

GaussDB

User Guide

Issue	01
Date	2025-08-18



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.

Huawei Cloud Computing Technologies Co., Ltd.

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

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

Contents

1 Buying a GaussDB Instance.....	1
2 GaussDB Instance Connection.....	8
2.1 Connecting to a GaussDB Instance.....	8
2.2 Connecting to an Instance Through DAS.....	10
2.3 Using gsql to Connect to an Instance.....	11
3 Instance Management.....	16
3.1 Configuring Security Group Rules for a GaussDB Instance.....	16
3.2 Binding and Unbinding an EIP for a GaussDB Instance.....	18
3.3 Modifying the Recycle Bin Policy for a GaussDB Instance.....	20
3.4 Exporting Information About All GaussDB Instances.....	21
3.5 Unsubscribing a Yearly/Monthly GaussDB Instance.....	21
3.6 Rebooting a GaussDB Instance.....	22
3.7 Deleting a GaussDB Instance.....	23
3.8 Rebuilding a GaussDB Instance.....	24
3.9 Rebooting a GaussDB Node.....	25
4 Instance Modifications.....	26
4.1 Changing the Name of a GaussDB Instance.....	26
4.2 Changing the CPU and Memory Specifications of a GaussDB Instance.....	27
4.3 Configuring Read Replicas.....	28
4.4 Scaling In and Out an Instance.....	31
4.4.1 Adding Coordinator Nodes for an Instance (Distributed).....	31
4.4.2 Adding Shards for an Instance (Distributed).....	32
4.4.3 Deleting Coordinator Nodes for an Instance (Distributed).....	33
4.5 Scaling Up Storage Space.....	34
4.5.1 Manually Scaling Up Storage Space for an Instance.....	34
4.5.2 Configuring Storage Autoscaling for an Instance.....	35
4.6 Changing the Deployment Model.....	38
4.6.1 Changing Standby DN to Log Nodes (for a Distributed Instance).....	38
4.6.2 Changing Log Nodes to Standby DNs (for a Distributed Instance).....	39
4.7 Performing a Primary/Standby DN Switchover.....	39
4.7.1 Changing the DN Failover Priority.....	39
4.7.2 Performing a Primary/Standby Switchover.....	41

5 Instance Upgrade.....	43
5.1 Overview.....	43
5.2 Hot Patch Update.....	46
5.3 In-place Upgrade.....	51
5.4 Gray Upgrade.....	55
6 Data Backup.....	64
6.1 Overview.....	64
6.2 Backup Execution.....	66
6.2.1 Configuring an Automated Backup Policy for GaussDB Instances.....	66
6.2.2 Configuring an Automated Backup Policy for GaussDB Tables.....	68
6.2.3 Creating a Manual Backup for GaussDB Instances.....	70
6.2.4 Creating a Manual Backup for GaussDB Tables.....	72
6.3 Backup Management.....	73
6.3.1 Exporting Backup Information About GaussDB Instances.....	73
6.3.2 Deleting a Manual Backup of a GaussDB Instance.....	74
7 Data Restoration.....	75
7.1 Restoring a Backup File to a GaussDB Instance.....	75
7.2 Restoring a GaussDB Database or Table Using a Backup File.....	78
7.3 Restoring a GaussDB Instance to a Specific Point in Time.....	82
7.4 Restoring a GaussDB Database or Table to a Specific Point in Time.....	86
8 Parameter Management.....	91
8.1 Modifying GaussDB Instance Parameters.....	91
8.2 Viewing Parameter Change History of a GaussDB Instance.....	93
8.3 Exporting Parameters of a GaussDB Instance.....	94
8.4 Creating a Custom Parameter Template for GaussDB Instances.....	95
8.5 Managing Parameter Templates for GaussDB Instances.....	96
9 Logs and Auditing.....	101
9.1 Downloading Slow Query Logs of a GaussDB Instance.....	101
9.2 Downloading Switchover/Failover Logs of a GaussDB Instance.....	102
9.3 Viewing GaussDB Operation Logs on CTS.....	102
9.4 Interconnecting with LTS and Querying Database Audit Logs.....	105
10 Quota Adjustment.....	107
10.1 Adjusting GaussDB Resource Quotas of an Enterprise Project.....	107
11 Disaster Recovery Management.....	110
11.1 Constraints.....	110
11.2 Creating a DR Relationship.....	111
11.3 Checking DR Task Statuses.....	113
11.4 Promoting the DR Instance to Primary.....	115
11.5 Stopping a DR Task.....	116
11.6 Deleting a DR Task.....	116

11.7 Re-creating a DR Task After the Primary Instance Is Faulty.....	117
11.8 Switching Roles of Primary and DR Instances.....	117
11.9 Re-establishing a DR Relationship.....	118
11.10 Performing a DR Drill.....	118
11.11 Enabling or Disabling Log Cache.....	119
11.12 Restrictions on the DR Instance.....	120
12 Managing GaussDB Tasks.....	127
13 Managing GaussDB Tags.....	129
14 Resetting the Administrator Password of a GaussDB Instance.....	132

1 Buying a GaussDB Instance

Scenarios

You can buy a GaussDB instance on the management console.

Procedure



- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Click  in the upper left corner of the page and choose **Databases > GaussDB**.
- Step 4** On the **Instances** page, click **Buy DB Instance**.
- Step 5** On the displayed page, configure parameters about the instance, and click **Next**.

Table 1-1 Basic information

Parameter	Description
Billing Mode	<p>GaussDB provides yearly/monthly billing and pay-per-use billing.</p> <ul style="list-style-type: none">• Yearly/Monthly: You pay upfront for the amount of time you expect to use the DB instance for. You will need to make sure you have a top-up account with a sufficient balance or have a valid payment method configured first.• Pay-per-use: You can start using the DB instance first and then pay as you go. Pricing is listed on a per-hour basis, but bills are calculated based on the actual usage duration.

Parameter	Description
Region	<p>A region where the tenant is located. You can change the region on the instance creation page, or go back to the Instances page and change it in the upper left corner.</p> <p>Products in different regions cannot communicate with each other over a private network. After the instance is created, you cannot change its region.</p>
DB Instance Name	<p>The instance name is case-sensitive, must start with a letter, and can contain 4 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.</p>
Edition Type	<p>GaussDB provides Basic edition and Enterprise edition.</p> <p>The basic edition lacks certain advanced features that are available in the enterprise edition. The basic edition delivers the same level of performance as the enterprise edition at a more affordable price. This edition is ideal for users who prioritize cost and do not need advanced features.</p>
DB Engine Version	<p>GaussDB supports the following versions: V2.0-3.226, V2.0-8.102, V2.0-8.103, and V2.0-8.210.</p> <p>NOTE</p> <ul style="list-style-type: none">Instances of V2.0-8.102 or later run on HCE.A DR relationship cannot be established between instances running HCE and those running EulerOS.
DB Instance Type	<ul style="list-style-type: none">Distributed: You can add nodes for distributed instances as needed to handle large volumes of concurrent requests.Centralized: Centralized instances are suitable for scenarios with small and stable volumes of data, where data reliability and service availability are extremely important.
Deployment Model	<ul style="list-style-type: none">Distributed<ul style="list-style-type: none">Independent: Database components are deployed on different nodes. This model is suitable for where high availability and stability are required and the instance scale is large.Centralized<ul style="list-style-type: none">HA (1 primary + 2 standby): 3-node deployment where there is one shard. The shard contains one primary DN and two standby DNs.Single: single-node deployment where there is only one CMS component and one DN. This model is only available for instances of V2.0-2.2 or later. <p>CAUTION</p> <ul style="list-style-type: none">Single: The availability (or SLA) cannot be guaranteed because the instance is deployed on a single server.

Parameter	Description
Transaction Consistency	<p>This parameter is available only to distributed instances.</p> <ul style="list-style-type: none">• Strong consistency: When an application updates data, every user can query all data that has been successfully committed, but performance is affected.• Eventual consistency: When an application updates data, the data users queried may be different, and some users may not obtain the most current value. The most current data may take a bit of time to become available for query by all users. However, DB instances with eventual consistency generally have higher performance. Eventual consistency cannot ensure strong read consistency of distributed transactions and consistency of transactions that depend on query results, such as <code>INSERT INTO SELECT * FROM</code>. Write operations that are split into multiple statements or involve in multiple nodes are not supported.
Replicas	<p>This parameter is only available for distributed instances.</p> <p>Total number of DNs each shard, primary and standby DNs combined. There are three replicas in a shard, indicating that there are one primary and two standby DNs in a shard.</p>
Shards	<p>This parameter is available only for distributed instances. It indicates the number of shards in an instance. A shard contains multiple DNs. The number of DNs in a shard depends on the value of Replicas, for example, if Replicas is set to 3, there are three DNs (one primary and two standby DNs) in a shard. The value ranges from 1 to 64.</p>
Coordinator Nodes	<p>This parameter is available only for distributed instances. It indicates the number of CNs in an instance. The value ranges from 1 to 64.</p> <p>It is recommended that at least two CNs be configured for an instance. Configuring only one CN results in low service reliability, making it best suited for test environments rather than production use.</p> <p>A CN provides the following functions:</p> <ul style="list-style-type: none">• It receives access requests from applications and returns execution results to clients.• It breaks down tasks and distributes task fragments to different DNs for parallel processing.
AZ	<p>An AZ is a physical region where resources have independent power supply and networks. AZs are physically isolated but interconnected through an internal network.</p> <p>A DB instance can be deployed in one AZ or three AZs. Three AZs are recommended.</p>

Parameter	Description
Time Zone	Select a time zone according to the region hosting your DB instance when you buy the instance.

Table 1-2 Specifications and storage

Parameter	Description
Instance Specifications	CPU and memory specifications of the instance. Different instance specifications support different numbers of database connections. For details, see "DB Instance Description" > "Instance Specifications" in <i>GaussDB Service Overview</i> .
Storage Type	The storage type determines the read/write speed of an instance. The higher the maximum throughput is, the higher the instance read/write speed can be.
Storage Space (GB)	The storage space contains the file system overhead required for inodes, reserved blocks, and database operations. When you create an instance, the storage for a single shard starts from 40 GB. You can increase the storage in increments of 4 GB. If too little storage is configured, when traffic volume is high, the storage can be used up quickly, and the instance will change to read-only. Select an amount of storage based on how much traffic you expect there to be.
Disk Encryption	<ul style="list-style-type: none">• Disable: Encryption is disabled.• Enable: Encryption is enabled, which improves data security but affects system performance. Key Name: If disk encryption is enabled, you need to select or create a key.

Table 1-3 Network

Parameter	Description
VPC	A virtual network where your GaussDB instances are located. A VPC isolates networks for different workloads. You need to create or select the required VPC. After the GaussDB instance is created, the VPC cannot be changed.

