

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: >!<"&'=
- The time required for creating a manual backup depends on the amount of data.

Step 4 View and manage the created backup on the **Backups & Restorations** page.

----End

1.7.2 Deleting a Manual Backup of a FlexusRDS for MySQL Instance

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 In the instance list, click the target instance name.

Step 2 Click **Backups and Restorations**.

Step 3 Locate a manual backup and click **Delete** in the **Operation** column.

The following backups cannot be deleted:

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

Step 4 In the displayed dialog box, click **Yes**.

----End

1.7.3 Downloading a Full Backup of a FlexusRDS for MySQL Instance

Scenarios

You can download manual and automated full backup files in .qp format for local storage.

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
- In the instance list, click the target instance name.
- Step 2
- Click **Backups and Restorations**.
- Step 3
- Locate the backup to be downloaded and click **Download** in the **Operation** column.
- Step 4
- In the displayed dialog box, select **Use OBS Browser+** for **Download Method** and click **OK**.

Figure 1-20 Using OBS Browser+






Figure 1-21 Download guide



1.
- Download OBS Browser+ following step 1 provided on the download guide page.


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

Figure 1-22 Logging in to OBS Browser+



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

Figure 1-23 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 the Resources in It](#).

- 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 1-24 Downloading a backup




----End

Method 2: Using Current Browser

- Step 1** In the instance list, click the target instance name.
- Step 2** Click **Backups and Restorations**.
- Step 3** Locate the backup to be downloaded and click **Download** in the **Operation** column.
- Step 4** In the displayed dialog box, select **Use Current Browser** for **Download Method** and click **OK**.

Figure 1-25 Using the current browser




----End

Method 3: Using Download URL

- Step 1** In the instance list, click the target instance name.
- Step 2** Click **Backups and Restorations**.
- Step 3** Locate the backup to be downloaded and click **Download** in the **Operation** column.

- Step 4** 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 1-26 Using the download URL



- You can also run the following command to download backup files:
wget -O *FILE_NAME* --no-check-certificate "*DOWNLOAD_URL*"
The parameters in the command are as follows:
FILE_NAME: indicates 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: indicates the location of the backup file to be downloaded. If the location contains special characters, escape is required.

----End

1.7.4 Checking and Exporting Backup Information of a FlexusRDS for MySQL Instance

Scenarios

You can export backup information of FlexusRDS for MySQL instances to an Excel file for further analysis. The exported information includes the backup ID, backup name, backup type, backup method, backup start and end times, status, size, and description.

For details about how to export backup data, see [Downloading a Full Backup of a FlexusRDS for MySQL Instance](#).

Procedure

Step 1 In the instance list, click the target instance name.

Step 2 Click **Backups and Restorations**.

Step 3 Click  above the backup list to export backup information.

- If you want to export specified backup records, you can first select them and then export them. You can only select and export the backup records displayed on the current page.
- If you do not select any backup records, all backup records are exported by default. (A maximum of the first 5,000 backup records can be exported. If you want to export more, select the records and export them.)
- The backup information is exported to an Excel file for your further analysis.

Figure 1-27 Backup information

	A	B	C	D	E	F	G	H
1	Backup ID	Backup Name	Backup Type	Backup Method	Backup Time	Status	Size	Description
2	f08c648504944ba88ce2d143c27869d5br01	mysql-flexusrds-71b8	Automated	Physical backup	Jun 12, 2024 10:32:09	Completed	5.93 MB	--
3								

----End

1.7.5 Restoring a FlexusRDS for MySQL Instance

1.7.5.1 Restoring a FlexusRDS for MySQL Instance from Backups

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, the backup file is downloaded from OBS and then restored to the DB instance at an average speed of 100 MB/s.

Constraints

- Constraints on restoring data to a new DB instance:
 - You can restore data to a new instance only when your account balance is greater than or equal to \$0 USD. You will pay for the new instance specifications.
 - The storage space of the new instance should be no less than that of the original instance.
 - 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.
- Constraints on restoring data to the original DB instance:
 - 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 data on it and cause the DB instance to be unavailable during the restoration.

Procedure

Step 1 In the instance list, click the target instance name.


Step 2 Click **Backups and Restorations**.

Step 3 Locate the backup to be restored and click **Restore** in the **Operation** column.

Step 4 Select a restoration method and click **OK**.

- Create New Instance

Figure 1-28 Restoring data to a new instance



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

- The DB engine version of the new instance is the same as that of the original instance.
- You can select a new instance class for the new instance, but the storage of the new instance must be no less than that of the original instance.
- For details about the configuration items on the purchase page, see [Buying a FlexusRDS for MySQL Instance](#).

NOTE

To change the instance class, choose a new one for the new instance on the **Create New Instance** page.

- Restore to Original

Figure 1-29 Restoring data to the original instance

Restore Data

When you restore the 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 100 MB/s.

DB Instance	Name/ID	Backup Name	DB Engine Version
	flexusrds-71b8 dcad90a5d32b4998a0f7f34ddda9b954in...	mysql-flexusrds-71b8-2024...	MySQL 8.0.28

Restoration Method:

☒ 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.

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

Step 5 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 instance status changes from **Creating** to **Available**.
The new DB instance is independent from the original one. After the new instance is created, a full backup will be automatically triggered.
- Restore to Original
In the instance list, the status of the original DB instance changes from **Restoring** to **Available**. After the restoration is complete, a full backup will be automatically triggered.

----End

1.7.5.2 Restoring a FlexusRDS for MySQL Instance to a Point in Time

Scenarios

You can use automated backups to restore an instance to a specific point in time.

When you enter the time point that you want to restore the DB instance to, FlexusRDS 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 100 MB/s.

Constraints

- Do not run the **reset master** command on instances within their lifecycle. Otherwise, an exception may occur during the point-in-time recovery (PITR).
- Constraints on restoring data to a new DB instance:
 - You can restore data to a new instance only when your account balance is greater than or equal to \$0 USD. You will pay for the new instance specifications.
 - The storage space of the new instance should be no less than that of the original instance.
 - 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, contact customer service.
- Constraints on restoring data to the original DB instance:
 - Restoring to the original DB instance will overwrite data on it and cause the DB instance to be unavailable during the restoration.

Procedure

Step 1 In the instance list, click the target instance name.

Step 2 Click **Backups and Restorations**.

Step 3 Click **Restore** above the backup list.

Step 4 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

Figure 1-30 Restoring data to a new instance

Restore Data ×

Information: When you restore the 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 100 MB/s. ×

Restore To: Apr 28, 2024 📅

Time Range: Apr 28, 2024 00:00:00 – Apr 28, 2024 21:15:38 GMT+08:00 ▼

Time Point: 21:15:38 🕒

Restoration Method: **Create New Instance** Restore to Original

OK **Cancel**

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

- The DB engine version of the new instance is the same as that of the original instance.

- You can select a new instance class for the new instance, but the storage of the new instance must be no less than that of the original instance.
- For details about the configuration items on the purchase page, see [Buying a FlexusRDS for MySQL Instance](#).

NOTE

To change the instance class, choose a new one for the new instance on the **Create New Instance** page.

- Restore to Original

Figure 1-31 Restoring data to the original instance

Restore Data

When you restore the 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 100 MB/s.

Restore To: Apr 28, 2024

Time Range: Apr 28, 2024 00:00:00 – Apr 28, 2024 21:01:48 GMT+08:00

Time Point: 21:01:48

Restoration Method:

☒ 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.

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

Step 5 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 instance status changes from **Creating** to **Available**.

The new DB instance is independent from the original one. After the new instance is created, a full backup will be automatically triggered.

- Restore to Original

In the instance list, 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.

----End

1.8 Parameters

1.8.1 Suggestions on Parameter Tuning for a FlexusRDS for MySQL Instance

Parameters are key configuration items in a database system. Improper parameter settings may adversely affect database performance. This section describes some important parameters for your reference. For details, visit the [MySQL official website](#).

For details on how to modify FlexusRDS parameters on the console, see [Modifying Parameters of a FlexusRDS for MySQL Instance](#).

Sensitive Parameters

- **innodb_flush_log_at_trx_commit**

Default value: **1**

Function: Controls the balance between strict ACID compliance for commit operations and higher performance. The default setting of **1** is required for full ACID compliance. Logs are written and flushed to disks at each transaction commit. If the value is set to **0**, logs are written and flushed to disks once per second. If the value is set to **2**, logs are written at each transaction commit and flushed to disks every two seconds.

Impact: If this parameter is not set to **1**, data security is not guaranteed. If the system fails, data may be lost.

Recommended value for POC: **2**

- **sync_binlog**

Default value: **1**

Function: Controls how often the MySQL server synchronizes binary logs to the disk. The default setting of **1** requires synchronization of the binary log to the disk at each transaction commit. If the value is set to **0**, synchronization of the binary log to the disk is not controlled by the MySQL server but relies on the OS to flush the binary log to the disk. This setting provides the best performance. However, if a power failure occurs or the OS crashes, all binary log information in **binlog_cache** will be lost.

Impact: If this parameter is not set to **1**, data security is not guaranteed. If the system fails, binary logs may be lost.

Recommended value for POC: **1000**

- **innodb_buffer_pool_size**

Default value: Varies depending on the DB instance classes.

Function: Specifies the size of the InnoDB buffer pool. The InnoDB buffer pool is used to cache table and index data. Increasing the value of this parameter reduces disk I/O.

Impact: Setting this parameter to a large value may cause system breakdown. Exercise caution when changing this parameter value.

Recommended value for POC: 70% to 75% of the memory for your DB instances with 32 GB memory or above

Performance Parameters

- The values of **innodb_spin_wait_delay** and **query_alloc_block_size** are determined by the DB instance specifications. If you increase their values, database performance may be affected.
- **max_connections**: indicates the total number of clients that can be concurrently connected. The default value of this parameter depends on the system architecture. System built-in connections occupy some connections specified by this parameter. To prevent concurrent connection conflicts, you are advised not to set this parameter to a value less than 30. This parameter cannot be set to a value smaller than the number of current connections.
- The default values of the following parameters are determined by the DB instance specifications: **innodb_buffer_pool_size**, **max_connections**, and **back_log**. These parameter values are **default** before being specified.
- The values of **innodb_io_capacity_max** and **innodb_io_capacity** are determined by the storage type. These parameter values are **default** before being specified.

Associated Parameters

- **character_set_server**: If you change the value of this parameter, the system changes the value of **collation_server** accordingly.
The parameters **character_set_server** and **collation_server** are correlated with each other. For example, for MySQL 5.7, when **character_set_server** is **latin1**, the default value of **collation_server** is **latin1_swedish_ci**. The **collation_server** value must start with **latin1**.
- **innodb_io_capacity**: The value of this parameter must be less than or equal to the value of **innodb_io_capacity_max**. For example, if **innodb_io_capacity_max** is set to **2000**, the maximum value of **innodb_io_capacity** is **2000**.

Constraints on Parameter Modification

- When the **innodb_adaptive_hash_index** and **innodb_buffer_pool_size** parameters are modified at the same time, the value of **innodb_adaptive_hash_index** will fail to be changed from **OFF** to **ON**.
- If **innodb_buffer_pool_instances** is set to **2**, the value of **innodb_buffer_pool_size** must be greater than or equal to 1 (unit: GB).

Other Parameters

- **max_prepared_stmt_count**: limits the upper limit of prepared statements. Too many prepared statements consume server memory resources. If this parameter is set to a small value, your DB instance may be vulnerable to the denial of service (DoS) attacks. You are advised to change this parameter value based on service requirements.

- The values of the following parameters will be adjusted based on kernel rules:
 - **key_cache_age_threshold**: automatically adjusted to a multiple of 100.
 - **join_buffer_size** and **key_cache_block_size**: automatically adjusted to multiples of 128.
 - **query_prealloc_size**, **innodb_log_buffer_size**, **max_allowed_packet**, and **thread_stack**: automatically adjusted to multiples of 1024.
 - **read_buffer_size**, **read_rnd_buffer_size**, **binlog_cache_size**, and **binlog_stmt_cache_size**: automatically adjusted to multiples of 4096.
- **innodb_strict_mode**: restricts the InnoDB check policy. The default value is **OFF**.
- **binlog_rows_query_log_events**: controls whether to write original SQL statements into binlogs. If this parameter is set to **ON**, database performance may deteriorate when a large amount of data is updated. Before you change the parameter value, consider the compatibility with tools such as Otter.

1.8.2 Modifying Parameters of a FlexusRDS for MySQL Instance

Scenarios

You can change parameter values in a custom parameter template and apply it to optimize database performance.

Modifying a Single Parameter

- Step 1** In the instance list, click the target instance name.
- Step 2** Click the **Parameters** tab.
- Step 3** In the parameter list, locate the parameter you want to modify and click **Modify** in the **Operation** column.

NOTICE

After you modify a parameter, check the value in the **Effective upon Reboot** column.

- If the value is **Yes** and the instance status in the instance list 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 the value is **No**, the modifications take effect immediately.
-

Figure 1-32 Parameters

Parameter Name	Effective upon Reboot	Value	Allowed Values	Description	Operation
auto_increment_increment	No	1	1-65,535	auto_increment_increment and auto_increment_offset control operations in the AUTO_INCREMENT.	Modify
auto_increment_offset	No	1	1-65,535	auto_increment_increment and auto_increment_offset control operations in the AUTO_INCREMENT.	Modify
back_log	Yes	1000	1-65,535	Number of valid connection requests that MySQL can have. This comes into play when the number of connections exceeds the 'max_connections' value.	Modify
binlog_cache_size	No	32768	4,096-16,777,216	The cache capacity used to store SQL statements for the binary log in a transaction.	Modify
binlog_checksum	No	CRC32	NONE, CRC32	When this variable is enabled, the primary server writes a checksum for each event in the binary log.	Modify
binlog_rows_query_log_events	No	OFF	ON, OFF	Once the parameter is enabled, a MySQL 5.6.2 or later server writes information to log events.	Modify
binlog_stmt_cache_size	No	32768	4,096-16,777,216	This variable determines the size of the cache for the binary log to hold non-transactional statements.	Modify
block_encryption_mode	No	aes-128-ecb	aes-128-ecb, aes-192-ecb, aes-256-ecb, aes-128-ctr, aes-192-ctr, aes-256-ctr, aes-128-gcm, aes-192-gcm, aes-256-gcm, aes-128-ccm, aes-192-ccm, aes-256-ccm, aes-128-ccm3, aes-192-ccm3, aes-256-ccm3, aes-128-ccm4, aes-192-ccm4, aes-256-ccm4	Controls the block encryption mode for block-based algorithms such as AES. It affects encryption and decryption.	Modify
bulk_insert_buffer_size	No	8388608	0-16,445,744,073,709,551,615	Limits the size of the MyISAM cache tree in bytes per thread.	Modify
character_set_server	Yes	utf8mb4	utf8latin1, gbk, utf8mb4	The server's default character set.	Modify

- To save the modifications, click **Confirm**. In the displayed dialog box, click **Yes**.
- To cancel the modifications, click **Cancel**.

To view the change history, click **Change History** above the parameter list. The change history of the last seven days is displayed.

----End

Modifying Parameters in Batches

Step 1 In the instance list, click the target instance name.

Step 2 Click the **Parameters** tab.

Step 3 Switch on the batch modification switch . A maximum of 30 parameters can be modified at a time.

NOTICE


After you modify a parameter, check the value in the **Effective upon Reboot** column.

- If the value is **Yes** and the instance status in the instance list 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 the value is **No**, the modifications take effect immediately.

Figure 1-33 Modifying parameters



- To save your modifications, click **Save**. In the displayed dialog box, click **Yes**.
- To cancel your modifications, click **Cancel**. In the displayed dialog box, click **Yes**.

To view the change history, click **Change History** above the parameter list. The change history of the last seven days is displayed.

----End

1.8.3 Exporting the Parameter List of a FlexusRDS for MySQL Instance


Scenarios

You can also export the parameter information (including parameter names, values, and descriptions) of a DB instance to a CSV file for viewing and analyzing details.

Procedure

- Step 1
- In the instance list, click the target instance name.
- Step 2
- Click the **Parameters** tab.
- Step 3
- Click **Export** above the parameter list.

Figure 1-34 Export a parameter list



Step 4 In the displayed dialog box, enter a file name and click **OK**.

 NOTE

The file name can contain 4 to 81 characters.

----End

1.9 Monitoring

1.9.1 Viewing Monitoring Metrics of a FlexusRDS for MySQL Instance

Scenarios

This section describes how to view monitoring metrics of FlexusRDS for MySQL instances and configure alarm rules. You can customize objects to be monitored and notification policies so that you can closely monitor your instances.

Viewing Metrics

- Step 1** In the instance list, click the target instance name.
- Step 2** Click the **Monitoring** tab and view the monitoring metrics of the instance.

Figure 1-35 Metrics




Table 1-6 Supported metrics

Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
CPU Usage	CPU usage of the monitored object	0-100%	FlexusRDS for MySQL instance	1 minute

Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
Memory Usage	Memory usage of the monitored object	0-100%	FlexusRDS for MySQL instance	1 minute
Storage Space Usage	Storage usage of the monitored object	0-100%	FlexusRDS for MySQL instance	1 minute
IOPS	Average number of I/O requests processed by the system in a specified period	≥ 0 counts/s	FlexusRDS for MySQL instance	1 minute
Network Input Throughput	Incoming traffic in bytes per second	≥ 0 bytes/s	FlexusRDS for MySQL instance	1 minute
Network Output Throughput	Outgoing traffic in bytes per second	≥ 0 bytes/s	FlexusRDS for MySQL instance	1 minute
Total Connections	Total number of connections that attempt to connect to the MySQL server	≥ 0 counts	FlexusRDS for MySQL instance	1 minute
Current Active Connections	Number of connections that are not in the sleep state	≥ 0 counts	FlexusRDS for MySQL instance	1 minute

----End

Configuring Alarm Rules

Step 1 On the metric card, click + to go to the page for configuring alarm rules.

Figure 1-36 Selecting a metric



Step 2 On the **Create Alarm Rule** page, configure required parameters.

- **Name:** The system generates a random name for the alarm rule. You can change it as needed.
- **Description:** Add supplementary information about the rule.
- **Alarm Policy:** Specify the conditions for triggering an alarm. An alarm will be triggered if the metric data in the specified number of consecutive periods reaches the specified threshold.
- **Alarm Notification:** Specify whether to notify users when alarms are triggered. Notifications can be sent by email, text message, or HTTP/HTTPS message.

For details about how to configure alarm notification, see [Creating an Alarm Rule](#) in *Cloud Eye User Guide*.

Step 3 Click **Create**.

----End

1.10 Logs

1.10.1 Viewing Operation Logs of a FlexusRDS for MySQL Instance

FlexusRDS for MySQL allows you to view logs of key operations for future query, audit, and backtracking.

Operation logs generated within the last 7 days can be viewed.

Viewing Log Details

Step 1 In the instance list, click the target instance name.

Step 2 Click the **Logs** tab. On the **Operation Logs** tab page, view operation logs.

- You can select a log level in the upper right corner to view logs of the selected level.

NOTE

You can view logs of the following operations on FlexusRDS instances:

- resetPassword
 - instanceRestore
 - instanceRestart
 - updateParameterGroup
 - backupsDownload
- You can click  in the upper right corner to view operation logs generated in different time segments.

----End

1.10.2 Viewing and Downloading Error Logs of a FlexusRDS for MySQL Instance

FlexusRDS for MySQL 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 In the instance list, click the target instance name.

Step 2 Click the **Logs** tab. On the **Error Logs** tab page, 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 FlexusRDS instances, the following levels of logs are displayed:


- All log levels
 - ERROR
 - WARNING
 - NOTE
- Error logs are displayed in log loading mode. There is no upper limit on the number of log records displayed within the query time range, and the total number of log records is not displayed.
 - You can click  in the upper right corner to view logs generated in different time segments.
 - 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** In the instance list, click the target instance name.
- Step 2** Click the **Logs** tab and click **Download Logs** on the right.
- Step 3** Locate the log file whose status is **Preparation completed** and click **Download** in the **Operation** column.

Figure 1-37 Downloading an Error Log



- 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.
- 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 redownload the log, click **OK**.
- The downloaded logs contain only the logs of the primary node.

-----End

1.10.3 Viewing and Downloading Slow Query Logs of a FlexusRDS for MySQL Instance

Scenarios

Slow query logs record statements that exceed the **long_query_time** value (1 second by default). You can view log details and statistics 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 7 days can be viewed.

FlexusRDS for MySQL supports the following statement types:

- All statement types
- SELECT
- INSERT
- UPDATE
- DELETE
- CREATE

Parameter Description

Table 1-7 Parameters related to 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 query without indexes. The default value is OFF .
log_throttle_queries_not_using_indexes	Limits the number of SQL statements without indexes per minute that can be written to the slow query log. The default value is 0 .

Viewing Log Details

Step 1 In the instance list, click the target instance name.

Step 2 Click the **Logs** tab. On the **Slow Query Logs** tab page, view details about slow SQL statements.


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 this parameter is changed from 1s to 0.1s, FlexusRDS 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.
- Slow query logs are displayed in log loading mode. There is no upper limit on the number of log records displayed within the query time range, and the total number of log records is not displayed.
- 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** In the instance list, click the target instance name.
- Step 2** Click the **Logs** tab and click **Download Logs** on the right.
- Step 3** In the displayed dialog box, click **Slow Query Logs**.
- Step 4** Locate the log file whose status is **Preparation completed** and click **Download** in the **Operation** column.

Figure 1-38 Downloading a Slow Query Log

- 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 redownload the log, click **OK**.
- The downloaded logs contain only the logs of the primary node.

----End

1.11 Interconnection with CTS

1.11.1 FlexusRDS Operations Supported by CTS

With Cloud Trace Service (CTS), you can record operations associated with FlexusRDS instances for later query, audit, and backtrack operations.

Table 1-8 FlexusRDS operations that can be recorded by CTS

Operation	Resource Type	Trace Name
Creating a DB instance or restoring data to a new instance	instance	createInstance
Enabling autoscaling	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
Binding or unbinding an EIP	instance	setOrResetPublicIP
Adding a tag	instance	createTag
Deleting a tag	instance	deleteTag
Editing a tag	instance	modifyTag
Deleting a DB instance	instance	deleteInstance

Operation	Resource Type	Trace Name
Downloading a backup (using OBS)	backup	downLoadSnapshot
Downloading a backup (using a browser)	backup	backupsDownLoad
Deleting a frozen DB instance	all	rdsUnsubscribeInstance
Freezing a DB instance	all	rdsfreezeInstance
Renewing a DB instance	all	bssUpdateMetadata

1.11.2 Viewing Traces

Scenarios

After Cloud Trace Service (CTS) is enabled, it starts recording operations on cloud resources. You can check the records generated over the last 7 days on the CTS console.

This section describes how to check the operation records generated over the last 7 days on the CTS console.

For details about how to view audit logs, see [Querying Real-Time Traces](#).

Prerequisites

Before using CTS, you need to enable it. For details, see [Enabling CTS](#).

Procedure

Step 1 [Log in to the management console](#).

Step 2 In the upper left corner of the page, click  and choose **Management & Governance** > **Cloud Trace Service** to go to the **Trace List** page on the CTS console.

Step 3 In the time range drop-down list above the trace list, select a desired query time range: **Last 1 hour**, **Last 1 day**, or **Last 1 week**. You can also select **Custom** to specify a custom time range within the last seven days.

Step 4 The search box above the trace list supports advanced queries. Combine one or more filters to refine your search.

Table 1-9 Trace filtering parameters

Parameter	Description
Trace Name	Name of a trace. The entered value is case-sensitive and requires an exact match. Fuzzy matching is not supported.
Trace Source	Cloud service name abbreviation. The entered value is case-sensitive and requires an exact match. Fuzzy matching is not supported.
Resource Name	Name of a cloud resource involved in a trace. The entered value is case-sensitive and requires an exact match. Fuzzy matching is not supported.
Resource ID	ID of a cloud resource involved in a trace. The entered value is case-sensitive and requires an exact match. Fuzzy matching is not supported.
Trace ID	Value of the trace_id parameter for a trace reported to CTS. The entered value requires an exact match. Fuzzy matching is not supported.
Resource Type	Type of a resource involved in a trace. The entered value is case-sensitive and requires an exact match. Fuzzy matching is not supported.
Operator	User who triggers a trace. Select one or more operators from the drop-down list. If the value of trace_type in a trace is SystemAction , the operation is triggered by the service and the trace's operator may be empty.
Trace Status	Select one of the following options from the drop-down list: ● normal: The operation succeeded. ● warning: The operation failed. ● incident: The operation caused a fault that is more serious than a normal failure, for example, causing other faults.
Enterprise Project ID	ID of the enterprise project to which a resource belongs. To check enterprise project IDs, go to the Enterprise Project Management Service (EPS) console and choose Project Management in the navigation pane.

Parameter	Description
Access Key	Temporary or permanent access key ID. To check access key IDs, hover over your username in the upper right corner of the console and select My Credentials from the pop-up list. On the displayed page, choose Access Keys in the navigation pane.

Step 5 Check the traces that meet the search criteria.

Step 6 Click a trace name. Details about the trace are displayed in the dialog box on the right.

Step 7 Click **Export** to export all traces in the query result as an .xlsx file.

For details about key fields in a trace structure, see [Trace Structure](#) and [Example Traces](#) in the *Cloud Trace Service User Guide*.

----End

1.12 FlexusRDS for MySQL Tags

Scenarios

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.

- Log in to the management console 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.
- Each DB instance can have up to 20 tags.

Editing a Tag

Step 1 In the instance list, click the target instance name.

Step 2 Click the **Tags** tab and click **Edit Tag**.

Step 3 In the displayed dialog box on the right, click **Add Tag**, enter a tag key and value, and click **OK**.

- 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 (optional) can consist of up to 255 characters and can include letters, digits, spaces, and the following characters: _ . : / = + - @.

Step 4 After a tag has been added, you can view and manage it on the **Tags** page.

----End

Deleting a Tag

Step 1 In the instance list, click the target instance name.

Step 2 Click the **Tags** tab and click **Edit Tag**.