Parameter	Description
Subnet	<p>A subnet provides dedicated network resources that are logically isolated from other networks for network security. Subnets take effect only within a specific AZ. Dynamic Host Configuration Protocol (DHCP) is enabled by default for subnets in which you plan to create GaussDB instances and cannot be disabled. GaussDB supports automatic IP address allocation during instance creation.</p>
Security Group	<p>Controls access to and from an instance. When you buy an instance, the selected security group must meet the following requirements:</p> <ul style="list-style-type: none">• If you need to change the security group when buying a distributed instance, ensure that the TCP ports in the inbound rule include the following: 40000-60480, 20050, 5000-5001, 2379-2380, 6000, 6001, 6500, and <code><database_port>-(<database_port> + 100)</code>. (For example, if the database port is 8000, the security group must contain ports 8000 to 8100.) Additionally, ensure that the outbound security group rules allow all outbound traffic.• If you need to change the security group when buying a centralized instance, ensure that the TCP ports in the inbound rule include the following: 20050, 5000-5001, 2379-2380, 6000, 6500, and <code><database_port>-(<database_port> + 100)</code>. (For example, if the database port is 8000, the TCP ports for the security group must include 8000-8100.) <p>The security group enhances security by controlling access to GaussDB from other services. Ensure that the security group you selected allows your client to access the instance.</p>
Single Private IP Address	<p>Specifies whether to enable the single private IP address policy. If this policy is enabled, only one private IP address is assigned to an instance and is bound to the primary node. The private IP address does not change after a primary/standby switchover. If this policy is disabled, each node is bound to a private IP address, and the private IP address changes after a primary/standby switchover.</p> <p>The constraints on the single private IP address policy are as follows:</p> <ul style="list-style-type: none">• This policy is only available for centralized instances of version V2.0-3.207 or later.• This policy is configurable only during instance creation and cannot be modified afterwards.

Table 1-4 Database configuration

Parameter	Description
Administrator	Database administrator. The default username is root .
Administrator Password	Enter a strong password and periodically change it to improve security, preventing security risks such as brute force cracking. The password must contain: <ul style="list-style-type: none">• 8 to 32 characters.• At least three types of the following: uppercase letters, lowercase letters, digits, and special characters. Supported special characters: ~!@#%^*_-=+?, Keep the password secure. Password retrieval is not supported.
Confirm Password	Enter the administrator password again.

Table 1-5 Parameter templates

Parameter	Description
Parameter Template	A template of parameters for creating an instance. The template contains engine configuration values that are applied to one or more instances. After creating an instance, you can modify the parameter template.
Enterprise Project	If the instance has been associated with an enterprise project, select the target project from the Enterprise Project drop-down list. You can also go to the Enterprise Project Management console to create a project. .

 **NOTE**

The performance of your GaussDB instance depends on its settings. Hardware items include the instance specifications, storage type, and storage space.

Step 6 Confirm the displayed details.

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

Step 7 After the creation task is submitted, check the instance status on the **Instances** page. When a GaussDB instance is being created, its status is **Creating**. Once created, the instance becomes **Available**. You can view and manage it on the **Instances** page.

An automated full backup is immediately triggered once your instance is created.

----End

2 GaussDB Instance Connection

2.1 Connecting to a GaussDB Instance

GaussDB instances can be connected using gsql or Data Admin Service (DAS).

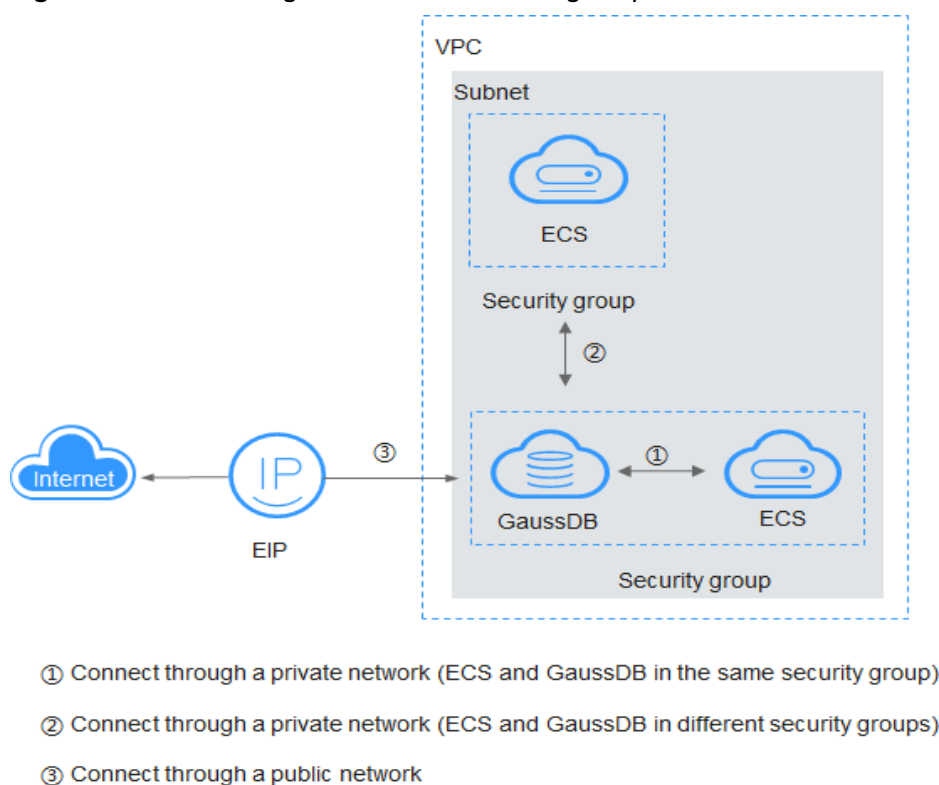
Table 2-1 GaussDB instance connection modes

Connect Through	IP Address	Description	Comments
DAS	Not required	DAS enables you to manage databases on a web-based console. It supports SQL execution, advanced database management, and intelligent O&M, simplifying database management and improving both efficiency and data security. The permissions required for connecting to a GaussDB instance through DAS are enabled by default.	Easy to use, secure, advanced, and intelligent

Connect Through	IP Address	Description	Comments
gsql	Private IP address/EIP	gsql is a client tool provided by GaussDB. You can use gsql to connect to the database and then enter, edit, and execute SQL statements in an interactive manner.	To achieve a higher data transmission rate and security level, migrate your applications to a server that is in the same subnet as your GaussDB instance and use a private IP address to access the instance. The bandwidth is not limited for private network connections.

[Figure 2-1](#) shows how an instance is connected.

Figure 2-1 Connecting to an instance through a private network and an EIP



 NOTE

- If the ECS and GaussDB instance are in the same VPC and security group, they can communicate with each other through the private network by default. In this case, you can connect to the instance through a private IP address.
- If the ECS and GaussDB instance are in the same VPC but different security groups, you need to set security group rules for both the GaussDB instance and ECS, and then connect to the instance through a private IP address.
 - GaussDB instance: Configure an **inbound** rule for the security group with which the GaussDB instance is associated. For details, see [Configuring Security Group Rules for a GaussDB Instance](#).
 - ECS: The default security group rule allows all outgoing data packets. In this case, you do not need to configure a security group rule for the ECS. If **not all outbound traffic is allowed** in the security group, you need to configure an **outbound** rule for the ECS to allow all outbound packets.
- If the ECS and GaussDB instance are in different VPCs, you can bind an EIP to the ECS and use the EIP to connect to the instance. Ensure that both the ECS and GaussDB instance have EIPs.
 - For details about how to bind an EIP to an ECS, see [Binding an EIP](#).
 - For details about how to bind an EIP to a GaussDB instance, see [Binding an EIP](#).
- Database connection specifications:
 - As concurrency increases, establishing a new connection takes longer, growing from milliseconds to seconds.
 - When the database is under high load (CPU usage $\geq 80\%$), the connection establishment time increases significantly, for example, from 210 ms to 1,400 ms for 50 concurrent connections.
 - To ensure connection establishment time remains below 500 ms, limit the connection rate to no more than 20 connections per second.

2.2 Connecting to an Instance Through DAS

Scenarios


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

This section describes how to connect to a GaussDB instance through DAS.

Procedure

Step 1 Log in to the management console.

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

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

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

Alternatively, click the DB instance name on the **Instances** page. On the displayed **Basic Information** page, click **Log In** in the upper right corner of the page.

- Step 5** Enter the correct database username and password, and click **Test Connection**. After the connection test is successful, click **Log In**.

Table 2-2 Parameter description

Parameter	Description
Login Username	Username of the GaussDB database account. The default administrator is root .
Database Name	Name of the database to be connected. The default management database is postgres .
Password	Password of the database user.
Show Executed SQL Statements	You are advised to enable Show Executed SQL Statements . With it enabled, you can view the executed SQL statements under SQL Operations > SQL History and execute them again without entering the SQL statements.

----End

2.3 Using gsql to Connect to an Instance

This section describes how to use the gsql client to connect to a GaussDB instance you have bought on the GaussDB management console.

- **Step 1: Buy an ECS**
- **Step 2: Query the IP Address and Port Number of the Instance to Be Connected**
- **Step 3: Test the Connectivity**
- **Step 4: Connect to the Database**
 - **Non-SSL connection**
 - **SSL connection**

Buying an ECS

If you want to connect to a database using the command-line interface (CLI), like gsql, you need to create an ECS and install gsql on it.

1. Log in to the management console and check whether there is an available ECS.
 - If there is, go to **3**.
 - If there is not, go to **2**.
2. Buy an ECS that runs EulerOS.
For details about how to create an ECS, see section "Purchasing an ECS" in *Elastic Cloud Server User Guide*.
3. On the **ECS Information** page of the target ECS, view the region and VPC of the ECS.

 **NOTE**



The ECS must run on EulerOS. gsql supports the following OS versions:

For x86 servers: EulerOS V2.0SP5 and Kylin V10 SP2

For Kunpeng servers: EulerOS V2.0SP8 and Kylin V10 SP1

4. On the **Basic Information** page of your GaussDB instance, view the region and VPC of the instance.
5. Check whether the ECS and GaussDB instance are in the same region and VPC.
 - If the ECS and GaussDB instance are in the same region and VPC, the DB instance can be connected through a private network. For details about how to obtain the private IP address, see [Querying the IP Address of the Instance to Be Connected](#).
 - If the ECS and DB instance are in different VPCs, the DB instance must be connected over a public network. For details about how to obtain the public IP address, see [Querying the IP Address of the Instance to Be Connected](#). Ensure that both the ECS and GaussDB instance have EIPs.
 - For details about how to bind an EIP to an ECS, see [Binding an EIP](#).
 - For details about how to bind an EIP to a GaussDB instance, see [Binding an EIP](#).

Querying the IP Address and Port Number of the Instance to Be Connected

1. Log in to the management console.
2. Click  in the upper left corner and select a region and project.
3. Click  in the upper left corner of the page and choose **Databases > GaussDB**.
4. On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.
5. In the **Connection Information** area, view the port number.

For a centralized instance, obtain the IP address of the primary DN. For a distributed instance, obtain the IP address of any CN.

 - If the ECS and GaussDB instance are in the same VPC, obtain the private IP address.
 - If the ECS and GaussDB instance are in different VPCs, obtain the EIP.