Step 3 In the displayed dialog box on the right, select the tag to be deleted and click **Delete**.

Step 4 Click **OK**.

After a tag has been deleted, it will no longer be displayed on the **Tags** page.

-----End

1.13 FlexusRDS for MySQL 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 FlexusRDS 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 [Log in to the management console](#).

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

Step 3 In the upper right corner of the page, choose **Resources** > **My Quotas**.

Figure 1-39 My quotas



Step 4 On the **Quotas** page, view the used and total quotas of each type of resources.

-----End

Increasing Quotas

- Step 1 Log in to the management console.
- Step 2 Click  in the upper left corner and select a region and project.
- Step 3 In the upper right corner of the page, choose **Resources > My Quotas**.
- Step 4 In the upper right corner of the page, click **Increase Quota**.

Figure 1-40 Increasing quotas

Service Quota 				Resource Type	Used Quota	Total Quota
Auto Scaling	AS group			AS configuration	0	0
				AS configuration	0	0
Image Management Service	Image			Image	0	0
				Image	0	0
Cloud Container Engine	Cluster			Cluster	0	0
				Cluster	0	0
FunctionGraph	Function			Function	0	0
				Function	0	0
CodeArts	Code storage(RD)			Code storage(RD)	0	0
				Code storage(RD)	0	0
Data	Data			Data	3	3
				Data	3	3
Elastic Volume Service	Disk capacity(CD)			Disk capacity(CD)	120	120
				Disk capacity(CD)	120	120
Storage Disaster Recovery Service	Storage			Storage	4	4
				Storage	4	4
Cloud Backup Service	Protection group			Protection group	0	0
				Protection group	0	0
Cloud Backup Service	Recovery plan			Recovery plan	0	0
				Recovery plan	0	0
Cloud Backup Service	Recovery capacity(CD)			Recovery capacity(CD)	0	0
				Recovery capacity(CD)	0	0
Elastic File Service	Storage			Storage	0	0
				Storage	0	0
Elastic File Service	File system			File system	0	0
				File system	0	0
Elastic File Service	File system capacity(CD)			File system capacity(CD)	0	0
				File system capacity(CD)	0	0
CDN	Domain name			Domain name	0	0
				Domain name	0	0
CDN	File URL refreshing			File URL refreshing	0	0
				File URL refreshing	0	0
CDN	Directory URL refreshing			Directory URL refreshing	0	0
				Directory URL refreshing	0	0

- Step 5 On the **Create Service Ticket** page, configure parameters as required.
- In the **Problem Description** area, enter the required quota and reason for the quota adjustment.
- Step 6 After all required parameters are configured, select the agreement and click **Submit**.
- End

2 Working with FlexusRDS for PostgreSQL

2.1 Permissions Management

2.1.1 Creating a User and Granting Permissions for FlexusRDS

This section describes how to use [Identity and Access Management \(IAM\)](#) for fine-grained permissions management for your FlexusRDS 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 FlexusRDS 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 FlexusRDS resources.

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


This section describes the procedure for granting permissions (see [Figure 2-1](#)).

Prerequisites

Before assigning permissions to user groups, you should learn about FlexusRDS system-defined policies and select the policies as required. For the system-defined permissions supported by FlexusRDS, see [FlexusRDS System-defined Policies](#). For the system policies of other services, see [System-defined Permissions](#).

Process Flow

Figure 2-1 Process for granting FlexusRDS 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**.

1. **Create a 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.

2. **Log in** and verify permissions.

Log in to the console by using the created user, and verify the user permissions.

- Go to the FlexusRDS console and click **Buy FlexusRDS Instance** in the upper right corner. 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. If a message appears indicating that you have insufficient permissions to access the service, the **RDS ReadOnlyAccess** policy has already taken effect.

2.1.2 FlexusRDS Custom Policies

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

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.

For details, see [Creating a Custom Policy](#). The following contains examples of common FlexusRDS custom policies.

Example Custom Policies

Example: Allowing users to create manual backups

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

2.2 Buying a FlexusRDS for PostgreSQL Instance

Scenarios

This section describes how to purchase a FlexusRDS for PostgreSQL instance on the FlexusRDS console.

FlexusRDS for PostgreSQL only supports the yearly/monthly billing mode. It allows you to tailor your compute resources and storage space to your business needs.

Prerequisites

- You have [created a Huawei ID and enabled Huawei Cloud services](#).
- Your account balance is greater than or equal to \$0 USD.

Procedure

- Step 1** Go to the [FlexusRDS console](#).
- Step 2** If this is your first time to create a DB instance on this console, click **Buy**.
- Step 3** Configure the instance information and click **Buy**.

Figure 2-2 Selecting an instance class




Figure 2-3 Selecting the required duration




Table 2-1 Basic information

Parameter	Description
Region	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.
DB Engine	Select PostgreSQL .

Parameter	Description
Instance Class	An instance class includes vCPUs, memory, storage, and DB instance type. • Storage: The purchased storage space. After a DB instance is purchased, you can configure storage autoscaling. The maximum allowed storage is 4,000 GB. For details, see Enabling Storage Autoscaling for a FlexusRDS for PostgreSQL Instance. • DB Instance Type – Primary/Standby: An HA architecture. 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: A single-node architecture. It is less expensive than a primary/standby DB pair.
Network	A Virtual Private Cloud (VPC) is a virtual network in which your instance is located. A VPC can isolate networks for different workloads. A subnet provides dedicated network resources that are logically isolated from other networks for security purposes. For details about how to create a VPC and subnet, see Creating a VPC and Subnet.
DB Instance Name	Must start with a letter and consist of 4 to 64 characters. Only letters (case-sensitive), digits, hyphens (-), underscores (_), and periods (.) are allowed. If you buy 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 <i>instance-0001</i>, <i>instance-0002</i>, and so on. If existing instances' suffixes have already reached up to <i>0010</i>, the new instance names will start from <i>instance-0011</i>.
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.
Auto-renew	• This option is not selected by default. • If you select this option, the auto-renew cycle is determined by the selected required duration.
Quantity	You can buy a maximum of 50 DB instances at a time. If you intend to create primary/standby DB instances and set Quantity to 1, a primary instance and a synchronous standby instance will be created.

Step 4 Confirm the order.

Figure 2-4 Order confirmation

Resource	Configuration	Requ...	Quantity	Price
FlexusRDS	Region: CN-Hong Kong DB Instance Name: flexusrds-c161 DB Engine: PostgreSQL DB Engine Version: 16 DB Instance Type: Single Instance Specifications: 2U4G Storage: 120 GB	1 year	1	

Agreement ☒ I have read and agree to the terms and conditions of the [Service Level Agreement \(SLA\)](#).

Price: [Redacted]
You will be charged based on the bill.

[Previous](#) [Pay Now](#)

- If you need to modify your settings, click **Previous**.
- If you do not need to modify your settings, click **Pay Now**.

Step 5 Select a payment method and complete the payment.

Step 6 To view and manage your instance, go to the instance list page.

- When your instance is being created, the status is **Creating**. The status changes to **Available** after the instance is created.
- Automated backup is enabled by default during instance creation and cannot be disabled. An automated full backup is immediately triggered once your DB instance is created.
- The default administrator account of your DB instance is **root**.
- During instance creation, the system randomly sets a password for the administrator account. You need to reset the password before you can connect to the instance. For details, see [Resetting the Administrator Password of a FlexusRDS for PostgreSQL Instance](#).
- The default database port is **5432** and cannot be changed.

----End

2.3 Connecting to a FlexusRDS for PostgreSQL Instance

2.3.1 Using DAS to Connect to a FlexusRDS for PostgreSQL Instance (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 remotely 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** In the instance list, locate the target DB instance and click **Log In** in the **Operation** column.

Figure 2-5 Logging in to an instance



Alternatively, click the instance name in the instance list. On the displayed page, click **Log In** in the upper right corner.

- Step 2** On the displayed login page, enter the username and password and click **Log In**.
----End

2.3.2 Using CLI to Connect to a FlexusRDS for PostgreSQL Instance

Scenarios

You can connect to your DB instance using the psql command-line interface (CLI) from a FlexusX instance with the PostgreSQL client installed.

Connecting to a FlexusRDS for PostgreSQL Instance Through a Client (Version 15 and Earlier Versions)

1. **Log in to the FlexusX instance in the same region as your FlexusRDS DB instance.**
2. Install the PostgreSQL client.

The PostgreSQL community provides **client installation methods** for different OSs. You can download and install the client using the installation tool of the OS. This installation method is simple but has requirements on the ECS OS. It is only available to the OSs supported by the PostgreSQL community.

In this example, Linux CentOS 7 is used. Use the default installation tool of the OS to install the client (version 15 or earlier).

Figure 2-6 Obtaining the installation tool

PostgreSQL Yum Repository

The PostgreSQL Yum Repository will integrate with your normal systems and patch management, and provide automatic updates for all supported versions of PostgreSQL throughout the support lifetime of PostgreSQL.

The PostgreSQL Yum Repository currently supports:

- Red Hat Enterprise Linux
- Rocky Linux
- AlmaLinux
- CentOS (7 and 6 only)
- Oracle Linux
- Fedora*

*Note: due to the shorter support cycle on Fedora, all supported versions of PostgreSQL are not available on this platform. We do not recommend using Fedora for server deployments.

To use the PostgreSQL Yum Repository, follow these steps:

1. Select version:

15

2. Select platform:

Red Hat Enterprise, CentOS, Scientific or Oracle version 7

3. Select architecture:

x86_64

4. Copy, paste and run the relevant parts of the setup script:

```
# Install the repository RPM:
sudo yum install -y https://download.postgresql.org/pub/repos/yum/reporepms/EL-7-x86_64/pgdg-redhat-repo-latest.noarch.rpm

# Install PostgreSQL:
sudo yum install -y postgresql15-server

# Optionally initialise the database and enable automatic start:
sudo /usr/pgdg-15/bin/postgresql-15-setup initdb
sudo systemctl enable postgresql-15
sudo systemctl start postgresql-15
```

Copy Script

Run the **installation** commands:

```
sudo yum install -y https://download.postgresql.org/pub/repos/yum/reporepms/EL-7-x86_64/pgdg-redhat-repo-latest.noarch.rpm
sudo yum install -y postgresql15-server
```

Check whether the installation is successful.

```
psql -V
```

Figure 2-7 Successful installation

```
Running transaction
Installing : postgresql15-libs-15.8-1PGDG.rhel7.x86_64 1/5
Installing : libstd-1.5.5-1.el7.x86_64 2/5
Installing : libicu-50.2-4.el7_7.x86_64 3/5
Installing : postgresql15-15.8-1PGDG.rhel7.x86_64 4/5
Installing : postgresql15-server-15.8-1PGDG.rhel7.x86_64 5/5
Verifying : libicu-50.2-4.el7_7.x86_64 1/5
Verifying : postgresql15-server-15.8-1PGDG.rhel7.x86_64 2/5
Verifying : libstd-1.5.5-1.el7.x86_64 3/5
Verifying : postgresql15-libs-15.8-1PGDG.rhel7.x86_64 4/5
Verifying : postgresql15-15.8-1PGDG.rhel7.x86_64 5/5

Installed:
  postgresql15-server.x86_64 0:15.8-1PGDG.rhel7

Dependency Installed:
  libicu.x86_64 0:50.2-4.el7_7      libstd.x86_64 0:1.5.5-1.el7      postgresql15.x86_64 0:15.8-1PGDG.rhel7      postgresql15-libs.x86_64 0:15.8-1PGDG.rhel7

Complete!
[root@ecs-4dc2 ~]# psql -V
psql (PostgreSQL) 15.8
[root@ecs-4dc2 ~]#
```

3. Run the following command on the FlexusX instance to connect to the FlexusRDS for PostgreSQL DB instance:

```
psql --no-readline -U<user>-h<host>-p<port>-d<datastore>-W
```

Example:

```
psql --no-readline -U root -h 192.168.0.44 -p 5432 -d postgres -W
```

Table 2-2 Parameter description

Parameter	Description
<host>	- Private network connection (recommended): Click the DB instance name and obtain the floating IP address on the Overview page. - Public network connection: Click the DB instance name and obtain the EIP on the Overview page. For details about how to bind an EIP to a DB instance, see Binding an EIP to a FlexusRDS for PostgreSQL Instance or Unbinding an EIP from a FlexusRDS for PostgreSQL Instance. To connect to your DB instance through an EIP, add the EIP and port 5432 to an inbound rule of security group sg-default-smb. For details, see Adding a Security Group Rule. NOTE If the FlexusX instance and the FlexusRDS for PostgreSQL DB instance are in the same region and VPC, you are advised to connect to the DB instance through a private network for higher security. If a DB instance cannot be accessed through a private network, connect to it through a public network.
<port>	5432
<user>	root
<datastore>	postgres

4. When the following information is displayed, enter the password of user **root**:
Enter password:

Connecting to a FlexusRDS for PostgreSQL Instance Using Source Code (No Restrictions on the Version)

1. [Log in to the FlexusX instance in the same region as your FlexusRDS DB instance](#).
2. Install the PostgreSQL client.

Installation from source code: This installation method has no restrictions on the RDS for PostgreSQL instance version and ECS OS.