Testing Connectivity

1. Log in to the ECS. For details, see section "Logging In to a Linux ECS Using VNC" in *Elastic Cloud Server User Guide*.
2. On the ECS, check whether it can connect to the target GaussDB instance using the IP address and port number obtained in [Querying the IP Address and Port Number of the Instance to Be Connected](#).

telnet *IP address Port number*

Example:

telnet 192.168.0.16 8000

NOTE

If the message "command not found" is displayed, install a Telnet client that matches the OS of the ECS.

- If the ECS can connect to the DB instance, no further action is required.
- If the communication fails, check the security group rules.
 - On the **Outbound Rules** page of the ECS, add the IP address and port of the GaussDB instance to the outbound rules.
 - If the ECS and GaussDB instance are in the same VPC, add the private IP address and port of the GaussDB instance to the outbound rules.
 - If the ECS and GaussDB instance are in different VPCs, add the EIP address and port of the GaussDB instance to the outbound rules.
 - On the **Inbound Rules** page of the GaussDB instance, add the IP address and port of the ECS to the inbound rules.
 - If the ECS and GaussDB instance are in the same VPC, add the private IP address and port of the ECS to the inbound rules.
 - If the ECS and GaussDB instance are in different VPCs, add the EIP address and port of the ECS to the inbound rules.

For details, see [Configuring Security Group Rules](#).

Connecting to a Database

GaussDB lets you select either SSL or non-SSL connections as required.

Non-SSL Connection

1. Log in as the **root** user to the ECS you have created.
2. Upload the client tool package and configure gsql environment variables.
 - a. Run the following command to create the **/tmp/tools** directory for storing the client tool package:

```
mkdir /tmp/tools
```
 - b. Download the **GaussDB_driver.zip** driver package of the required version, and upload it to the **/tmp/tools** directory of the created ECS.
 - c. Run the following commands to decompress the **GaussDB_driver.zip** driver package:




```
cd /tmp/tools
unzip GaussDB_driver.zip
```
 - d. Run the following commands to copy the decompressed **GaussDB-Kernel_***_EULER_64bit-Gsql.tar.gz** client tool package to the **/tmp/tools** directory:

NOTE

This section uses the gsql tool package suitable for the centralized instances running on Euler2.5_x86_64 as an example. The relative path of the tool package varies depending on where you decompressed it.

- ```
cd /tmp/tools/GaussDB_driver/Centralized/Euler2.5_X86_64/
cp GaussDB-Kernel_***_EULER_64bit-Gsql.tar.gz /tmp/tools
```
- e. Run the following commands to decompress the package:
- ```
cd /tmp/tools
tar -zxvf GaussDB-Kernel_***_EULER_64bit-Gsql.tar.gz
```
- f. Configure environment variables.
- Run the following command to open the `~/.bashrc` file:
- ```
vim ~/.bashrc
```
- Press **G** to move the cursor to the last line, press **i** to enter Insert mode, and type the following information. Then, press **Esc** to exit Insert mode, and run **:wq** to save the settings and exit.
- ```
export PATH=/tmp/tools/bin:$PATH
export LD_LIBRARY_PATH=/tmp/tools/lib:$LD_LIBRARY_PATH
```
- Run the following command to make the environment variables take effect permanently:
- ```
source ~/.bashrc
```
3. Enter the password when prompted to connect to the database.
- After an instance is created, a **postgres** database is generated by default. Database **postgres** is used as an example.
- ```
gsql -d postgres -h 10.0.0.0 -U root -p 8000
```
- Password for user root:
- postgres** is the name of the database you want to connect. **10.0.0.0** is the IP address of the instance obtained in [Querying the IP Address of the Instance to Be Connected](#). **root** is the username for logging in to the database. **8000** is the database port obtained in [Querying the Port Number of the Instance to Be Connected](#).
- If the following information is displayed, the connection is successful:
- ```
gaussdb=>
```

## SSL Connection

- Log in to the management console.
  - Click  in the upper left corner and select a region and project.
  - Click  in the upper left corner of the page and choose **Databases > GaussDB**.
  - On the **Instances** page, click the name of the target instance. In the **DB Information** area on the **Basic Information** page, click  next to the **SSL** field to download the root certificate or certificate bundle.
  - Upload the root certificate to the ECS or save it to the device to be connected to the GaussDB instance.
- Import the root certificate to the Linux ECS. You can use a connection tool (such as WinSCP or PuTTY) to upload the certificate to any directory on a Linux server.
- Connect to a GaussDB instance.
- A Linux ECS is used in this example. Run the following command to set environment variables on the ECS:
- ```
export PGSSLMODE=<sslmode>
export PGSSLROOTCERT=<ca-file-directory>
```
- ```
gsql -h <host> -p <port> -d <database> -U <user>
```

Example:

```
export PGSSLMODE="verify-ca"
export PGSSLROOTCERT="/home/Ruby/ca.pem"
```

```
gsql -h 10.0.0.0 -p 8000 -d postgres -U root
```

Password for user root:

**Table 2-3** Parameter description

| Parameter                        | Description                                                                                                                                                                                                                                                                                                                           |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>&lt;host&gt;</i>              | IP address of the DB instance. To obtain the IP address, click the instance name on the <b>Instances</b> page to go to the <b>Basic Information</b> page of the instance. The IP address can be found in the <b>Private IP Address</b> field of the <b>Connection Information</b> area if the DB instance is accessed through an ECS. |
| <i>&lt;port&gt;</i>              | Database port number. The default value is <b>8000</b> . To obtain the database port, click the instance name on the <b>Instances</b> page to go to the <b>Basic Information</b> page of the instance. The database port can be founded in the <b>Database Port</b> field of the <b>Connection Information</b> area.                  |
| <i>&lt;database&gt;</i>          | Name of the database to connect to. The default database is <b>postgres</b> .                                                                                                                                                                                                                                                         |
| <i>&lt;user&gt;</i>              | Username of the GaussDB database account. The default administrator is <b>root</b> .                                                                                                                                                                                                                                                  |
| <i>&lt;ca-file-directory&gt;</i> | Directory of the CA certificate for SSL connection.                                                                                                                                                                                                                                                                                   |
| <i>&lt;sslmode&gt;</i>           | SSL connection mode. Set it to <b>verify-ca</b> to verify that the server is trustworthy by checking the certificate chain.                                                                                                                                                                                                           |

7. Check the command output after you log in to the database. If information similar to the following is displayed, the SSL connection has been established.  
SSL connection (cipher: DHE-RSA-AES256-GCM-SHA384, bits: 256)

# 3 Instance Management

---

## 3.1 Configuring Security Group Rules for a GaussDB Instance

### Scenarios

A security group is a collection of access control rules for ECSs and GaussDB instances that are within the same VPC, have the same security requirements, and are mutually trusted.

To ensure database security and reliability, you need to configure security group rules to allow specific IP addresses and ports to access the GaussDB instances.

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

### Precautions

The default security group rule allows all outbound data packets. This means that ECSs and GaussDB instances associated with the same security group can access


each other by default. After a security group is created, you can add security group rules to control the access from and to the GaussDB instance.

- By default, you can create up to 500 security group rules.
- Ensure that each security group has no more than 50 rules.
- To access a GaussDB instance from resources outside the security group, configure an **inbound rule** for the security group associated with the instance.
- Outbound rules typically do not apply to DB instances. The rules are used only when a DB instance acts as a client.
- If a DB instance resides in a VPC but is not publicly accessible, you can also use a VPN connection to connect to it.
- If you need to change the security group when creating a distributed instance, ensure that the TCP ports in the inbound rule include the following: 40000-60480, 20050, 5000-5001, 2379-2380, 6000, 6500, and  $\text{<database\_port>}-(\text{<database\_port>} + 100)$ . (For example, if the database port is 8000, the security group must contain ports 8000 to 8100.) Additionally, ensure that the outbound rules allow all outbound traffic.
- If you need to change the security group when creating a centralized instance, ensure that the TCP ports in the inbound rule include the following: 20050, 5000-5001, 2379-2380, 6000, 6500, and  $\text{<database\_port>}-(\text{<database\_port>} + 100)$ . (For example, if the database port is 8000, the TCP ports for the security group must include 8000-8100.)
- The default value of **Source** is **0.0.0.0/0**, indicating that all IP addresses can access the GaussDB instance as long as they are associated with the same security group as the instance.

## Procedure

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** Configure security group rules.

In the **Security Group** field of the **Connection Information** area, click the security group name.

**Step 6** On the **Inbound Rules** tab, click **Add Rule**. In the displayed dialog box, configure the required parameters and click **OK**.

You can click **+** to add more inbound rules.

**Table 3-1** Inbound rule parameter description

| Parameter       | Description                                                                                                                                                                                                                                             | Example Value                                                                                                       |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| Protocol & Port | Network protocol. Currently, the value can be <b>All</b> , <b>TCP</b> , <b>UDP</b> , <b>ICMP</b> , <b>GRE</b> , or others.                                                                                                                              | TCP (Custom ports)                                                                                                  |
|                 | <b>Port</b> : port or port range over which the traffic can reach your ECS. The value ranges from 1 to 65535.                                                                                                                                           | When connecting to your instance through a private network, enter the port of the used to connect to your instance. |
| Type            | IP address type. <ul style="list-style-type: none"><li>IPv4</li><li>IPv6</li></ul>                                                                                                                                                                      | IPv4                                                                                                                |
| Source          | Source of the security group rule. The value can be a security group or an IP address. Examples: <ul style="list-style-type: none"><li>xxx.xxx.xxx.xxx/32 (IPv4 address)</li><li>xxx.xxx.xxx.0/24 (subnet)</li><li>0.0.0.0/0 (any IP address)</li></ul> | 0.0.0.0/0                                                                                                           |
| Description     | Provides supplementary information about the security group rule. This parameter is optional.<br><br>The description can contain up to 255 characters and cannot contain angle brackets (<) or (>).                                                     | -                                                                                                                   |

----End

## 3.2 Binding and Unbinding an EIP for a GaussDB Instance

### Scenarios

You can bind an EIP to a GaussDB instance for public access and can unbind the EIP from an instance as required.

### Precautions

- To ensure that the DB instance can be accessed, the security group used by the DB instance must allow access to the database port. For example, if the


database port is **1611**, ensure that the security group allows access to the port **1611**.

- 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.
- An EIP can be bound to only one node IP address of a DB instance.

## Binding an EIP

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the **Connection Information** area, locate the private IP address and click **Bind** in the **Operation** column.

**Step 6** In the displayed dialog box, all available unbound EIPs are listed. Select the required EIP and click **OK**.

If no available EIPs are displayed, click **View EIP** and obtain an EIP.

**Step 7** In the **Connection Information** area, check the operation result in the **EIP** column of the private IP address list.

To unbind the EIP from the instance, see [Unbinding an EIP](#).

----End

## Viewing Traffic Information

**Step 1** Bind an EIP to the instance.

**Step 2** Click **View Traffic** to check the result.


**Step 3** Check the curves of the following metrics:


Uplink Bandwidth, Downlink Bandwidth, Outbound Bandwidth Usage, Downlink Traffic, and Uplink Traffic

----End

## Unbinding an EIP

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

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

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

- Step 4** On the **Instances** page, click the instance that has been bound with an EIP.
- Step 5** In the **Connection Information** area, click **Unbind** in the **Operation** column of the relevant IP address.
- Step 6** In the displayed dialog box, click **Yes** to unbind the EIP.
- Step 7** In the **Connection Information** area, check the operation result in the **EIP** column of the private IP address list.

To bind an EIP to the instance again, see [Binding an EIP](#).

----End

## 3.3 Modifying the Recycle Bin Policy for a GaussDB Instance

GaussDB automatically moves deleted instances to the recycle bin. This section describes how to set the retention period of instances in the recycle bin.



### Precautions

- The recycle bin is enabled by default and cannot be disabled. The deleted instances can be retained for 7 days by default.
- The new recycling policy only applies to instances that were put in the recycle bin after the new policy was put into effect. Those already in the recycle bin follow the original policy.

### Billing

Instances in the recycle bin will not incur charges.

### Procedure

- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Click  in the upper left corner of the page and choose **Databases > GaussDB**.
- Step 4** In the navigation pane on the left, choose **Recycle Bin**.
- Step 5** Click **Modify Recycling Policy**. In the displayed dialog box, set the retention period for the deleted DB instances from 1 day to 7 days.
- Step 6** Click **OK**.

----End

## 3.4 Exporting Information About All GaussDB Instances



### Scenarios

You can export information about all instances in the list for review and analysis.



### Constraints

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

### Exporting All Instance Information

- Step 1** Log in to the management console.
  - Step 2** Click  in the upper left corner and select a region and project.
  - Step 3** Click  in the upper left corner of the page and choose **Databases > GaussDB**.
  - Step 4** On the **Instances** page, click **Export Instance Info**. By default, all instance information is exported.
  - Step 5** In the displayed dialog box, select the items to be exported and click **OK**.
  - Step 6** After the export task is complete, a CSV file is generated on the local PC.
- End

### Exporting Information About Selected Instances

- Step 1** Log in to the management console.
  - Step 2** Click  in the upper left corner and select the desired region and project.
  - Step 3** Click  in the upper left corner of the page and choose **Databases > GaussDB**.
  - Step 4** On the **Instances** page, select the instances whose information you want to export and click **Export Instance Info**.
  - Step 5** In the displayed dialog box, select the items to be exported and click **OK**.
  - Step 6** After the export task is complete, a CSV file is generated on the local PC.
- End

## 3.5 Unsubscribing a Yearly/Monthly GaussDB Instance

### Scenarios

To delete a DB instance billed on a yearly/monthly basis, you need to unsubscribe the order. Currently, DB instances cannot be unsubscribed in batches. You can

unsubscribe only one instance at a time. For details, see [Unsubscribing a Single Instance](#).

For pay-per-use DB instances, you need to delete them on the **Instances** page. For details, see [Deleting a GaussDB Instance](#).


## Precautions


- After an unsubscription request is submitted, resources and data will be deleted and cannot be retrieved.
- If you want to retain data, complete a manual backup before submitting the unsubscription request.
- A maximum of 100 resources can be unsubscribed at a time.
- Unsubscribed instances will be moved to the recycle bin, but will be permanently deleted after a length of time determined by the recycling policy. Automated backups are deleted with the instances, but manual backups are retained and still billed. To delete the manual backups, go to the **Backups** page on the console.

## Unsubscribing a Single Instance

Unsubscribe a yearly/monthly DB instance on the **Instances** page.

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

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

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

**Step 4** On the **Instances** page, locate the instance and choose **More > Unsubscribe** in the **Operation** column.

**Step 5** On the displayed page, confirm the order to be unsubscribed and select a reason. Then, click **Confirm**.

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

**Step 7** View the unsubscription result. After the DB instance order is successfully unsubscribed, the DB instance is no longer displayed in the instance list on the **Instances** page.

----End

## 3.6 Rebooting a GaussDB Instance

### Scenarios

You can reboot a DB instance for the modifications to take effect.


## Precautions

- You can reboot a DB instance only when its status is **Available**. Your database may be unavailable in some cases, for example, when some modifications are being made.
- Rebooting a DB instance will cause service interruptions. During this period, the DB instance status is **Rebooting**.
- An instance is not available when it is being rebooted. After the reboot completes, the cached memory will be automatically cleared. You are advised to reboot the instance during off-peak hours.
- To quickly reboot a DB instance, perform fewer operations on the DB instance.
- If there are a large number of slow SQL statements or sessions, or if the thread pool is full, the reboot process may take a longer time than usual.

## Procedure

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

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

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

**Step 4** On the **Instances** page, locate the instance you want to reboot and choose **More > Reboot** in the **Operation** column.

Alternatively, click the instance name to go to the **Basic Information** page. Click **Reboot** in the upper right corner of the page.

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

The instance status becomes **Rebooting**.

**Step 6** Refresh the DB instance list and view the status of the DB instance. If its status is **Available**, it has been rebooted.

----End

## 3.7 Deleting a GaussDB Instance

### Scenarios



On the **Instances** page, you can manually delete DB instances that fail to be created or unnecessary instances after service changes.

### Precautions

- Instance deletion cannot be undone. Exercise caution when performing this operation. Back up your data first if you want to keep it after deleting an instance.

- Instances cannot be deleted while an operation is being performed on them. They can only be deleted once the operations are complete.
- Instances in the recycle bin will not incur charges.
- When instances are deleted, manual backups are retained.

## Procedure

- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Click  in the upper left corner of the page and choose **Databases > GaussDB**.
- Step 4** On the **Instances** page, locate the instance you want to delete and click **More > Delete** in the **Operation** column.
- Step 5** In the displayed dialog box, click **OK**. Refresh the **Instances** page later to check that the deletion is successful.



----End

## 3.8 Rebuilding a GaussDB Instance

### Scenarios

When you delete an instance, the system automatically creates a full backup and stores it in the recycle bin. Within the retention period, you can restore (also known as rebuild in this context) this instance by creating a new instance using the backup file. However, note that the restored instance will have a different private IP address from the original one.

## Procedure

- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Click  in the upper left corner of the page and choose **Databases > GaussDB**.
- Step 4** In the navigation pane on the left, choose **Recycle Bin**.
- Step 5** Locate the instance to be rebuilt and click **Rebuild** in the **Operation** column.
- Step 6** On the displayed page, configure required parameters and submit the task.
- Step 7** After the rebuilding task is submitted, check the instance status on the **Instances** page. The rebuild is complete when the status shows **Available**.

----End

## 3.9 Rebooting a GaussDB Node

### Scenarios


If the status of a GaussDB instance node is abnormal, you can reboot the node to restore the node status. You can also reboot a node when it is in the **Available** state. A node is not available when it is being rebooted.


### Precautions

- You can reboot a node when the DB instance is in the following state or performing the following operations:
  - Backup and restoration failed
  - Changing the billing mode from pay-per-use to yearly/monthly
  - DR in progress for the primary instance in a streaming DR task
  - Caching logs for the primary instance in a streaming DR task
  - DR simulation in progress for the DR instance in a streaming DR task
  - DR in progress for the DR instance in a streaming DR task
  - DR instance promoted to primary in a streaming DR task
- Rebooting nodes will clear the cached memory in them. To prevent traffic congestion during peak hours, you are advised to reboot nodes during off-peak hours.
- Only nodes of centralized instances can be rebooted.
- A primary/standby switchover will be triggered if a primary node is rebooted.

### Procedure

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

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

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

**Step 4** Click the target instance name to go to the **Basic Information** page.

**Step 5** In the **Node List** area, click **Reboot** in the **Operation** column of a node. Confirm information about the node to be restarted, enter **reboot**, and click **OK**.

The node status becomes **Rebooting node**.

**Step 6** Refresh the instance basic information and check the reboot result. If the node status is **Available**, the reboot is successful.

----End

# 4 Instance Modifications

---

## 4.1 Changing the Name of a GaussDB Instance

### Scenarios

You can change the name of an instance.

### Constraints

You cannot perform the following operations when the instance name is being changed:

- Deleting the instance
- Creating a backup for the instance


### Precautions


- The new name of an instance can be the same as an existing instance name.
- Changing the name of a DB instance does not disassociate the associated tags from the instance.
- If a DB instance is renamed, backups of the DB instance are still retained.


### Procedure

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



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

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

**Step 4** On the **Instances** page, locate the instance whose name you want to edit and click  next to the instance name. Then, edit the name and click **OK**.

Alternatively, click the instance name to go to the **Basic Information** page. In the **DB Information** area, click  next to the **DB Instance Name** field to edit the instance name.

The name must start with a letter and consist of 4 to 64 characters. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), and underscores (\_).

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

**Step 5** View the new instance name.

----End

## 4.2 Changing the CPU and Memory Specifications of a GaussDB Instance

### Scenarios

You can change the CPU and memory specifications of an instance to match your workload requirements. This includes scaling up or down the specifications of your GaussDB instance.

#### NOTE


This function is only available to authorized users. To apply for the permissions needed, [submit a service ticket](#).


### Precautions

- Before you change the instance specifications, ensure that the instance is available. If the instance or node is abnormal, or the storage space is full, you cannot perform this operation.
- During the specification change for an HA (1 primary + 2 standby) instance, a primary/standby failover is triggered. During the failover, services are interrupted for about 1 minute.
- For a single-replica instance, changing instance specifications will reboot the instance and interrupt services for 5 to 10 minutes.
- After you change instance specifications, the DB instances will be rebooted and services will be interrupted. You are advised to perform this operation during off-peak hours.
- If the instance load is heavy, it takes a longer time to change its instance specifications.
- If there is only one coordinator node in an instance, services will be interrupted during the specification change. Exercise caution when performing this operation.

### Procedure

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

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

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

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

Alternatively, click the instance name to go to the **Basic Information** page. In the **DB Information** area, click **Change** in the **Instance Specifications** field.

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

**Step 6** Confirm the specifications and click **Submit**.

**Step 7** View the new instance specifications.

After the task is submitted, click **Go to Instance List**. On the **Instances** page, the DB instance status is **Changing instance specifications**. After a few minutes, view the new instance specifications on the **Basic Information** page.