The following uses an ECS using the Huawei Cloud EulerOS 2.0 image as an example to describe how to install a PostgreSQL 16.4 client.

Figure 2-8 Checking the ECS image

ECS Information	
ID	94ab74c6-2e8a-4a7a-b4f0-abfeb13e7581
Name	ecs-6dc2 
Description	-- 
Region	
AZ	AZ7
Specifications	General computing-plus 2 vCPUs 4 GiB c7.large.2
Image	Huawei Cloud EulerOS 2.0 Standard 64 bit Public image
VPC	default_vpc

- To use SSL connection, download OpenSSL on the ECS in advance.
`sudo yum install -y openssl-devel`
- Obtain the [code download link](#), run **wget** to download the installation package to the ECS, or [download the installation package](#) to the local PC and then upload it to the ECS.
`wget https://ftp.postgresql.org/pub/source/v16.4/postgresql-16.4.tar.gz`
- Decompress the installation package.
`tar xf postgresql-16.4.tar.gz`
- Compile and install the client.
`cd postgresql-16.4`
`./configure --without-icu --without-readline --without-zlib --with-openssl`
`make -j 8 && make install`

 NOTE

If **--prefix** is not specified, the default path is **/usr/local/pgsql**. The client can be installed in the simplest way.

Figure 2-9 Compilation and installation

```
make[4]: Leaving directory '/root/postgresql-16.4/src/port'
make -C ../../src/common all
make[4]: Entering directory '/root/postgresql-16.4/src/common'
make[4]: Nothing to be done for 'all'.
make[4]: Leaving directory '/root/postgresql-16.4/src/common'
make[3]: Leaving directory '/root/postgresql-16.4/src/interfaces/libpq'
make -C ../../src/port all
make[3]: Entering directory '/root/postgresql-16.4/src/port'
make[3]: Nothing to be done for 'all'.
make[3]: Leaving directory '/root/postgresql-16.4/src/port'
make -C ../../src/common all
make[3]: Entering directory '/root/postgresql-16.4/src/common'
make[3]: Nothing to be done for 'all'.
make[3]: Leaving directory '/root/postgresql-16.4/src/common'
/usr/bin/mkdir -p '/usr/local/pgsql/lib/pgxs/src/test/isolation'
/usr/bin/install -c pg_isolation_regress '/usr/local/pgsql/lib/pgxs/src/test/isolation/pg_isolation_regress'
/usr/bin/install -c isolationtester '/usr/local/pgsql/lib/pgxs/src/test/isolation/isolationtester'
make[2]: Leaving directory '/root/postgresql-16.4/src/test/isolation'
make -C test/perl install
make[2]: Entering directory '/root/postgresql-16.4/src/test/perl'
make[2]: Nothing to be done for 'install'.
make[2]: Leaving directory '/root/postgresql-16.4/src/test/perl'
/usr/bin/mkdir -p '/usr/local/pgsql/lib/pgxs/src'
/usr/bin/install -c -m 644 Makefile.global '/usr/local/pgsql/lib/pgxs/src/Makefile.global'
/usr/bin/install -c -m 644 Makefile.port '/usr/local/pgsql/lib/pgxs/src/Makefile.port'
/usr/bin/install -c -m 644 ./Makefile.shlib '/usr/local/pgsql/lib/pgxs/src/Makefile.shlib'
/usr/bin/install -c -m 644 ./nls-global.mk '/usr/local/pgsql/lib/pgxs/src/nls-global.mk'
make[1]: Leaving directory '/root/postgresql-16.4/src'
make -C config install
make[1]: Entering directory '/root/postgresql-16.4/config'
/usr/bin/mkdir -p '/usr/local/pgsql/lib/pgxs/config'
/usr/bin/install -c -m 755 ./install-sh '/usr/local/pgsql/lib/pgxs/config/install-sh'
/usr/bin/install -c -m 755 ./missing '/usr/local/pgsql/lib/pgxs/config/missing'
make[1]: Leaving directory '/root/postgresql-16.4/config'
```

- e. Add the following code to the **/etc/profile** file to configure environment variables:
- ```
export PATH=/usr/local/pgsql/bin:$PATH
export LD_LIBRARY_PATH=/usr/local/pgsql/lib:$LD_LIBRARY_PATH
source /etc/profile
```
- f. Test whether the psql is available.
- ```
psql -V
```

Figure 2-10 Testing psql

```
. /etc/bashrc
fi
export PATH=/usr/local/pgsql/bin:$PATH
export LD_LIBRARY_PATH=/usr/local/pgsql/lib:$LD_LIBRARY_PATH
[root@ecs-88a7 postgresql]# source /etc/profile
[root@ecs-88a7 postgresql]# psql -V
psql (PostgreSQL) 16.4
[root@ecs-88a7 postgresql]#
```

3. Run the following command on the FlexusX instance to connect to the FlexusRDS for PostgreSQL DB instance:

psql --no-readline -U<user>-h<host>-p<port>-d<datastore>-W

Example:

psql --no-readline -U root -h 192.168.0.44 -p 5432 -d postgres -W

Table 2-3 Parameter description

Parameter	Description
<host>	- Private network connection (recommended): Click the DB instance name and obtain the floating IP address on the Overview page. - Public network connection: Click the DB instance name and obtain the EIP on the Overview page. For details about how to bind an EIP to a DB instance, see Binding an EIP to a FlexusRDS for PostgreSQL Instance or Unbinding an EIP from a FlexusRDS for PostgreSQL Instance. To connect to your DB instance through an EIP, add the EIP and port 5432 to an inbound rule of security group sg-default-smb. For details, see Adding a Security Group Rule. NOTE If the FlexusX instance and the FlexusRDS for PostgreSQL DB instance are in the same region and VPC, you are advised to connect to the DB instance through a private network for higher security. If a DB instance cannot be accessed through a private network, connect to it through a public network.
<port>	5432
<user>	root
<datastore>	postgres

4. When the following information is displayed, enter the password of user **root**:
- Enter password:

2.4 Usage Rules

2.4.1 Instance Usage Suggestions

Database Connection

PostgreSQL uses a process architecture, providing a backend service process for each client connection.

- Set **max_connections** depending on the type of your application. Use the parameter settings provided on pg_tune as examples:
 - Set **max_connections** to **200** for web applications.
 - Set **max_connections** to **300** for OLTP applications.
 - Set **max_connections** to **40** for data warehouses.
 - Set **max_connections** to **20** for desktop applications.
 - Set **max_connections** to **100** for hybrid applications.

- Limit the maximum number of connections allowed for a single user based on workload requirements.

```
ALTER ROLE xxx CONNECTION LIMIT xxx;
```
- Set the number of active connections to two to three times the number of vCPUs.
- Avoid long transactions, which may block autovacuum and affect database performance.
- Periodically release persistent connections because maintaining them may generate a large cache and use up memory. You can configure parameters such as **idle_session_timeout** and **idle_in_transaction_session_timeout** to release idle connections.
- Check the application framework to prevent the application from automatically starting transactions without performing any operations.

Reliability and Availability

- Select primary/standby DB instances for production databases.
- Keep vCPU, memory, and storage usage less than 85% for production databases to prevent problems such as out of memory (OOM) and full storage.
- Deploy primary and standby instances in different AZs to improve availability.
- Set the time window for automated backup to off-peak hours. Do not disable full backup.
- Configure asynchronous replication between primary and standby DB instances to prevent workloads on the primary instance from being blocked due to a fault on the standby instance.
- Pay attention to the size of temporary files and the generation rate. Too many temporary files affect database performance, slow down database startup, and even cause service unavailability.
- Do not create too many objects in one instance. Generally, the number of tables in a single instance does not exceed 20,000, and that in a single database does not exceed 4,000. This prevents service unavailability caused by long-time table file scanning during database startup.

Logical Replication

- Keep the name of a logical replication slot less than 40 bytes to prevent full backup failures.
- Delete replication slots that are no longer used for logical replication to prevent database bloat.
- Replication slots will be lost after a primary/standby switchover is performed (due to an instance class change, a minor version upgrade, or a host failure). When this occurs, you need to create replication slots again.
- Use failover slots for PostgreSQL 12.6 and later minor versions, and all minor versions of PostgreSQL 13 and 14 to prevent replication slot loss after a primary/standby switchover or instance reboot.
- When using logical replication, avoid long transactions and commit discarded two-phase transactions in a timely manner to prevent stacked WAL logs from occupying too much storage space.

- When using logical replication, avoid using a large number of sub-transactions (such as savepoints and exception clauses in a transaction) to prevent high memory usage.
- When using Data Replication Service (DRS) to synchronize or migrate data, delete the logical replication slots contained in the databases that are rarely accessed or add heartbeat tables to periodically push the replication slots to prevent stacked WAL logs.

Database Age

- Definition of database age:
 - Database age is a PostgreSQL-specific concept. It refers to the latest transaction ID minus oldest transaction ID in the database.
 - As defined in the Multi-Version Concurrency Control (MVCC) mechanism of PostgreSQL, the maximum age allowed for a database is 2 billion transactions old. When a database reaches the maximum age, it will be forcibly shut down. In this case, contact technical support to vacuum the database.
 - To view the age of a database, run the following SQL statement:
select datname, age(datfrozenxid) from pg_database;
- You are advised to use the **db_max_age** metric to monitor the database age and set the alarm threshold to 1 billion on the Cloud Eye.

Stability

- Commit or roll back two-phase transactions in a timely manner to prevent database bloat.
- Change the table structure, for example, adding fields or indexes, during off-peak hours.
- To create indexes during peak hours, use the CONCURRENTLY syntax to avoid blocking the DML of the table.
- Before modifying the structure of a table during peak hours, perform a verification test to prevent the table from being rewritten.
- Configure a lock wait timeout duration for DDL operations to avoid blocking operations on related tables.
- Partition your database if its capacity exceeds 2 TB.
- If a frequently accessed table contains more than 20 million records or its size exceeds 10 GB, split the table or create partitions.
- To prevent replication exceptions on the standby instance or read replicas, control the data write speed of the primary instance under 50 MB/s. That's because the standby instance or read replica replays WAL logs in a single process at a maximum speed of 50 MB/s to 70 MB/s.

Routine O&M

- Periodically download and view slow query logs on the **Logs** page to identify and resolve performance issues in a timely manner.
- Periodically check the resource usage of your instance. If the service pressure fluctuates greatly, you are advised to configure resource alarms and upgrade

the instance specifications when necessary. High write pressure will slow down database reboots and affect service availability.

- Run the **SELECT** statement before deleting or modifying a record.
- After a large amount of data is deleted or updated in a table, run **VACUUM** on the table.
- Note the number of available replication slots and ensure that at least one replication slot is available for database backup.
- Remove any replication slots that are no longer used to prevent the replication slots from blocking log reclaiming.
- Do not use unlogged tables because data in these tables will be lost after a database exception (such as OOM or underlying faults) or primary/standby switchover.
- Do not run **VACUUM FULL** on system catalogs. If necessary, run **VACUUM**. Running **VACUUM FULL** on system catalogs causes the instance to reboot and the instance cannot be connected for a long time.

Security

- Avoid enabling access to your database from the Internet. If you do need to enable Internet access, bind an EIP to your DB instance and configure a whitelist.
- Use SSL to connect to your DB instance.

2.4.2 Database Usage Rules

Naming

- The names of objects (such as databases, tables, and indexes) must be no more than 63 bytes. Note that some characters (such as Chinese characters) may occupy multiple bytes.
- Do not use reserved database keywords in object names or start an object name with **pg** or a digit.
- A database name can contain 1 to 63 characters. Only letters, digits, and underscores (**_**) are allowed. It cannot start with **pg** or a digit and cannot be the same as PostgreSQL template database names. PostgreSQL template databases include **postgres**, **template0**, and **template1**.

Table Design

- The table structure must be designed in advance to avoid frequent structure changes, such as adding fields or changing data types.
- There cannot be more than 64 fields in a single table.
- Create partitioned tables for the tables whose data needs to be deleted periodically. For example, you can create partitions by time and delete data from the partitions using **DROP** or **TRUNCATE**.
- Use appropriate data types for table fields. For example, do not use the character type for numeric or date data.
- When using the numeric data type, ensure that the values are within allowed ranges and meet precision requirements.

Index Design

- Design primary keys or unique keys for tables that require logical replication.
- When creating a foreign key, specify the action for deleting or updating the foreign key, for example, ON DELETE CASCADE.
- Create indexes for fields that are frequently used (such as fields for data query and sorting).
- Create partial indexes for queries using fixed conditions.
- Create expression indexes for queries using conditional expressions.
- A single table cannot contain too many indexes because indexes also occupy storage. For example, there should be fewer than 5 single-column indexes and fewer than 3 composite indexes.

SQL Design

- Specify the required fields to be returned in a query.
- Only use IS NULL or IS NOT NULL to determine whether a field is NULL.
- Use NOT EXISTS instead of NOT IN in a query.
- Use UNION ALL instead of UNION to concatenate result sets.
- Use TRUNCATE instead of DELETE to delete an entire table.
- Submit data changes in large transactions in batches to prevent high pressure during transaction commit or rollback.
- When creating a function, define the volatility of the function as the strictest category, instead of the default VOLATILE. Too many concurrent calls of VOLATILE functions may result in failures to establish new connections.