----End

## 4.3 Configuring Read Replicas

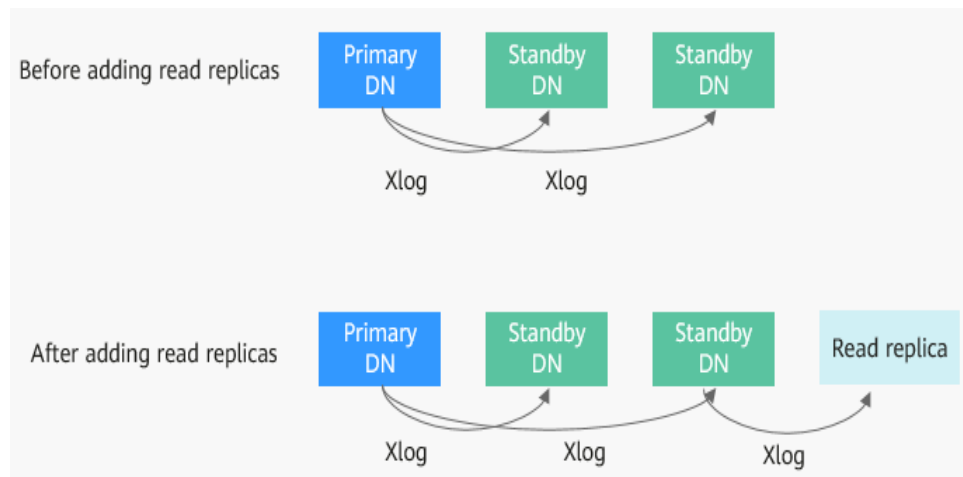
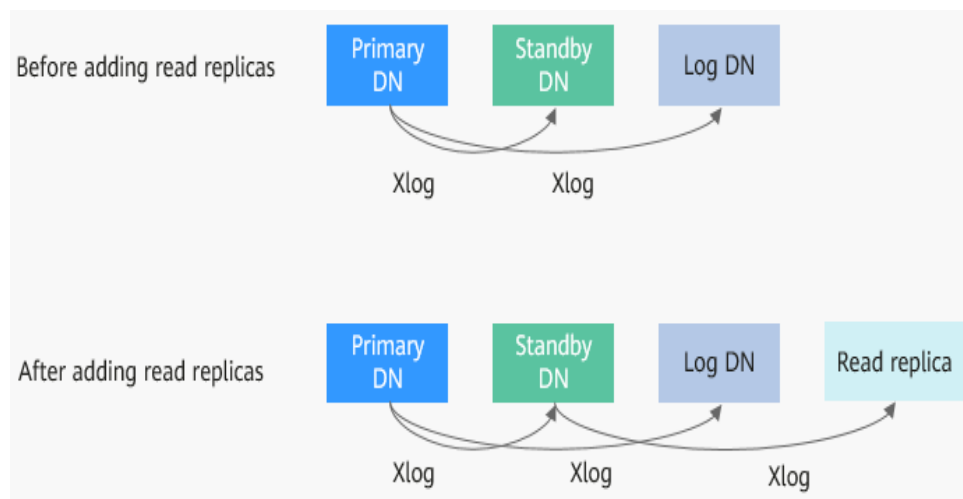
### Scenarios

In a centralized instance, a surge in read requests can consume excessive compute resources. This heavy load may interfere with transaction log (Xlog) replay on standby nodes, leading to increased replication lag between primary and standby nodes. If the primary node fails, the accumulated Xlogs may delay failover. To mitigate this, you can add read replicas to offload and distribute read workloads.

You can also delete read replicas or adjust their specifications to better match changing read demands.

### Technical Details

The primary node handles both read and write operations, while standby and log nodes synchronize Xlogs from it. When read traffic increases, high compute resource usage can hinder Xlog replay on standby nodes. Adding read replicas helps offload read requests. They synchronize Xlogs from standby nodes, reducing the primary node's workload and maintaining stable Xlog replay.

**Figure 4-1** Adding read replicas for a 1 primary + 2 standby instance**Figure 4-2** Adding read replicas for a 1 primary + 1 standby + 1 log instance


## Precautions

- Read replica-related operations are supported for centralized (1 primary + 2 standby) instances of V2.0-2.7.1 or later.
- Adding read replicas temporarily stops backups. Once the replicas are ready, the system re-enables backups and runs a full backup automatically.

## Adding a Read Replica

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the **DB Information** area, click **Add** in the **Read Replicas** field.

**Step 6** Select the specifications and parameter template of the read replica, and click **Add Read Replica**.

**Step 7** Click **Next**, confirm the information, and click **Submit**.


**Step 8** Check the result of adding a read replica.


On the **Instances** page, the instance status is **Creating read replicas**. After the instance status becomes **Available**, click the instance name. On the **Basic Information** page, check the number of read replicas.

----End

## Deleting a Read Replica

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the **DB Information** area, click **Delete** in the **Read Replicas** field.

**Step 6** Select the read replica to be deleted and click **Next**.

**Step 7** Confirm the information and click **Submit**.


**Step 8** Check the result of deleting a read replica.


On the **Instances** page, the instance status is **Deleting a read replica**. After the instance status becomes **Available**, click the instance name. On the **Basic Information** page, check the number of read replicas.

----End

## Changing the Specifications of Read Replicas

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the **DB Information** area, click **Change** in the **Read Replicas** field.

**Step 6** On the displayed page, select the target read replica and new specifications, and click **Next**.

**Step 7** Confirm the specifications and click **Submit**.

**Step 8** Check the specification change result.

After the task is submitted, click **Go to Instance List**. On the **Instances** page, the instance status is **Changing instance specifications**. After a few minutes, check the new instance specifications on the **Basic Information** page.

-----End

## 4.4 Scaling In and Out an Instance

### 4.4.1 Adding Coordinator Nodes for an Instance (Distributed)

#### Scenarios

As the instance deployment time and data increase, the database performance and storage will gradually reach the bottleneck. Adding nodes can improve the instance performance and storage capacity. You can only add nodes for distributed GaussDB instances that are deployed independently.


#### Precautions

- You can flexibly add coordinator nodes (CNs) or shards as needed. It is recommended that the number of CNs of a GaussDB instance do not exceed twice the number of shards.
- Instances can be scaled out only when they are in the **Available** state.

#### Procedure

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

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

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

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

**Step 5** On the **Basic Information** page, click **Add** in the **Coordinator Nodes** field in the **DB Information** area.

**Step 6** Specify the number of coordinator nodes to be added and the AZ.

If single-AZ deployment is specified during the instance creation, CNs are only added to the AZ you specified.

**Step 7** Click **Next**.

**Step 8** Confirm the information and click **Submit**.

**Step 9** Check the result.

On the **Instances** page, check that the instance is in the **Adding CNs** state during the scale-out. After the instance status changes to **Available**, click the instance

name. On the **Basic Information** page, check that the number of CNs matches the required quantity.

----End

## 4.4.2 Adding Shards for an Instance (Distributed)

### Scenarios

As the instance deployment time and data increase, the database performance and storage will gradually reach the bottleneck. In this case, you need to add hosts to improve the instance performance and storage capability. This function is available only for distributed GaussDB instances that are deployed independently.


### Precautions

- Instances can be scaled out only when they are in the **Available** state. When shards are being added, you can still query and insert data, query services are not interrupted, and the data insertion performance is not affected. The performance of join queries on local tables across node groups during redistribution may be affected.

### Procedure

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

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

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

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

**Step 5** On the **Basic Information** page, click **Add** in the **Shards** field in the **DB Information** area.

**Step 6** Specify the number of shards to be added. Click **Next**.

**Step 7** Confirm the information and then click **Submit**.

#### NOTE

By default, a shard contains three replicas (a primary DN and two standby DNs). Each time you add a shard, three replicas will be added.

**Step 8** Check the result.

On the **Instances** page, check that the instance is in the **Adding shards** state during the scale-out. After the instance status changes to **Available**, click the instance name. On the **Basic Information** page, check that the number of shards matches the required quantity.

----End

## 4.4.3 Deleting Coordinator Nodes for an Instance (Distributed)

### Scenarios

As service demand decreases, some coordinator nodes (CNs) are left idle. To improve resource utilization, you can delete unnecessary coordinator nodes. This function is available only for distributed GaussDB instances that are deployed independently.

#### NOTE


This function is only available to authorized users. To apply for the permissions needed, [submit a service ticket](#).


### Precautions

- Deleting CNs does not interrupt ongoing services.
- You can only delete the CNs of instances that were deployed independently.
- At least one CN needs to be reserved for each DB instance.
- Before deleting a CN, ensure that the CN is not in a JDBC connection configuration, or the high availability of the JDBC connection may be affected.
- DDL operations will be rolled back when CNs are being deleted.
- PITR backup is suspended during the deletion and is automatically restored after deletion is complete.
- After the deletion is complete, a full backup is performed automatically.
- Before you delete CNs, you need to ensure that the instance status and all CNs are normal.
- Main processes are running on the main CN (that is, the CN whose component ID is cn\_5001), so this CN cannot be deleted for scale-in. You can call the "Querying the Components of a DB Instance" API to query cn\_5001. If the CN to be deleted is cn\_5001, the system will randomly select another CN to delete.

### Procedure

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

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

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

**Step 4** On the **Instances** page, click the name of the instance for which you want to delete CNs.

**Step 5** In the **DB Information** area of the **Basic Information** page, delete CNs.

1. Click **Delete** next to **Coordinator Nodes**.
2. Select the coordinator nodes to be deleted.

3. Click **Next**.
4. Confirm the information about the CNs to be deleted and click **Submit**.
5. Check the result.

On the **Instances** page, check that the instance is in the **Deleting coordinator nodes** state during the scale-in. After the instance status changes to **Available**, click the instance name. On the **Basic Information** page, check that the number of CNs matches the required quantity.

----End

## 4.5 Scaling Up Storage Space

### 4.5.1 Manually Scaling Up Storage Space for an Instance

#### Scenarios

As more data is added, you may start to run out of space. If the kernel system detects that the storage usage exceeds the predefined threshold, the instance is set to read-only and no data can be written to it. (The default threshold is 85%. You can [modify the instance parameter `cms:datastorage\_threshold\_value\_check`](#) to change the usage threshold.) This section describes how to scale up the storage space of a DB instance. Services will not be interrupted during storage scale-up.

#### Precautions


- Within the maximum allowed range, the storage usage cannot exceed the predefined threshold (85% by default) after the scale-up.
- If any node becomes faulty, [submit a service ticket](#) to contact the O&M engineers for troubleshooting before the scale-up.
- The storage space must be a multiple of (Number of shards x 4 GB).
- If a DR relationship has been established for the instance and the instance functions as the primary one, ensure that the storage of its DR instances is greater than or equal to that of the primary instance during storage scale-up. Otherwise, disaster recovery may fail if the data volume of the primary instance is greater than that of the DR instance.


#### Constraints

- The maximum allowed storage for a single shard is 24 TB by default. There is no limit on the number of scale-ups.
- The DB instance is in the **Scaling up** state when its storage space is being scaled up and the backup services are not affected.
- Do not reboot or delete the instance whose storage is being scaled up.
- Storage space can only be scaled up, not down.

## Procedure

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

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

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

**Step 4** On the **Instances** page, locate the instance you want to scale up and click **More > Scale Storage Space** in the **Operation** column.

Alternatively, click the instance name to go to the **Basic Information** page. In the **Storage/Backup Space** area, click **Scale**.

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

When you scale up the storage space, ensure that the usage of the new storage space is less than the predefined threshold (85% by default). An instance can be restored from read-only to the read/write state only when the storage usage is lower than the specified threshold.

**Step 6** Confirm settings.

- If you need to modify your settings, click **Previous**.
- If the settings are correct, click **Submit**.

**Step 7** View the storage scale-up results.

During the scale-up, the status of the instance on the **Instances** page is **Scaling up**. This process may take 3 to 5 minutes. Once the scale-up is complete, click the instance name to go the **Basic Information** page and you can see the new storage space.

----End

## 4.5.2 Configuring Storage Autoscaling for an Instance

### Scenarios

You can enable autoscaling for a GaussDB instance so that its storage can be automatically scaled up when the disk usage reaches the specified threshold.

### Precautions


- Autoscaling is unavailable for distributed instances using combined deployment.
- All nodes in the target instance must be in an available state.
- Storage autoscaling is mutually exclusive with the following operations: manually scaling up storage space, adding nodes, changing the disk type, deleting an instance, checking snapshots, updating agents, and storage autoscaling. That is, storage autoscaling cannot be performed when any of the preceding operations is ongoing, even if the autoscaling policy is configured.
- Autoscaling for centralized instances is at the instance level.

- Autoscaling for distributed instances is at the shard level.
- If the storage sizes of shards in a distributed instance are different after the scale-up, no shards can be added or deleted until all shards are scaled to the same size.
- During storage autoscaling, the storage space is increased in increments of 40 GB.
  - If the space to increase exceeds the upper limit you have specified, only the space size equal to the upper limit will be increased.
  - If the space to increase exceeds the upper limit specified by the system, only the space size equal to the upper limit will be increased.
- An alarm will be generated when autoscaling fails. This alarm will be automatically cleared when the disk usage is lower than the specified threshold.
- If a yearly/monthly DB instance has pending orders, autoscaling will fail.
- If your account balance is insufficient, autoscaling will fail.

## Procedure

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the **Storage/Backup Space** area, click **Auto-Scale**.

**Step 6** In the displayed **Configure Storage Autoscaling** dialog box, set the following parameters:

**Table 4-1** Parameters

| Parameter                             | Description                                                                                                                                                                                                                    |
|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Autoscaling                           | Specifies whether to enable automatic scale-out. By default, automatic scale-out is disabled.                                                                                                                                  |
| Trigger If Available Storage Drops To | The storage will be automatically scaled up if the available storage drops to or below the threshold specified by this parameter. The default value is <b>20%</b> . The value can be <b>20%</b> , <b>25%</b> , or <b>50%</b> . |

| Parameter         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Autoscaling Limit | <p>Upper limit of the storage space in GB that can be automatically scaled to. The value of this parameter must be greater than the current storage of the instance.</p> <p>Value range:</p> <ul style="list-style-type: none"><li>Centralized instances: [<i>Current storage</i> + 40 GB, 24,000 GB]</li><li>Distributed instances: [<i>Current storage</i> + 40 GB, 24,000 GB x <i>Number of shards</i>]</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Scaling Method    | <p>The value can be <b>By percentage</b> or <b>By fixed amount</b>.</p> <ul style="list-style-type: none"><li>If <b>By percentage</b> is selected, the storage space to be expanded increases each time.</li><li>If <b>By fixed amount</b> is selected, a fixed volume of storage will be expanded each time.</li></ul> <p>Evaluate your workloads and costs and select a method as required.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Scale Up By       | <p>Size of the storage space to be expanded each time, which depends on the selected scaling method. Storage can be scaled by percentage or a fixed amount.</p> <ul style="list-style-type: none"><li>If <b>By percentage</b> is selected for <b>Scaling Method</b>, the default value of this parameter is <b>20%</b>, and the value range is [1%, 100%]. If the available storage drops to or below the specified threshold, the storage will be automatically scaled up by a percentage specified by this parameter (in increments of 40 GB). For example, if the current storage space of an instance is 40 GB and this parameter is set to <b>20%</b>, the volume to be expanded is 8 GB, which will be rounded up to 40 GB.</li><li>If <b>By fixed amount</b> is selected for <b>Scaling Method</b>, the default value of this parameter is <b>40 GB</b>, and the value range is [40 GB, <i>Current storage space</i> + 40 GB]. If the available storage drops to or below the specified threshold, the storage will be automatically scaled up by a fixed amount specified by this parameter.</li></ul> |

**Step 7** Click **OK**.

----End

## 4.6 Changing the Deployment Model

### 4.6.1 Changing Standby DNs to Log Nodes (for a Distributed Instance)

#### Scenarios


Each shard of a distributed GaussDB instance consists of one primary DN and two standby DNs. However, this architecture may not meet your application requirements. GaussDB provides the flexibility to convert one standby DN in each shard into a log node.


#### Precautions

- This function is only available for distributed instances of V2.0-3.200.0 or later that have one primary DN and two standby DNs in each shard.
- PITR backup is suspended during the process of changing standby DNs to log nodes and automatically resumes after the operation is complete.
- After standby DNs are changed to log nodes, a full backup is automatically performed.

#### Procedure

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the **Node List** area, click **Change to Log Nodes**.

**Step 6** In the **Change Standby Data Nodes to Log Nodes** dialog box, select an AZ and click **OK**.

**Step 7** Check the change result.

After the task is submitted, click **Back to DB Instance List**. On the **Instances** page, the instance status is **Changing to log node**. After the task is complete, go to the **Basic Information** page of the instance and check that the deployment model is changed to 1 primary + 1 standby + 1 log.

----End

## 4.6.2 Changing Log Nodes to Standby DNs (for a Distributed Instance)

### Scenarios


Each shard of a distributed GaussDB instance consists of one primary DN, one standby DN, and one log node. However, this architecture may not meet your application requirements. GaussDB provides the flexibility to convert the log node in each shard into a standby DN.


### Constraints

- This function is only available for distributed instances of V2.0-3.200.0 or later that have one primary DN, one standby DN, and one log node in each shard.
- The statuses of the instance, cluster, and DNs must all be normal. If there is an error, rectify the fault before performing this operation.
- After the deployment model of a yearly/monthly distributed instance is changed to 1 primary + 2 standby, you need to pay for the order in the order center.

### Procedure

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the **Node List** area, click **Change to 1 Primary + 2 Standby**.

**Step 6** In the **Change Log Nodes to Standby Data Nodes** dialog box, click **OK**.

**Step 7** Check the change result.

After the task is submitted, click **Back to DB Instance List**. On the **Instances** page, the instance status is **Changing to standby DN**. After the task is complete, go to the **Basic Information** page of the instance and check that the deployment model is changed to 1 primary + 2 standby.

-----End

## 4.7 Performing a Primary/Standby DN Switchover

### 4.7.1 Changing the DN Failover Priority

#### Scenarios

GaussDB provides failover priority on availability or reliability. You can change the failover priority of a GaussDB instance on the **Basic Information** page. Reliability

applies to scenarios that require high data consistency, and availability applies to scenarios that require uninterrupted online services.

#### NOTE

This function is only available to authorized users. To apply for the permissions needed, [submit a service ticket](#).


## Precautions

- This function is only available to distributed instances.
- If you select **Availability** for the failover priority, exercise caution when changing the following database parameters:
  - **recovery\_time\_target**: Specifies the time for the standby node to write and replay logs. The value ranges from 0 to 3600, in seconds. The default value is **60**. **0** indicates that log flow control is disabled. A value from 1 to 3600 indicates that a standby node can write and replay logs within the period specified by this parameter, so that the standby node can quickly assume the primary role. If **recovery\_time\_target** is set to a small value, the performance of the primary node is affected. If it is set to a large value, the log flow is not effectively controlled. You are advised to retain the default value.
  - **audit\_system\_object**: Specifies whether to audit the CREATE, DROP, and ALTER operations on GaussDB database objects. GaussDB database objects include databases, users, schemas, and tables. The value of this parameter ranges from 0 to 536,870,911. The default value is **67121159**. You can change the value of this parameter to audit only the operations on required database objects. In the scenario where the leader node is forcibly selected, you are advised to set **audit\_system\_object** to the maximum value and audit all DDL objects.

## Procedure

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

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

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

**Step 4** Click the name of the target instance to go to the **Basic Information** page.

**Step 5** Click **Change** in the **Failover Priority** field of the **DB Information** area.

**Step 6** In the displayed dialog box, select **Reliability** or **Availability** as required.

- **Reliability**: Data consistency is given priority during a failover. This is recommended for applications with highest priority for data consistency.
- **Availability**: Database availability is given priority during a failover. This is recommended for applications that require their databases to provide uninterrupted online services.

**Step 7** Click **OK**.

**Step 8** After you change some parameters, manually reboot the instance for the changes to take effect. For details, see [Rebooting a GaussDB Instance](#).

The failover priority cannot be changed when the DB instance is in the **Rebooting** state.

----End

## 4.7.2 Performing a Primary/Standby Switchover

### Scenarios

GaussDB supports primary/standby DN switchover in a shard of an instance when the instance is available. You can promote a standby DN to the primary DN in a shard.

### Constraints

- This operation cannot be performed when the node status is abnormal.
- Only one standby node can be specified as the primary node in a shard.
- Single-node instances do not support primary/standby DN shard switchovers.
- During a primary/standby switchover, the following operations cannot be performed:
  - Rebooting a DB instance
  - Switching AZs
  - Changing CPU and memory specifications of an instance
  - Repairing a node
  - Replacing a node
  - Adding nodes
  - Backing up and restoring an instance


### Precautions

- Services may be interrupted for several seconds or minutes during the switchover. You are advised to perform this operation during off-peak hours.
- Switchovers will not change private IP addresses of an instance.

### Procedure

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

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

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

**Step 4** Click the target instance name to go to the **Basic Information** page.

**Step 5** In the **Node List** area, click **Switch Primary and Standby DNs**.

**Step 6** Select an AZ to view the DN shard of the primary DN in the selected AZ. Select the standby DN to be promoted to primary, enter **rearrange**, and click **OK**.

- If there is no primary DN in the selected AZ, shard information is not displayed.
- Services may be interrupted for several seconds or minutes during the switchover. You are advised to perform this operation during off-peak hours.
- Primary/Standby switchovers can be performed on a maximum of 30 shards at a time.

----End

# 5 Instance Upgrade

## 5.1 Overview

You can manually upgrade the GaussDB kernel version of a single instance or multiple instances in batches using in-place upgrade, gray upgrade, or hot patch update to improve performance, add new functions, and fix bugs.

### Checking the Current Kernel Version

To check the version of an instance, go to the **Basic Information** page of the instance and check the value of **DB Engine Version** in the **DB Information** area.

### Upgrade Methods

The following table describes the upgrade methods supported by GaussDB.