Security

- Do not assign the public role to the owner of an application database object. Assign a specific role to the owner.
- A database password must meet complexity requirements.
- Allocate a unique database account for each service.
- When accessing an object, explicitly specify the schema of the object to avoid accessing objects with the same name in other schemas.

2.5 Data Migration

2.5.1 Migrating Databases to FlexusRDS for PostgreSQL Using DRS

You can migrate data from RDS for PostgreSQL, self-managed PostgreSQL databases, PostgreSQL databases built on other clouds, self-managed Oracle databases, RDS for MySQL, self-managed MySQL databases, or MySQL databases built on other clouds to FlexusRDS for PostgreSQL, or from one FlexusRDS for PostgreSQL instance to another FlexusRDS for PostgreSQL instance.

You are advised to use Data Replication Service (DRS) to migrate data because it is easy to use and can complete a migration task in minutes. It facilitates data

transfer between databases, helping you reduce DBA labor costs and hardware costs.

DRS Migration Solutions

The method for migrating databases to FlexusRDS for PostgreSQL using DRS is the same as that for migrating databases to RDS for PostgreSQL. For details, see [Migration Solution Overview](#).

DRS Migration Billing

- Real-time migration supports only the pay-per-use billing mode.
Real-time migration tasks are free of configuration and data transmission fees in the first seven days, lowering your costs for migrating data to the cloud.
- Real-time synchronization and DR support pay-per-use and yearly/monthly billing modes.
Real-time migration and synchronization will provide long-term discounts, lowering your costs for data transfers.

For more information, see [Data Replication Service Billing](#).

2.5.2 Migrating Data to FlexusRDS for PostgreSQL Using psql

Preparing for Data Migration

You can access your FlexusRDS for PostgreSQL DB instance through an EIP or from a FlexusX instance.

1. Prepare a FlexusX instance for accessing your FlexusRDS for PostgreSQL DB instance or prepare a device for accessing your FlexusRDS for PostgreSQL DB instance through an EIP.
To connect to a FlexusRDS for PostgreSQL instance through an EIP, [bind an EIP](#) to the instance.
2. Install a PostgreSQL client of the same version as your FlexusRDS for PostgreSQL instance on the prepared FlexusX instance or device.

NOTE

A PostgreSQL database or client comes with pg_dump and psql.

Exporting Data

Before migrating an existing PostgreSQL database to FlexusRDS for PostgreSQL, you need to export data first.

NOTICE

- The export tool must match the DB engine version.
 - Database migration is performed offline. Before the migration, you must stop any applications using the source database.
-

Step 1 Log in to the FlexusX instance or device that can access the FlexusRDS for PostgreSQL instance.

Step 2 Use the `pg_dump` tool to export the source database into an SQL file.

```
pg_dump--username=<DB_USER> --host=<DB_ADDRESS> --port=<DB_PORT> --  
format=plain --file=<BACKUP_FILE><DB_NAME>
```

- *DB_USER* indicates the database username.
- *DB_ADDRESS* indicates the database address.
- *DB_PORT* indicates the database port.
- *BACKUP_FILE* indicates the name of the file to which the data will be exported.
- *DB_NAME* indicates the name of the database to be migrated.

Enter the database password as prompted.

NOTE

If the exported SQL file uses INSERT statements, you can easily edit and modify the file. However, the speed of importing data may be slower than that of using COPY statements. You are advised to select a right statement format as needed.

- If both the source and destination databases are PostgreSQL databases, you are advised to export COPY statements (default). For details, see [Example 1: Exporting the source database to an SQL file \(COPY\)](#).
- If either of the source and destination databases is a non-PostgreSQL database, you are advised to export INSERT statements. For details, see [Example 2: Exporting the source database to an SQL file \(INSERT\)](#).

For more information, see [pg_dump options](#).

Examples:

- Example 1: Exporting the source database to an SQL file (COPY)
**\$ pg_dump --username=root --host=192.168.151.18 --port=5432 --
format=plain --file=backup.sql my_db**
Password for user root:
- Example 2: Exporting the source database to an SQL file (INSERT)
**\$ pg_dump --username=root --host=192.168.151.18 --port=5432 --
format=plain --inserts --file=backup.sql my_db**
Password for user root:
- Example 3: Exporting all table structures from the source database to an SQL file
**\$ pg_dump --username=root --host=192.168.151.18 --port=5432 --
format=plain --schema-only --file=backup.sql my_db**
Password for user root:
- Example 4: Exporting all table data from the source database to an SQL file
**\$ pg_dump --username=root --host=192.168.151.18 --port=5432 --
format=plain --data-only --file=backup.sql my_db**
Password for user root:

After the commands in any of the above examples are executed, a **backup.sql** file will be generated as follows:

```
[rds@localhost ~]$ ll backup.sql
-rw-r-----. 1 rds rds 2714 Sep 21 08:23 backup.sql
```

Step 3 Use `pg_dump` to export tables from the source database to an SQL file.

`pg_dump --username=<DB_USER> --host=<DB_ADDRESS> --port=<DB_PORT> --format=plain --file=<BACKUP_FILE> <DB_NAME> --table=<TABLE_NAME>`

- *DB_USER* indicates the database username.
- *DB_ADDRESS* indicates the database address.
- *DB_PORT* indicates the database port.
- *BACKUP_FILE* indicates the name of the file to which the data will be exported.
- *DB_NAME* indicates the name of the database to be migrated.
- *TABLE_NAME* indicates the name of the specified table in the database to be migrated.

Enter the database password as prompted.

Examples:

- Example 1: Exporting one table from the source database to an SQL file
`$ pg_dump --username=root --host=192.168.151.18 --port=5432 --format=plain --file=backup.sql my_db --table=test`
Password for user root:
- Example 2: Exporting multiple tables from the source database to an SQL file
`$ pg_dump --username=root --host=192.168.151.18 --port=5432 --format=plain --file=backup.sql my_db --table=test1 --table=test2`
Password for user root:
- Example 3: Exporting all tables starting with `ts_` from the source database to an SQL file
`$ pg_dump --username=root --host=192.168.151.18 --port=5432 --format=plain --file=backup.sql my_db --table=ts_*`
Password for user root:
- Example 4: Exporting all tables except those starting with `ts_` from the source database to an SQL file
`$ pg_dump --username=root --host=192.168.151.18 --port=5432 --format=plain --file=backup.sql my_db -T=ts_*`
Password for user root:

After the commands in any of the above examples are executed, a **backup.sql** file will be generated as follows:

```
[rds@localhost ~]$ ll backup.sql
-rw-r-----. 1 rds rds 2714 Sep 21 08:23 backup.sql
```

----End

Importing Data

Step 1 Log in to the FlexusX instance or device that can access the FlexusRDS for PostgreSQL instance.

Step 2 Ensure that the destination database to which data is to be imported exists.

If the destination database does not exist, run the following command to create a database:

```
# psql --host=<DB_ADDRESS>--port=<DB_PORT>--username=root--  
dbname=postgres-c "create database<DB_NAME>";
```

- *DB_ADDRESS* indicates the IP address of the DB instance.
- *DB_PORT* indicates the database port of the DB instance.
- *DB_NAME* indicates the name of the database to which data is to be imported.

Step 3 Import the exported file to FlexusRDS for PostgreSQL.

```
# psql --host=<DB_ADDRESS> --port=<DB_PORT>--username=root--  
dbname=<DB_NAME>--file=<BACKUP_DIR>/backup.sql
```

- *DB_ADDRESS* indicates the IP address of the DB instance.
- *DB_PORT* indicates the database port of the DB instance.
- *DB_NAME* indicates the name of the database to which data is to be imported. Ensure that the database exists.
- *BACKUP_DIR* indicates the directory where the **backup.sql** file is stored.

Enter the password for the DB instance when prompted.

Example:

```
# psql --host=172.16.66.198 --port=5432 --username=root --dbname=my_db --  
file=backup.sql
```

Password for user root:

Step 4 View the import result.

```
my_db=> \l my_db
```

In this example, the database named **my_db** has been imported.

```
my_db=> \l my_db  
List of databases  
Name | Owner | Encoding | Collate | Ctype | Access privileges  
-----+-----+-----+-----+-----+-----  
my_db | root | UTF8 | en_US.UTF-8 | en_US.UTF-8 |  
(1 row)
```

----End

2.5.3 Migrating Data to FlexusRDS for PostgreSQL Using the Export and Import Functions of DAS

Scenarios

To back up or migrate data, you can use Data Admin Service (DAS) to export data from the source database first and then import the data to from your local PC or OBS bucket to the destination database.

For more information, see [Data Import and Export](#).

Constraints

- Only one file that is no larger than 1 GB can be imported at a time.
- Binary fields such as BINARY, VARBINARY, TINYBLOB, BLOB, MEDIUMBLOB, and LONGBLOB are not supported.
- If there are more than 100,000 tables in a FlexusRDS for PostgreSQL instance, an error will be reported when you export data using the **Export Database** function of DAS. In this case, use the **Export SQL Result** function instead.

Exporting Data

Step 1 On the **Instances** page, locate the DB instance and click **Log In** in the **Operation** column.

Step 2 On the displayed login page, enter the username and password and click **Log In**.

Step 3 On the top menu bar, choose **Import and Export** > **Export**.

Step 4 On the displayed page, click **Create Task** and choose **Export Database** or **Export SQL Result** as required. The following takes database export as an example.

Alternatively, click **Quick Export** and select the target database. On the displayed page, select a storage path and click **OK**.

Figure 2-11 Quick export

Quick Export

Export Database

Storage ?

No OBS bucket? [Create OBS Bucket](#)

Creating an OBS bucket is free of charge, but storing files in it will incur fees.

OK Cancel

Step 5 On the displayed page, set parameters as required in areas **Basic Information** and **Advanced Settings**. Then, select the tables to be exported on the right.

Figure 2-12 Creating an export task

Export Database

Basic Information

Database

create_new_db1

☐ Export all tables

Allowed Rows

10,000

File Type

SQL

CSV

Object to Export

Data

Structure

Data and structure

Charset

UTF8

GBK

Storage

abcsdfdf

No OBS bucket? [Create OBS Bucket](#)

Options

☒ Combine INSERT statements. (Combine INSERT statements into files, with each file smaller than 5 MB.)

☐ Generate a file for each table. (Downloading table files in the details slows down the export.)

Remarks

Advanced Settings

Tables

Selected Tables: 0

Enter a table name.

	Table Name	Column	WHERE Clause
☐	asd	Edit	Edit
☐	dddd	Edit	Edit
☐	new_db1_tb1	Edit	Edit
☐	rule_aml_main	Edit	Edit
☐	test1	Edit	Edit

10 / page

Total Records: 5

< 1 >

NOTE

- In a SQL result export task, the executed SQL statements cannot exceed 5 MB.

Figure 2-13 Export SQL Result

Export SQL Result

Basic Information

Database

create_new_db1

Allowed Rows

10,000

File Type

SQL-Insert

CSV

Charset

UTF8

GBK

Storage

abcsdfdf

No OBS bucket? [Create Bucket](#)

Options

☒ Combine INSERT statements. (Combine INSERT statements into files, with each file smaller than 5 MB.)

☐ Generate one file for each result.

SQL to Execute

SELECT * FROM "asd"

Remarks

Advanced Settings

OK

Cancel

- Databases are classified into user databases and system databases. System databases cannot be exported. If system database data is required, deploy system database services in a created user database, so that you can export the system database data from the user database.
- DAS connects to your standby database to export data. This prevents the primary database from being affected by data export. However, if the standby database has a high replication delay, the exported data may not be the latest.

Step 6 After settings are complete, click **OK**.

Step 7 In the task list, view the task ID, type, status, and progress.

Step 8 Click **Details** in the **Operation** column to view task details.

Figure 2-14 Task list

Create Task

Quick Export

Create Task

Time Range:

Start Date

End Date

Search

Task ID	Task Type	Database	Started	Ended	File Size	File Type	Status	Elapsed Time	Exports # Rows	Progress	Remarks	Operation
c488c596b0843276b5386a08932251	Quick E	db_01	2020-09-07 20:18:45	2020-09-07 20:16:55	4.53 MB	SQL	successful	10 secs	202415	100%		Details Download
c20b47326d747e1947b147752a1a17e1c2	Database	create_new_db1	2020-09-03 18:50:45	2020-09-03 18:52:14	16.36 MB	SQL	successful	1 mins	19090	100%		Details Download
7a970a2c9e0d48e409f6962e18e08d86ca	Database	create_new_db1	2020-09-03 16:47:05	2020-09-03 16:47:22	3.04 MB	SQL	successful	17 secs	2414	100%		Details Download

---End

Importing Data

Step 1 On the top menu bar, choose **Import and Export > Import**.

Step 2 Import a file from your local PC or an OBS bucket.

Figure 2-15 Creating an import task

Create Task

X

Import Type

sql

CSV

File Source

Upload file

Choose from OBS

Attachment Storage ?

407154

▼

No OBS bucket? [Create OBS Bucket](#)

Creating an OBS bucket is free of charge, but storing files in it will incur fees.

Attachment

+

Click here to upload a file, or drag one here. (.sql)

Upload only one attachment that is no larger than 1 GB.