**Table 5-1** Upgrade methods

| Upgrade Method            | Action      | Type           | Scenario            | Rollback Method                                                          | Impact on Services                            | Suggestions |
|---------------------------|-------------|----------------|---------------------|--------------------------------------------------------------------------|-----------------------------------------------|-------------|
| <a href="#">Hot patch</a> | Auto-commit | Online upgrade | Fix product issues. | <ul style="list-style-type: none"><li>Automatic</li><li>Manual</li></ul> | No service is interrupted during the upgrade. | None        |

| Upgrade Method   | Action | Type            | Scenario                                                                                           | Rollback Method | Impact on Services                                                         | Suggestions                            |
|------------------|--------|-----------------|----------------------------------------------------------------------------------------------------|-----------------|----------------------------------------------------------------------------|----------------------------------------|
| In-place upgrade | N/A    | Offline upgrade | <ul style="list-style-type: none"><li>• Add new functions.</li><li>• Fix product issues.</li></ul> | Automatic       | Services are interrupted for about 30 minutes during the in-place upgrade. | Stop all workloads during the upgrade. |

| Upgrade Method | Action          | Type           | Scenario                                                                                           | Rollback Method                                                              | Impact on Services                                                                                                                                                                                                                                                                     | Suggestions                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------|-----------------|----------------|----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Gray upgrade   | Auto-commit     | Online upgrade | <ul style="list-style-type: none"><li>• Add new functions.</li><li>• Fix product issues.</li></ul> | Automatic                                                                    | Services are interrupted for about 10s during the upgrade of primary DN and during the upgrade of CNs. During upgrade commit, primary/standby distribution balancing may be performed. Services may be interrupted for different periods of time based on factors such as data volume. | Add the service reconnection mechanism. It is recommended that the retry interval be 1s. During low-pressure periods (less than 3,000 TPS + 4,000 QPS for each shard), the total retry duration is 25s. During high-pressure periods (less than 6,000 TPS + 10,000 QPS for each shard), the total retry duration is 100s. The upgrade is not recommended when the pressure is out of the acceptable range. |
|                | Rolling upgrade | Online upgrade | <ul style="list-style-type: none"><li>• Add new functions.</li><li>• Fix product issues.</li></ul> | <ul style="list-style-type: none"><li>• Automatic</li><li>• Manual</li></ul> | If the AZ to be upgraded contains primary DNs, services will be interrupted for about 10s during the upgrade of each primary DN. If the AZ to be upgraded contains CNs, services will be interrupted for about 10s during the upgrade of each CN.                                      |                                                                                                                                                                                                                                                                                                                                                                                                            |

## 5.2 Hot Patch Update

### Scenarios

You can install a hot patch for your GaussDB instance to rectify product issues. A hot patch can be loaded without interrupting services and can be used to resolve some emergent database kernel problems online without affecting services. Hot patch update supports manual rollback.

### Precautions

- During the update, hot patch packages will be downloaded and decompressed, which occupies certain disk space. It is recommended that the disk usage on the DN be less than or equal to the disk usage threshold minus 10%.

#### NOTE

To check the current DN disk usage, go to the metric monitoring page on the management console.




To obtain the disk usage threshold, [submit a service ticket](#) to contact technical support.

- Version upgrade is unavailable if instance nodes are in an abnormal state.
- If a hot patch conflicts with the backup, the differential backup and full backup of the instance will be stopped during hot patch installation.
- During an upgrade or rollback, the following operations cannot be performed: scaling up storage, changing specifications, backing up data, resetting passwords, rebooting instances, and deleting instances.
- You are advised to perform an upgrade during off-peak hours because there are more idle CPU, disk, and memory resources.
- Hot patch update is available only when there is a hot patch for installation. If no hot patch is available, hot patch update is not displayed.
- Hot patch update and rollback can be performed in batches for different patch versions of a single instance. During the update, hot patches are installed in ascending order of version numbers. During the rollback, hot patches are rolled back in descending order of version numbers.
- If the upgrade fails, the system automatically rolls back the instance to the source version. You can [submit a service ticket](#) to contact technical support, and the engineers will help you upgrade the instance if necessary.
- After the upgrade is complete, you can manually roll back the upgrade.
- A maximum of 30 instances can be selected at a time for batch upgrade.

### Step 1: Perform a Pre-upgrade Check


Before an upgrade, check the instance status and whether monitoring metrics such as the CPU usage, memory usage, and disk usage of the instance are normal.

1. Check instance status.

- a. Log in to the management console.
  - b. Click  in the upper left corner and select a region and project.
  - c. Click  in the upper left corner of the page and choose **Databases > GaussDB**.
  - d. On the **Instances** page, check whether the target instance is available.  
If the instance is in an abnormal state, [submit a service ticket](#) to contact technical support, and the engineers will help you analyze and handle the abnormal instance.
2. Check monitoring metrics.
- a. Click  in the upper left corner of the page, and choose **Management & Governance > Cloud Eye**.
  - b. In the navigation pane, choose **Cloud Service Monitoring > GaussDB**.
  - c. On the **Cloud Service Monitoring** page, click the target instance to go to the metric monitoring page.
    - On the **DB Instance** tab, view the value of **Instance Disk Usage** to check whether the disk usage is insufficient.
    - On the **Node** tab, view the value of **CPU Usage** to check whether the CPU usage remains high for a long time.
    - On the **Node** tab, view the value of **Memory Usage** to check whether the memory usage increases sharply.If any metric is abnormal, [submit a service ticket](#) to contact technical support, and the engineers will help you analyze and handle the abnormal metrics.

## Step 2: Perform the Upgrade

### [Method 1: upgrading a single instance]


1. Click  in the upper left corner of the page and choose **Databases > GaussDB**.
2. On the **Instances** page, click **More** in the **Operation** column of the target instance and choose **Upgrade**.
3. In the **Upgrade Instance** dialog box, select **Hot patch** for **Upgrade Method**, enter **confirm**, and click **OK**.

#### NOTE

All available patch versions are displayed in the **Target Version** area. If multiple patches are to be installed, they will be installed in ascending order of version numbers after the upgrade task is submitted.

4. Check the upgrade result on the **Instances** page.
  - During the upgrade, the instance status is **Upgrading version**.
  - After the upgrade is complete, the instance status changes to **Available**. Go to the **Basic Information** page of the instance and check whether **DB Engine Version** is the target version in the **DB Information** area.

**[Method 2: upgrading instances in batches]**

1. Click  in the upper left corner of the page and choose **Databases > GaussDB**.
2. On the **Instances** page, select the target instances and click **Batch Upgrade**.
3. In the **Batch Upgrade** dialog box, select **Hot patch** for **Upgrade Method**.
4. Enter **confirm** and click **OK**.

 **NOTE**

All available patch versions are displayed in the **Target Version** area. If multiple patches are to be installed, they will be installed in ascending order of version numbers after the upgrade task is submitted.

5. Check the upgrade result on the **Instances** page.
  - During the upgrade, the instance status is **Upgrading version**.
  - After the upgrade is complete, the instance status changes to **Available**. Go to the **Basic Information** page of the instance and check whether **DB Engine Version** is the target version in the **DB Information** area.

**Step 3: Verify the Upgrade**

After the upgrade is complete, check the instance status, backup creation status, and instance connectivity, and whether you can add, delete, update, and query data in the instance.

1. On the **Instances** page, check whether **Status** of the target instance is **Available**.
2. On the **Instances** page, click the name of the target instance. On the **Basic Information** page that is displayed, check that the target versions are displayed in the **Upgraded Hot Patch Version** and **Upgraded Kernel Hot Patch Version** fields in the **DB Information** area.
3. Check that the instance is properly connected and you can add, delete, update, and query data in the instance.
  - a. Log in to the database. For details, see [Connecting to an Instance Through DAS](#).
  - b. Go to the **SQL Query** page.
  - c. Create a database.

```
CREATE DATABASE database_name;
```

In this example, run the following command to create a database named **db\_tpcds**:

```
CREATE DATABASE db_tpcds;
```

Switch to the newly created database in the upper left corner.
  - d. Create a table and add, delete, update, and query data in the table.
    - i. Create a schema.

```
CREATE SCHEMA myschema;
```
    - ii. Create a table named **mytable** that has only one column. The column name is **firstcol** and the column type is integer.

- ```
CREATE TABLE myschema.mytable (firstcol int);
```
- iii. Insert data into the table.
- ```
INSERT INTO myschema.mytable VALUES (100);
```
- iv. View data in the table.
- ```
SELECT * FROM myschema.mytable;
```
- | | firstcol |
|---|----------|
| 1 | 100 |
- v. Update data in the table.
- ```
UPDATE myschema.mytable SET firstcol = 200;
```
- vi. View the data in the table again.
- ```
SELECT * FROM myschema.mytable;
```
- | | firstcol |
|---|----------|
| 1 | 200 |
- vii. Drop the table.
- ```
DROP TABLE myschema.mytable;
```

## Rollback Operations

If a rollback is required after the upgrade, perform the following operations to roll back an instance to the source version.

### [Method 1: Rolling Back a Single Instance]

- Step 1** In the **Upgrade Instance** dialog box, select **Hot patch** for **Upgrade Method**.
- Step 2** Select **Rollback** for **Action**.
- Step 3** Select the target version, enter **confirm**, and click **OK**.
- Step 4** On the **Instances** page, check the rollback status. After the rollback is complete, the instance status changes to **Available**.
- Step 5** On the **Instances** page, click the name of the target instance. On the **Basic Information** page that is displayed, check that the target versions are not displayed in the **Upgraded Hot Patch Version** and **Upgraded Kernel Hot Patch Version** fields in the **DB Information** area.
- Step 6** Check that the instance is properly connected and you can add, delete, update, and query data in the instance.
1. Log in to the database. For details, see [Connecting to an Instance Through DAS](#).
  2. Go to the **SQL Query** page.
  3. Create a database.  

```
CREATE DATABASE database_name;
```

In this example, run the following command to create a database named **db\_tpcds**:

```
CREATE DATABASE db_tpcds;
```

Switch to the newly created database in the upper left corner.

4. Create a table and add, delete, update, and query data in the table.
  - a. Create a schema.  
`CREATE SCHEMA myschema;`
  - b. Create a table named **mytable** that has only one column. The column name is **firstcol** and the column type is integer.  
`CREATE TABLE myschema.mytable (firstcol int);`
  - c. Insert data into the table.  
`INSERT INTO myschema.mytable values (100);`
  - d. View the data in the table.  
`SELECT * FROM myschema.mytable;`

| firstcol |
|----------|
| 100      |
  - e. Update data in the table.  
`UPDATE myschema.mytable SET firstcol = 200;`
  - f. View the data in the table again.  
`SELECT * FROM myschema.mytable;`

| firstcol |
|----------|
| 200      |
  - g. Drop the table.  
`DROP TABLE myschema.mytable;`

----End

## [Method 2: Rolling Back Instances in Batches]

- Step 1** On the **Instances** page, select the target instances and click **Batch Upgrade**.
- Step 2** In the **Batch Upgrade** dialog box, select **Hot patch** for **Upgrade Method** and **Rollback** for **Action**, select a target version, enter **confirm**, and click **OK**.
- Step 3** On the **Instances** page, check the rollback status. After the rollback is complete, the instance status changes to **Available**.
- Step 4** On the **Instances** page, click the name of the target instance. On the **Basic Information** page that is displayed, check that the target versions are not displayed in the **Upgraded Hot Patch Version** and **Upgraded Kernel Hot Patch Version** fields in the **DB Information** area.
- Step 5** Check that the instance is properly connected and you can add, delete, update, and query data in the instance.
  1. Log in to the database. For details, see [Connecting to an Instance Through DAS](#).
  2. Go to the **SQL Query** page.
  3. Create a database.  
`CREATE DATABASE database_name;`  
In this example, run the following command to create a database named **db\_tpcds**:

**CREATE DATABASE db\_tpcds;**

Switch to the newly created database in the upper left corner.