Database

db_4eb3_0000

▼

Charset

Auto Detect

UTF8

GBK

Options

☐ Ignore errors, that is, skip the step where the SQL statement fails to be executed.

☒ Delete the uploaded file upon an import success.

Remarks

Create

Cancel

- From your local PC

In the upper left corner, click **Create Task**. On the displayed page, select an import type, select **Upload file** for **File Source**, set the attachment storage, and upload the file. Then, set other parameters as required.

For security purposes, imported files are stored in OBS buckets.

 NOTE

- To keep your data secure, provide your own OBS bucket to store the attachments you upload. In this way, DAS automatically connects to your OBS bucket for in-memory reading. No data is stored on DAS.
- If you select **Delete the uploaded file upon an import success.**, the file you uploaded will be automatically deleted from the OBS bucket after being imported to the destination database.
- From an OBS bucket
In the upper left corner, click **Create Task**. On the displayed page, select an import type, select **Choose from OBS** for **File Source**, and select a file from the bucket. Then, set other parameters as required.

 NOTE

The file uploaded from an OBS bucket will not be deleted upon an import success.

Step 3 After setting the import parameters, click **Create**. Confirm the information again before you click **OK** because original data may be overwritten after data import.

Step 4 View the import progress in the task list or check task details.

----End

2.6 Instance Management

2.6.1 Changing the Name of a FlexusRDS for PostgreSQL Instance

Scenarios

You can change the name of a DB instance as required.

Procedure

Step 1 In the instance list, locate the instance that you want to edit name for and click  next to the instance name. Then, change the name and click **OK**.

Alternatively, click the target instance name. On the displayed page, click  under the **DB Instance Name** field and change the instance name.

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

- To submit the change, click .

- To cancel the change, click **✖**.

Step 2 Check the result in the instance list.

----End

2.6.2 Rebooting FlexusRDS for PostgreSQL Instances

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 one or more DB instances at a time on the console.

Constraints

- If the DB instance status is **Abnormal**, the reboot may fail.
- Rebooting a DB instance will reboot the DB engine service, causing service interruptions. During this period, the 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.
- After a primary/standby DB instance is rebooted, it takes about one minute to establish the replication relationship. During this period, some operations, such as changing the instance class, cannot be performed.

Procedure

Step 1 In the instance list, select one or more DB instances (maximum: 50) to be rebooted and click **Reboot** above the instance list.

Step 2 In the displayed dialog box, enter **YES** and click **OK**.

Figure 2-16 Rebooting a DB instance**Reboot DB Instance**

-  • The DB instances will be unavailable when they are being rebooted.
- Rebooting instances will clear their cached memory. To prevent traffic congestion during peak hours, reboot the instances during off-peak hours.
- The DB instances will be unavailable when they are being rebooted. Rebooting instances will clear their cached memory. To prevent traffic congestion during peak hours, reboot the instances during off-peak hours.

You are rebooting instance flexusrds-a734.

DB Instance Name	Status	DB Instance Type	Storage (GB)
flexusrds-a734 156726a437304c6c9c3559...	 Available	Single 2 vCPUs 4 GB	120

To confirm reboot, enter "YES" below. [Auto Enter](#)

CancelOK

Step 3 View the instance status. If the status is **Available**, the instance has been rebooted successfully.

----End

2.6.3 Resetting the Administrator Password of a FlexusRDS for PostgreSQL Instance

Scenarios

You can reset the administrator password only through a primary instance. The new password is applied immediately without rebooting the instance.

If you forget the password of the administrator account **root**, you can reset the password.

Precautions

- If the password you provide is regarded as a weak password by the system, you will be prompted to enter a stronger password.
- If you change the administrator password of a primary instance, the administrator password of the standby instance 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.

Procedure

- Step 1** In the instance list, locate the target instance and click **Reset Password** in the **Operation** column.
- Step 2** In the displayed dialog box, enter a new password and confirm the password.

Figure 2-17 Resetting a password

Reset Password ✕

DB instance ID	68481a20f1234aedb6b34d7b60732613in03
DB Instance Name	flexusrds-53e5
New Password	
Confirm Password	

Cancel OK

NOTICE

Keep this password secure. The system cannot retrieve it.

The 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.

- To submit the new password, click **OK**.
- To cancel the reset operation, click **Cancel**.

----End

2.6.4 Enabling Storage Autoscaling for a FlexusRDS for PostgreSQL Instance

Scenarios

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

Constraints

- If your account balance is insufficient, storage autoscaling will fail.
- The storage space can be autoscaled up only when your instance status is **Available** or **Storage full**.
- For a primary/standby DB instance, autoscaling the storage for the primary node will also autoscale the storage for the standby node.
- If a yearly/monthly DB instance has pending orders, it will not be autoscaled.

Procedure

Step 1 In the instance list, click the target instance name to go to the **Overview** page.

Step 2 On the **Overview** page, click **Configure** under the **Configure Autoscaling** field.

Step 3 In the displayed dialog box, click  and configure the required parameters.

Figure 2-18 Configuring autoscaling

Configure Autoscaling ✕

Enable Autoscaling

☒

After autoscaling is enabled, the storage added to your instance will be billed. Ensure that your account balance is sufficient. [Learn more](#) 

Trigger If Available Storage Drops To

If available storage drops to or below this value, your storage will autoscale.

Increment (%)

Enter an integer.

Autoscaling Limit (GB)

Storage can autoscale to no more than 4,000 GB.

If available storage drops below 10% or 10 GB, your storage will autoscale by 20% (in increments of 10 GB) of your allocated storage.

CancelOK

Table 2-4 Parameter description

Parameter	Description
Enable Autoscaling	If you select this option, autoscaling is enabled.
Trigger If Available Storage Drops To	If the available storage drops to a specified threshold (10% , 15% , or 20%), autoscaling is triggered.
Increment	Autoscaling increment, as a percentage. The default value range is from 5% to 50% .
Autoscaling Limit	The value range is from 120 to 4000 , in GB. The limit must be no less than the storage of the DB instance.

Step 4 Click **OK**.

----End

2.6.5 Binding an EIP to a FlexusRDS for PostgreSQL Instance or Unbinding an EIP from a FlexusRDS for PostgreSQL Instance

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.

Precautions

- You can buy an EIP [on the network console](#) and bind it to a FlexusRDS for PostgreSQL instance. One EIP can be bound to only one instance. For pricing details, see [Elastic IP pricing details](#).
- If a DB instance has already been bound with an EIP, you must unbind the EIP from the instance first before binding a new EIP to it.

Binding an EIP to a DB Instance

Step 1 In the instance list, click the target instance name to go to the **Overview** page.

Step 2 On the displayed **Overview** page, click **Bind** under the **EIP** field.

Step 3 In the displayed dialog box, all available EIPs are listed. Select the required EIP and click **Yes**.

Step 4 View the EIP that has been bound to the DB instance.

----End

Unbinding an EIP from a DB Instance

Step 1 In the instance list, click the target instance name to go to the **Overview** page.

Step 2 On the displayed **Overview** page, click **Unbind** under the **EIP** field.

Step 3 In the displayed dialog box, click **Yes**.

----End

2.6.6 Renewing FlexusRDS for PostgreSQL Instances

Scenarios

You can renew one or multiple yearly/monthly DB instances at a time.


Procedure

Step 1 In the instance list, select the target DB instance and click **Renew** above the instance list.

You can also click the target instance name to go to the **Overview** page to renew the instance.

Step 2 In the displayed dialog box, confirm the instance to be renewed.

Figure 2-19 Renewing an instance



Step 3 Click **OK** to go to the renewal page and renew the instance.

----End

2.6.7 Unsubscribing from a FlexusRDS for PostgreSQL Instance

Scenarios

To delete a DB instance billed on the yearly/monthly basis, you need to unsubscribe the order. For unsubscription fees, see [Unsubscription Rules](#).

Constraints

- A DB instance cannot be unsubscribed when any operations are being performed on it. It can be unsubscribed only after the operations are complete.

- If a backup of a DB instance is being restored, the instance cannot be unsubscribed.

Procedure

Step 1 In the instance list, select the target instance and click **Unsubscribe** above the instance list.

Step 2 In the displayed dialog box, enter **YES**.

Step 3 Click **OK**.

After you unsubscribe from an instance order, the instance will be deleted and it is no longer displayed in the instance list.

----End

2.7 Backups and Restorations

2.7.1 Creating a Manual Backup for a FlexusRDS for PostgreSQL Instance

Scenarios

FlexusRDS for PostgreSQL allows you to create manual backups for an available DB instance. You can use these backups to restore data.

Constraints

- You can create manual backups only when your account balance is no less than \$0 USD.
- The backup name must be unique.
- When you delete a DB instance, its automated backups are also deleted but its manual backups are retained.

Billing

Backups are saved as packages in OBS buckets. For the billing details, see [How Is FlexusRDS Backup Data Billed?](#)

Procedure

Step 1 In the instance list, click the target instance name.

Step 2 Click **Backups & Restorations** and then click **Create Backup**.

Step 3 In the displayed dialog box, enter a backup name and description, and click **OK**.

Figure 2-20 Creating a backup

✕

Create Backup

i

- When the DB instance is being backed up, data is copied and then compressed and uploaded to OBS at an average speed of 60 MB/s.
- Creating a backup increases the disk I/O load. Perform this operation during off-peak hours.

✕

DB instance ID
68481a20f1234aedb6b34d7b60732613in03

DB Instance Name
flexusrds-53e5

Backup Name

backup-6ba2

?

Description?

0/256

Cancel

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: >|<"&'=
- The time required for creating a manual backup depends on the amount of data.

Step 4 View and manage the created backup on the **Backups & Restorations** page.

----End

2.7.2 Deleting a Manual Backup of a FlexusRDS for PostgreSQL Instance

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 In the instance list, click the target instance name.

Step 2 Click **Backups and Restorations**.

Step 3 Locate a manual backup and click **Delete** in the **Operation** column.

The following backups cannot be deleted:

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

Step 4 In the displayed dialog box, click **Yes**.

-----End

2.7.3 Downloading a Full Backup of a FlexusRDS for PostgreSQL Instance

Scenarios

You can download manual and automated full backup files in .qp format for local storage.

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 In the instance list, click the target instance name.

Step 2 Click **Backups and Restorations**.

Step 3 Locate the backup to be downloaded and click **Download** in the **Operation** column.

Step 4 In the displayed dialog box, select **Use OBS Browser+** for **Download Method** and click **OK**.

Figure 2-21 Using OBS Browser+







Figure 2-22 Download guide



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 2-23 Logging in to OBS Browser+


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

Figure 2-24 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 the Resources in It](#).

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 2-25 Downloading a backup




----End

Method 2: Using Current Browser

- Step 1** In the instance list, click the target instance name.
- Step 2** Click **Backups and Restorations**.
- Step 3** Locate the backup to be downloaded and click **Download** in the **Operation** column.
- Step 4** In the displayed dialog box, select **Use Current Browser** for **Download Method** and click **OK**.

Figure 2-26 Using the current browser




----End

Method 3: Using Download URL

- Step 1** In the instance list, click the target instance name.
- Step 2** Click **Backups and Restorations**.
- Step 3** Locate the backup to be downloaded and click **Download** in the **Operation** column.

- Step 4** 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 2-27 Using the download URL



- You can also run the following command to download backup files:

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

The parameters in the command are as follows:

FILE_NAME: indicates 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: indicates the location of the backup file to be downloaded. If the location contains special characters, escape is required.

----End

2.7.4 Checking and Exporting Backup Information of a FlexusRDS for PostgreSQL Instance

Scenarios

You can export backup information of FlexusRDS for PostgreSQL instances to an Excel file for further analysis. The exported information includes the backup ID, backup name, backup type, backup method, backup start and end times, status, size, and description.

For details about how to export backup data, see [Downloading a Full Backup of a FlexusRDS for PostgreSQL Instance](#).

Procedure

Step 1 In the instance list, click the target instance name.

Step 2 Click **Backups and Restorations**.

Step 3 Click  above the backup list to export backup information.

- If you want to export specified backup records, you can first select them and then export them. You can only select and export the backup records displayed on the current page.
- If you do not select any backup records, all backup records are exported by default. (A maximum of the first 5,000 backup records can be exported. If you want to export more, select the records and export them.)
- The backup information is exported to an Excel file for your further analysis.

Figure 2-28 Backup information

	A	B	C	D	E	F	G	H
1	Backup ID	Backup Name	Backup Type	Backup Method	Backup Time	Status	Size	Description
2	6a276aef43ba40bcacbdccfa33063d2br03	backup-5dc6	Manual	Physical backup	Jan 08, 2025 16:18:12 — Jan 08, 2025 16:19:24 GMT+08:00	Completed	3.61 MB	--

----End

2.7.5 Restoring a FlexusRDS for PostgreSQL Instance

2.7.5.1 Restoring a FlexusRDS for PostgreSQL Instance from Backups

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, the backup is downloaded from OBS and then restored to the DB instance at an average speed of 40 MB/s.

Constraints

- Constraints on restoring data to a new DB instance:
 - You can restore data to a new instance only when your account balance is greater than or equal to \$0 USD. You will pay for the new instance specifications.
 - The storage space of the new instance should be no less than that of the original instance.
 - 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 In the instance list, click the target instance name.

Step 2 Click **Backups and Restorations**.

Step 3 Locate the backup to be restored and click **Restore** in the **Operation** column.

Step 4 Select a restoration method and click **OK**.

- Create New Instance

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

- The DB engine version of the new instance is the same as that of the original instance.
- You can select a new instance class for the new instance, but the storage of the new instance must be no less than that of the original instance.
- For details about the configuration items on the purchase page, see [Buying a FlexusRDS for PostgreSQL Instance](#).

 NOTE

To change the instance class, choose a new one for the new instance on the **Create New Instance** page.

Step 5 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 instance status changes from **Creating** to **Available**.

The new DB instance is independent from the original one. After the new instance is created, a full backup will be automatically triggered.

----End

2.7.5.2 Restoring a FlexusRDS for PostgreSQL Instance to a Point in Time

Scenarios

You can use automated backups to restore an instance to a specific point in time.

When you enter the time point that you want to restore the DB instance to, FlexusRDS 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 40 MB/s.

Constraints

- Do not run the **reset master** command on instances within their lifecycle. Otherwise, an exception may occur during the point-in-time recovery (PITR).
- Constraints on restoring data to a new DB instance:
 - You can restore data to a new instance only when your account balance is greater than or equal to \$0 USD. You will pay for the new instance specifications.
 - The storage space of the new instance should be no less than that of the original instance.

- 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, contact customer service.

Procedure

Step 1 In the instance list, click the target instance name.

Step 2 Click **Backups and Restorations**.

Step 3 Click **Restore** above the backup list.

Step 4 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 version of the new instance is the same as that of the original instance.
- You can select a new instance class for the new instance, but the storage of the new instance must be no less than that of the original instance.
- For details about the configuration items on the purchase page, see [Buying a FlexusRDS for PostgreSQL Instance](#).

NOTE

To change the instance class, choose a new one for the new instance on the **Create New Instance** page.

Step 5 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 instance status changes from **Creating** to **Available**.

The new DB instance is independent from the original one. After the new instance is created, a full backup will be automatically triggered.

----End

2.8 Parameters

2.8.1 Suggestions on Parameter Tuning for a FlexusRDS for PostgreSQL Instance

Parameters are key configuration items in a database system. Improper parameter settings may adversely affect database performance. This section describes some important parameters for your reference. For details, visit the [PostgreSQL official website](#).

For details on how to modify PostgreSQL parameters on the console, see [Modifying Parameters of a FlexusRDS for PostgreSQL Instance](#).

Sensitive Parameters

The following parameters can result in system security and stability issues if set improperly:

- The **search_path** parameter must be set to a schema sequence where schemas are separated by commas (.). Ensure that the schemas exist. Otherwise, the database performance will be affected.
- If you enable the parameter **log_duration**, SQL statements containing sensitive information may be recorded in logs. You are advised to disable this parameter.
- **log_min_duration_statement** specifies how many milliseconds a query has to run before it has to be logged. The unit is millisecond. Setting this parameter to **0** means that all statements are recorded. Setting this parameter to **-1** means that no statement is recorded. For details, see [Viewing and Downloading Slow Query Logs of a FlexusRDS for PostgreSQL Instance](#).
- The **temp_file_limit** parameter limits the total size (in KB) of all temporary files when writing temporary files to the disk is triggered in a session. The value ranges from -1 to 2,147,483,647. The value **-1** indicates that the total size of the temporary files is not limited.
 - This parameter is only available for PostgreSQL 12, 13, 14, and 15.
 - To prevent temporary files from occupying too much disk space and causing service exceptions, do not set this parameter to **-1**.
 - If the parameter value is changed to a larger value for temporary use but is not changed to the original value after the use, the disk space will be continuously used to store temporary files. If the disk space is used up, services will be interrupted and the DB instance will become unavailable.
- The **max_pred_locks_per_transaction** and **max_locks_per_transaction** parameters need to be set based on the values of **max_connections** and **max_prepared_transactions**. Too large values may cause instance unavailability.

Performance Parameters

The following parameters can result in system security and stability issues if set improperly:

- If **log_statement** is set to **ddl**, **mod**, or **all**, the operations for creating and deleting database users (including passwords and other sensitive information) are recorded. This operation affects database performance. Exercise caution when setting this parameter.
- Enabling the following parameters will affect the database performance: **log_hostname**, **log_duration**, **log_connections**, and **log_disconnections**. Exercise caution when enabling these parameters.
- The **shared_buffers** parameter is recommended to be a value ranging from 25% to 40% of the system memory. The maximum value of this parameter cannot exceed 80% of the system memory to avoid affecting database performance.
- The **max_worker_processes** parameter should be set based on the values of **max_parallel_workers** and **max_parallel_workers_per_gather**. If the

max_worker_processes value is too large, the database performance will be affected.

2.8.2 Modifying Parameters of a FlexusRDS for PostgreSQL Instance

Scenarios

You can change parameter values in a custom parameter template and apply it to optimize database performance.

Modifying a Single Parameter

- Step 1** In the instance list, click the target instance name.
- Step 2** Click the **Parameters** tab.
- Step 3** In the parameter list, locate the parameter you want to modify and click **Modify** in the **Operation** column.

NOTICE

After you modify a parameter, check the value in the **Effective upon Reboot** column.

- If the value is **Yes** and the instance status in the instance list 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 the value is **No**, the modifications take effect immediately.

Figure 2-29 Parameters

Parameter Name	Effective upon Reboot	Value	Allowed Values	Description	Operation
archive_command	No	/usr/pgsq/bin/rds_wal_archive_%p	--	Sets the shell command that will be called to archive a ...	Modify
archive_timeout	No	300	0-2,147,483,647	Forces a switch to the next xlog file if a new file has not ...	Modify
array_nulls	No	on	on,off	Enable input of NULL elements in arrays.	Modify
auth_delay	No	3000	0-2,147,483	The number of milliseconds to wait before reporting an ...	Modify
authentication_timeout	No	60	1-600	Sets the maximum allowed time to complete client auth...	Modify
auto_explain_log_min_duration	No	10000	-1-60,000	Sets the minimum execution time above which plans will...	Modify
autovacuum	No	on	on,off	Starts the autovacuum subprocess.	Modify

- To save the modifications, click **Confirm**. In the displayed dialog box, click **Yes**.
- To cancel the modifications, click **Cancel**.

To view the change history, click **Change History** above the parameter list. The change history of the last seven days is displayed.

----End

Modifying Parameters in Batches

Step 1 In the instance list, click the target instance name.

Step 2 Click the **Parameters** tab.

Step 3 Switch on the batch modification switch . A maximum of 30 parameters can be modified at a time.

NOTICE

After you modify a parameter, check the value in the **Effective upon Reboot** column.

- If the value is **Yes** and the instance status in the instance list 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 the value is **No**, the modifications take effect immediately.

- To save your modifications, click **Save**. In the displayed dialog box, click **Yes**.
- To cancel your modifications, click **Cancel**. In the displayed dialog box, click **Yes**.

To view the change history, click **Change History** above the parameter list. The change history of the last seven days is displayed.

----End

2.8.3 Exporting the Parameter List of a FlexusRDS for PostgreSQL Instance

Scenarios

You can also export the parameter information (including parameter names, values, and descriptions) of a DB instance to a CSV file for viewing and analyzing details.


Procedure

Step 1 In the instance list, click the target instance name.

Step 2 Click the **Parameters** tab.

Step 3 Click **Export** above the parameter list.

Figure 2-30 Export a parameter list



Step 4 In the displayed dialog box, enter a file name and click **OK**.

NOTE

The file name can contain 4 to 81 characters.

----End

2.9 Monitoring

2.9.1 Viewing Monitoring Metrics of a FlexusRDS for PostgreSQL Instance

Scenarios

This section describes how to view monitoring metrics of FlexusRDS for PostgreSQL instances and configure alarm rules. You can customize objects to be monitored and notification policies so that you can closely monitor your instances.

Viewing Metrics

Step 1 In the instance list, click the target instance name.

Step 2 Click the **Monitoring** tab and view the monitoring metrics of the instance.

Figure 2-31 Metrics




Table 2-5 Supported metrics

Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
Inactive Logical Replication Slots	Number of inactive logical replication slots	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
CPU Usage	CPU usage of the monitored object	0-100%	FlexusRDS for PostgreSQL instance	1 minute
Temporary Files per Minute	Number of temporary files generated within 1 minute	≥ 0 Count/min	FlexusRDS for PostgreSQL instance	1 minute
Memory Usage	Memory usage of the monitored object	0-100%	FlexusRDS for PostgreSQL instance	1 minute
Temporary File Size per Minute	Size of temporary files generated within 1 minute	≥ 0 Byte/min	FlexusRDS for PostgreSQL instance	1 minute
IOPS	Average number of I/O requests processed by the system in a specified period	≥ 0 counts/s	FlexusRDS for PostgreSQL instance	1 minute

Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
Synchronous Replication Blocking Time	Time during which synchronous replication between the primary and standby nodes is blocked	$\geq 0s$	FlexusRDS for PostgreSQL instance	1 minute
Network Input Throughput	Incoming traffic in bytes per second	≥ 0 bytes/s	FlexusRDS for PostgreSQL instance	1 minute
Stream Replication Status of Standby Instance or Read Replica	Stream replication status of the standby instance or a read replica. The value 0 indicates abnormal stream replication. A 1 indicates normal stream replication. A 2 means that this node is the primary instance.	Count	FlexusRDS for PostgreSQL instance	1 minute
Network Output Throughput	Outgoing traffic in bytes per second	≥ 0 bytes/s	FlexusRDS for PostgreSQL instance	1 minute
Inodes	Used data disk inodes	≥ 0 counts	FlexusRDS for PostgreSQL instance	5 minutes
Storage Space Usage	Storage usage of the monitored object	0-100%	FlexusRDS for PostgreSQL instance	1 minute
Transaction Logs Usage	Storage space usage of transaction logs	≥ 0 MB	FlexusRDS for PostgreSQL instance	1 minute
Replication Slot Usage	Storage space usage of replication slot files	≥ 0 MB	FlexusRDS for PostgreSQL instance	1 minute
Database Connections in Use	Number of database connections in use	≥ 0 counts	FlexusRDS for PostgreSQL instance	1 minute

Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
Maximum Used Transaction IDs	Maximum number of transaction IDs that have been used	≥ 0 counts	FlexusRDS for PostgreSQL instance	1 minute
Transaction Logs Generation	Size of transaction logs generated per second	≥ 0 MB/s	FlexusRDS for PostgreSQL instance	1 minute
Oldest Replication Slot Lag	Lagging size of the most lagging replica in terms of WAL data received	≥ 0 MB	FlexusRDS for PostgreSQL instance	1 minute
Replication Lag	Replication lag delay	≥ 0 ms	FlexusRDS for PostgreSQL instance	1 minute
Connections in Use	Number of connections in use (excluding built-in connections used for monitoring and O&M)	≥ 0 Counts	FlexusRDS for PostgreSQL instance	1 minute
Active Connections of Users	Number of active connections used by users (excluding built-in active connections used for monitoring and O&M)	≥ 0 counts	FlexusRDS for PostgreSQL instance	1 minute
TPS	Average execution times of submitted and rollback transactions per second	≥ 0 counts/s	FlexusRDS for PostgreSQL instance	1 minute
Connection Usage	Percent of used PostgreSQL connections to the total number of connections	0-100%	FlexusRDS for PostgreSQL instance	1 minute
Active Connections	Number of active connections to a DB instance	≥ 0	FlexusRDS for PostgreSQL instance	1 minute

Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
Idle Transaction Connections	Number of connections with idle transactions	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
Operation Rows	Number of rows that are being inserted, deleted, updated, or queried	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
Oldest Active Transaction Duration	Length of time since the start of the transaction that has been active longer than any other current transaction	≥ 0 ms	FlexusRDS for PostgreSQL instance	1 minute
Oldest Two-Phase Commit Transaction Duration	Length of time since the start of the transaction that has been prepared for two-phase commit longer than any other current transaction	≥ 0 ms	FlexusRDS for PostgreSQL instance	1 minute
Disk I/O Usage	I/O usage of disks The disk I/O usage is the percentage of the time that the disk processes I/O requests to the total time. NOTE If the disk I/O usage reaches 100%, data is being written to the disk during the statistical period. The disk performance is determined by multiple metrics, such as IOPS, disk throughput, and read/write latency.	0-100%	FlexusRDS for PostgreSQL instance	1 minute
Sessions Waiting for Locks	Number of sessions that are waiting for locks	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
Swap-In Rate	Volume of data written from the swap partition to the memory per second	≥ 0 KB/s	FlexusRDS for PostgreSQL instance	1 minute

Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
Swap-Out Rate	Volume of data written from the memory to the swap partition per second	≥ 0 KB/s	FlexusRDS for PostgreSQL instance	1 minute
Total Swap Size	Total size of the swap partition	≥ 0 MB	FlexusRDS for PostgreSQL instance	1 minute
Swap Usage	Usage of the swap partition	0-100%	FlexusRDS for PostgreSQL instance	1 minute
Maximum Database Age	Maximum age of the current database, which is the value of max(age(datfrozenxid)) in the pg_database table	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
User-mode CPU Time Percentage	Percentage of time that the CPU is in user mode	0-100%	FlexusRDS for PostgreSQL instance	1 minute
Kernel-mode CPU Time Percentage	Percentage of time that the CPU is in kernel mode	0-100%	FlexusRDS for PostgreSQL instance	1 minute
Disk I/O Wait Time Percentage	Percentage of time that the CPU is waiting for disk I/O operations to complete	0-100%	FlexusRDS for PostgreSQL instance	1 minute
Read I/O Latency	Average latency (in milliseconds) of disks responding to read requests	≥ 0 ms	FlexusRDS for PostgreSQL instance	1 minute
Write I/O Latency	Average latency (in milliseconds) of disks responding to write requests	≥ 0 ms	FlexusRDS for PostgreSQL instance	1 minute

Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
Read IOPS	Average number of read I/O requests processed by the system in a specified period	≥ 0 counts/s	FlexusRDS for PostgreSQL instance	1 minute
Write IOPS	Average number of write I/O requests processed by the system in a specified period	≥ 0 counts/s	FlexusRDS for PostgreSQL instance	1 minute
Number of SQL Statements Executed for More Than 1s	Number of slow SQL statements whose execution time is longer than 1s	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
Number of SQL Statements Executed for More Than 3s	Number of slow SQL statements whose execution time is longer than 3s	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
Number of SQL Statements Executed for More Than 5s	Number of slow SQL statements whose execution time is longer than 5s	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
Number of SQL Statements Executed for More Than log_min_duration_statement	Number of slow SQL statements whose execution time is longer than the value of log_min_duration_statement. You can change the value of this metric as required.	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
WAL Size	Size of WAL logs	≥ 0 GB	FlexusRDS for PostgreSQL instance	5 minutes

Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
Audit Log Size	Size of audit logs	≥ 0 GB	FlexusRDS for PostgreSQL instance	5 minutes
Memory Hit Rate	Memory hit rate	$\geq 0\%$	FlexusRDS for PostgreSQL instance	1 minute
Transactions Alive for More Than 1s	The number of long-running transactions that are kept alive for more than 1s	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
Transactions Alive for More Than 3s	The number of long-running transactions that are kept alive for more than 3s	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
Transactions Alive for More Than 5s	The number of long-running transactions that are kept alive for more than 5s	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
Transactions Idle for More Than 1s	The number of long-running transactions that have been idle for more than 1s	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
Transactions Idle for More Than 3s	The number of long-running transactions that have been idle for more than 3s	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
Transactions Idle for More Than 5s	The number of long-running transactions that have been idle for more than 5s	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
Two-Phase Transactions Running for More Than 1s	The number of long-running transactions that have been running for more than 1s in two phases	≥ 0	FlexusRDS for PostgreSQL instance	1 minute


Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
Two-Phase Transactions Running for More Than 3s	The number of long-running transactions that have been running for more than 3s in two phases	≥ 0	FlexusRDS for PostgreSQL instance	1 minute
Two-Phase Transactions Running for More Than 5s	The number of long-running transactions that have been running for more than 5s in two phases	≥ 0	FlexusRDS for PostgreSQL instance	1 minute

----End

Configuring Alarm Rules

Step 1 On the metric card, click + to go to the page for configuring alarm rules.

Figure 2-32 Selecting a metric



Step 2 On the **Create Alarm Rule** page, configure required parameters.

- **Name:** The system generates a random name for the alarm rule. You can change it as needed.

- **Description:** Add supplementary information about the rule.
- **Alarm Policy:** Policy for triggering an alarm. An alarm is triggered when the metric configured for this alarm reaches the preset threshold in consecutive periods.
- **Alarm Notification:** Specify whether to notify users when alarms are triggered. Notifications can be sent by email, text message, or HTTP/HTTPS message.

For details about how to configure alarm notification, see [Creating an Alarm Rule](#) in *Cloud Eye User Guide*.

Step 3 Click **Create**.

----End

2.10 Logs

2.10.1 Viewing and Downloading Error Logs of a FlexusRDS for PostgreSQL Instance

FlexusRDS for PostgreSQL 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 In the instance list, click the target instance name.

Step 2 Click the **Logs** tab. On the **Error Logs** tab page, 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 FlexusRDS instances, the following levels of logs are displayed:


- All log levels
- ERROR
- FATAL
- PANIC
- Error logs are displayed in log loading mode. There is no upper limit on the number of log records displayed within the query time range, and the total number of log records is not displayed.
- You can click  in the upper right corner to view logs generated in different time segments.
- 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** In the instance list, click the target instance name.
- Step 2** Click the **Logs** tab and click **Download Logs** on the right.
- Step 3** Locate the log file whose status is **Preparation completed** and click **Download** in the **Operation** column.

Figure 2-33 Downloading an Error Log



- 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.
- Log files of 40 MB to 100 MB can be downloaded.
- 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 redownload the log, click **OK**.
- The downloaded logs contain only the logs of the primary node.

----End

2.10.2 Viewing and Downloading Slow Query Logs of a FlexusRDS for PostgreSQL Instance

Scenarios

Slow query logs record statements that exceed the **log_min_duration_statement** value. 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 30 days can be viewed.

FlexusRDS for PostgreSQL supports the following statement types:

- All statement types
- SELECT
- INSERT
- UPDATE
- DELETE
- CREATE
- DROP
- ALTER
- DO
- CALL
- COPY
- WITH
- OTHER

Parameter Description

Table 2-6 Parameters related to slow queries

Parameter	Description
log_min_duration_statement	Specifies how many milliseconds a query has to run before it has to be logged. If this parameter is set to a smaller value, the number of log records increases, which increases disk I/O utilization and deteriorates the SQL performance.
log_statement	Specifies the SQL statement type. The value is ddl by default and cannot be changed.
cron.log_statement	Specifies whether to log all cron statements before running them. The default value is on .

Viewing Log Details

- Step 1** In the instance list, click the target instance name.
- Step 2** Click the **Logs** tab. On the **Slow Query Logs** tab page, view details about slow SQL statements.

NOTE


- You can view the slow query log records of a specified execution statement type or a specific time period.
- The **log_min_duration_statement** parameter determines when a slow query log is recorded. However, changes to this parameter do not affect already recorded logs. If this parameter is changed from 1,000 ms to 100 ms, FlexusRDS 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 1,500 ms SQL statement that was recorded when the threshold was 1,000 ms will not be deleted now that the new threshold is 2,000 ms.
- Currently, a maximum of 2,000 slow log records can be displayed.

----End

Downloading a Slow Query Log

- Step 1** In the instance list, click the target instance name.
- Step 2** Click the **Logs** tab and click **Download Logs** on the right.
- Step 3** In the displayed dialog box, click **Slow Query Logs**.
- Step 4** Locate the log file whose status is **Preparation completed** and click **Download** in the **Operation** column.

Figure 2-34 Downloading a Slow Query Log



- 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.
- 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 redownload the log, click **OK**.

- The downloaded logs contain only the logs of the primary node.

----End

2.11 Interconnection with CTS

2.11.1 FlexusRDS Operations Supported by CTS

With Cloud Trace Service (CTS), you can record operations associated with FlexusRDS instances for later query, audit, and backtrack operations.

Table 2-7 FlexusRDS operations that can be recorded by CTS

Operation	Resource Type	Trace Name
Creating a DB instance or restoring data to a new instance	instance	createInstance
Enabling autoscaling	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
Binding or unbinding an EIP	instance	setOrResetPublicIP
Adding a tag	instance	createTag
Deleting a tag	instance	deleteTag
Editing a tag	instance	modifyTag
Deleting a DB instance	instance	deleteInstance
Downloading a backup (using OBS)	backup	downloadSnapshot
Downloading a backup (using a browser)	backup	backupsDownload
Deleting a backup	backup	deleteManualSnapshot
Deleting a frozen DB instance	all	rdsUnsubscribeInstance
Freezing a DB instance	all	rdsfreezeInstance

Operation	Resource Type	Trace Name
Renewing a DB instance	all	bssUpdateMetadata

2.11.2 Viewing Traces

Scenarios

After Cloud Trace Service (CTS) is enabled, it starts recording operations on cloud resources. You can check the records generated over the last 7 days on the CTS console.

This section describes how to check the operation records generated over the last 7 days on the CTS console.

For details about how to view audit logs, see [Querying Real-Time Traces](#).

Prerequisites

Before using CTS, you need to enable it. For details, see [Enabling CTS](#).

Procedure

Step 1 [Log in to the management console](#).

Step 2 In the upper left corner of the page, click  and choose **Management & Governance > Cloud Trace Service** to go to the **Trace List** page on the CTS console.

Step 3 In the time range drop-down list above the trace list, select a desired query time range: **Last 1 hour**, **Last 1 day**, or **Last 1 week**. You can also select **Custom** to specify a custom time range within the last seven days.

Step 4 The search box above the trace list supports advanced queries. Combine one or more filters to refine your search.

Table 2-8 Trace filtering parameters

Parameter	Description
Trace Name	Name of a trace. The entered value is case-sensitive and requires an exact match. Fuzzy matching is not supported.
Trace Source	Cloud service name abbreviation. The entered value is case-sensitive and requires an exact match. Fuzzy matching is not supported.
Resource Name	Name of a cloud resource involved in a trace. The entered value is case-sensitive and requires an exact match. Fuzzy matching is not supported.

Parameter	Description
Resource ID	ID of a cloud resource involved in a trace. The entered value is case-sensitive and requires an exact match. Fuzzy matching is not supported.
Trace ID	Value of the trace_id parameter for a trace reported to CTS. The entered value requires an exact match. Fuzzy matching is not supported.
Resource Type	Type of a resource involved in a trace. The entered value is case-sensitive and requires an exact match. Fuzzy matching is not supported.
Operator	User who triggers a trace. Select one or more operators from the drop-down list. If the value of trace_type in a trace is SystemAction , the operation is triggered by the service and the trace's operator may be empty.
Trace Status	Select one of the following options from the drop-down list: • normal: The operation succeeded. • warning: The operation failed. • incident: The operation caused a fault that is more serious than a normal failure, for example, causing other faults.
Enterprise Project ID	ID of the enterprise project to which a resource belongs. To check enterprise project IDs, go to the Enterprise Project Management Service (EPS) console and choose Project Management in the navigation pane.
Access Key	Temporary or permanent access key ID. To check access key IDs, hover over your username in the upper right corner of the console and select My Credentials from the pop-up list. On the displayed page, choose Access Keys in the navigation pane.

Step 5 Check the traces that meet the search criteria.

Step 6 Click a trace name. Details about the trace are displayed in the dialog box on the right.

Step 7 Click **Export** to export all traces in the query result as an .xlsx file.

For details about key fields in a trace structure, see [Trace Structure](#) and [Example Traces](#) in the *Cloud Trace Service User Guide*.

----End

2.12 FlexusRDS for PostgreSQL Tags

Scenarios

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.

- Log in to the management console 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.
- Each DB instance can have up to 20 tags.

Editing a Tag

Step 1 In the instance list, click the target instance name.

Step 2 Click the **Tags** tab and click **Edit Tag**.

Step 3 In the displayed dialog box on the right, click **Add Tag**, enter a tag key and value, and click **OK**.

- 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 (optional) can consist of up to 255 characters and can include letters, digits, spaces, and the following characters: `_ . : / = + - @`.

Step 4 After a tag has been added, you can view and manage it on the **Tags** page.

-----End

Deleting a Tag

Step 1 In the instance list, click the target instance name.

Step 2 Click the **Tags** tab and click **Edit Tag**.

Step 3 In the displayed dialog box on the right, select the tag to be deleted and click **Delete**.

Step 4 Click **OK**.

After a tag has been deleted, it will no longer be displayed on the **Tags** page.

-----End

2.13 FlexusRDS for PostgreSQL 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 FlexusRDS 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 [Log in to the management console](#).

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

Step 3 In the upper right corner of the page, choose **Resources > My Quotas**.

Figure 2-35 My quotas



Step 4 On the **Quotas** page, view the used and total quotas of each type of resources.

----End

Increasing Quotas


Step 1 [Log in to the management console](#).

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

Step 3 In the upper right corner of the page, choose **Resources > My Quotas**.

Step 4 In the upper right corner of the page, click **Increase Quota**.

Figure 2-36 Increasing quotas

The image shows the 'Quotas' page in the Huawei Cloud management console. At the top right, there is a red button labeled 'Increase Quota'. Below it is a table with the following columns: 'Service', 'Resource Type', 'Used Quota', and 'Total Quota'. The table lists various services and their corresponding resource types and quotas.

Service	Resource Type	Used Quota	Total Quota
Elastic Cloud Server	EC2s	0	20
	vCPUs	0	200
	Memory (MB)	0	1,638,400
Dedicated Host	s3	0	5
	c3	0	5
	c6	0	100
Image Management Service	c6_pro	0	100
	Images	0	50
FunctionGraph	fFunction	0	400
	Code storage(MB)	0	10,240

Step 5 On the **Create Service Ticket** page, configure parameters as required.

In the **Problem Description** area, enter the required quota and reason for the quota adjustment.

Step 6 After all required parameters are configured, select the agreement and click **Submit**.

----End