4. Create a table and add, delete, update, and query data in the table.

- a. Create a schema.

```
CREATE SCHEMA myschema;
```

- b. Create a table named **mytable** that has only one column. The column name is **firstcol** and the column type is integer.

```
CREATE TABLE myschema.mytable (firstcol int);
```

- c. Insert data into the table.

```
INSERT INTO myschema.mytable VALUES (100);
```

- d. View data in the table.

```
SELECT * FROM myschema.mytable;
```

```
| firstcol |
-----+
1 | 100 |
```

- e. Update data in the table.

```
UPDATE myschema.mytable SET firstcol = 200;
```

- f. View the data in the table again.

```
SELECT * FROM myschema.mytable;
```

```
| firstcol |
-----+
1 | 200 |
```

- g. Drop the table.

```
DROP TABLE myschema.mytable;
```

----End

## 5.3 In-place Upgrade

### Scenarios

You can use in-place upgrade to upgrade your instance if a version upgrade is required for new functions or issue rectification. During an in-place upgrade, all nodes are upgraded at the same time, and all services are interrupted.

### Precautions

- The DN disk usage cannot be greater than the configured disk usage threshold minus 10%.

 **NOTE**

To check the current DN disk usage, go to the metric monitoring page on the management console.

To obtain the disk usage threshold, [submit a service ticket](#) to contact technical support.

- Version upgrade is unavailable if instance nodes are in an abnormal state.
- Before the upgrade, ensure that there is a public schema in each database. Otherwise, the upgrade will fail.

- During an upgrade, the following operations cannot be performed: scaling up storage, changing specifications, backing up data, resetting passwords, rebooting instances, and deleting instances.
- If this method is used for a major version upgrade, log archiving will be disabled before the upgrade, and you cannot use archive logs for Point-In-Time Recovery (PITR), which may result in data loss.

 **NOTE**

Example major version upgrade: from V2.0-1.x to V2.0-2.x or from V2.0-2.x to V2.0-2.y

- If the upgrade fails, the system automatically rolls back the instance to the source version. You can [submit a service ticket](#) to contact customer service, and the engineers will help you upgrade the instance if necessary.
- Services are interrupted for about 30 minutes during the in-place upgrade.
- After the upgrade is complete, an automated backup task will be triggered and log archiving will be enabled. However, for a single-replica instance upgraded to V2.0-3.0 or later from earlier versions, automated backup is disabled by default and will not be triggered. An automated backup task will also not be triggered in the case of minor version upgrades.



 **NOTE**


Example minor version upgrade: from V2.0-1.a.x to V2.0-1.a.y or from V2.0-2.a.x to V2.0-2.a.y

- In-place upgrade does not require manual rollback.

## Step 1: Perform a Pre-upgrade Check

Before an upgrade, check the instance status and whether monitoring metrics such as the CPU usage, memory usage, and disk usage of the instance are normal.

1. Check instance status.
  - a. Log in to the management console.
  - b. Click  in the upper left corner and select a region and project.
  - c. Click  in the upper left corner of the page and choose **Databases > GaussDB**.
  - d. On the **Instances** page, check whether the target instance is available.


If the instance is in an abnormal state, [submit a service ticket](#) to contact technical support, and the engineers will help you analyze and handle the abnormal instance.
2. Check monitoring metrics.
  - a. Click  in the upper left corner of the page, and choose **Management & Governance > Cloud Eye**.
  - b. In the navigation pane, choose **Cloud Service Monitoring > GaussDB**.
  - c. On the **Cloud Service Monitoring** page, click the target instance to go to the metric monitoring page.
    - On the **DB Instance** tab, view the value of **Instance Disk Usage** to check whether the disk usage is insufficient.

- On the **Node** tab, view the value of **CPU Usage** to check whether the CPU usage remains high for a long time.
- On the **Node** tab, view the value of **Memory Usage** to check whether the memory usage increases sharply.


If any metric is abnormal, [submit a service ticket](#) to contact technical support, and the engineers will help you analyze and handle the abnormal metrics.

## Step 2: Perform the Upgrade

### [Method 1: upgrading a single instance]

1. Click  in the upper left corner of the page and choose **Databases > GaussDB**.
2. On the **Instances** page, click **More** in the **Operation** column of the target instance and choose **Upgrade**.  
  
Alternatively, click the target instance name to go to the **Basic Information** page. In the **DB Information** area, click **Upgrade Instance** in the **DB Engine Version** field.
3. In the **Upgrade Instance** dialog box, select **In-place** for **Upgrade Method**, select the target version, enter **confirm**, and click **OK**.
4. Check the upgrade result on the **Instances** page.
  - During the upgrade, the instance status is **Upgrading version**.
  - After the upgrade is complete, the instance status changes to **Available**. Go to the **Basic Information** page of the instance and check whether **DB Engine Version** is the target version in the **DB Information** area.

### [Method 2: upgrading instances in batches]

1. Click  in the upper left corner of the page and choose **Databases > GaussDB**.
2. On the **Instances** page, select the target instances and click **Batch Upgrade**.
3. In the **Batch Upgrade** dialog box, select **In-place** for **Upgrade Method**, select the target version, enter **confirm**, and click **OK**.
4. Check the upgrade result on the **Instances** page.
  - During the upgrade, the instance status is **Upgrading version**.
  - After the upgrade is complete, the instance status changes to **Available**. Go to the **Basic Information** page of the instance and check whether **DB Engine Version** is the target version in the **DB Information** area.

## Step 3: Verify the Upgrade

After the upgrade is complete, check the instance status, backup creation status, and instance connectivity, and whether you can add, delete, update, and query data in the instance.

1. On the **Instances** page, check whether **Status** of the target instance is **Available**.
2. On the **Instances** page, click the name of the target instance. On the **Basic Information** page that is displayed, check whether the value of **DB Engine Version** in the **DB Information** area is the target version.
3. Check that the automated backup triggered after the upgrade is successfully created.
  - a. On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.
  - b. In the navigation pane, choose **Backups**. Check that a backup has been created and the backup status is **Completed**.
4. Check that the instance is properly connected and you can add, delete, update, and query data in the instance.

- a. Log in to the database. For details, see [Connecting to an Instance Through DAS](#).

- b. Go to the **SQL Query** page.

- c. Create a database.

```
CREATE DATABASE database_name;
```

In this example, run the following command to create a database named **db\_tpcds**:

```
CREATE DATABASE db_tpcds;
```

Switch to the newly created database in the upper left corner.

- d. Create a table and add, delete, update, and query data in the table.

- i. Create a schema.

```
CREATE SCHEMA myschema;
```

- ii. Create a table named **mytable** that has only one column. The column name is **firstcol** and the column type is integer.

```
CREATE TABLE myschema.mytable (firstcol int);
```

- iii. Insert data into the table.

```
INSERT INTO myschema.mytable VALUES (100);
```

- iv. View data in the table.

```
SELECT * FROM myschema.mytable;
```

```
| firstcol |
-----+
1 | 100 |
```

- v. Update data in the table.

```
UPDATE myschema.mytable SET firstcol = 200;
```

- vi. View the data in the table again.

```
SELECT * FROM myschema.mytable;
```

```
| firstcol |
-----+
1 | 200 |
```

- vii. Drop the table.

```
DROP TABLE myschema.mytable;
```

## 5.4 Gray Upgrade

### Scenarios

You can use gray upgrade to upgrade your GaussDB instance if a version upgrade is required for new functions or issue rectification. You can either select auto-commit after the upgrade or perform a rolling upgrade.

- In the auto-commit mode, all standby DN nodes are upgraded first and then primary DN nodes and CN nodes in sequence. After the upgrade is complete, the upgrade is automatically committed.
- The rolling upgrade mode is also known as the upgrade observation mode. In this mode, the system enters the observation state after the upgrade is complete. During this period, you can observe the service status and either commit or roll back the upgrade based on service status.
  - Distributed instances are upgraded by shard. For details, see [Upgrading a Distributed Instance](#).
  - Centralized instances are upgraded by AZ. For details, see [Upgrading a Centralized Instance](#).

### Procedure

| Step                                                | Description                                                                                                                                                                           |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">Step 1: Perform a Pre-upgrade Check</a> | Before an upgrade, check the instance status and whether monitoring metrics such as the CPU usage, memory usage, and disk usage of the instance are normal.                           |
| <a href="#">Step 2: Perform the Upgrade</a>         | Select either auto-commit after the upgrade or perform a rolling upgrade. You can upgrade a single instance or multiple instances in batches as required.                             |
| <a href="#">Step 3: Verify the Upgrade</a>          | After the upgrade is complete, check the instance status, backup creation status, and instance connectivity, and whether you can add, delete, update, and query data in the instance. |

### Precautions

- The DN disk usage cannot be greater than the configured disk usage threshold minus 10%.

#### NOTE

To check the current DN disk usage, go to the metric monitoring page on the management console.

To obtain the disk usage threshold, [submit a service ticket](#) to contact technical support.

- Version upgrade is unavailable if instance nodes are in an abnormal state.

- Before the upgrade, ensure that there is a public schema in each database. Otherwise, the upgrade will fail.
- The rolling upgrade mode supports manual rollback, but the auto-commit mode does not support manual rollback.
- During an upgrade or rollback, the following operations cannot be performed: scaling up storage, changing specifications, backing up data, resetting passwords, rebooting instances, and deleting instances.
- You are advised to perform an upgrade during off-peak hours because there are more idle CPU, disk, and memory resources.
- If upgrade auto-commit is used for a major version upgrade, log archiving will be disabled before the upgrade, and you cannot use archive logs for PITR, which may result in data loss.
- If rolling upgrade is used for a major version upgrade, full backup cannot be triggered during the upgrade, and differential backup may fail. Manual full backups cannot be created until the upgrade operations in all AZs are complete during the rolling upgrade and observation period. Archive logs are still generated before the upgrade is committed, and you can use archive logs for PITR to prevent data loss. In the upgrade commit phase, log archiving is disabled.
- If the upgrade fails, the system automatically rolls back the instance to the source version. You can [submit a service ticket](#) to contact technical support, and the engineers will help you upgrade the instance if necessary.
- Services are interrupted for about 10 seconds during the upgrade of primary DN and during the upgrade of CNs.
- After the upgrade is complete, an automated backup task will be triggered and log archiving will be enabled. However, an automated backup task will not be triggered in the case of minor version upgrades. Log archiving is only available for instances of versions later than V2.0-2.2.



 **NOTE**

Example minor version upgrade: from V2.0-1.a.x to V2.0-1.a.y or from V2.0-2.a.x to V2.0-2.a.y


Example major version upgrade: from V2.0-1.x to V2.0-2.x or from V2.0-2.x to V2.0-2.y

## Step 1: Perform a Pre-upgrade Check

Before an upgrade, check the instance status and whether monitoring metrics such as the CPU usage, memory usage, and disk usage of the instance are normal.

1. Check instance status.
  - a. Log in to the management console.
  - b. Click  in the upper left corner and select a region and project.
  - c. Click  in the upper left corner of the page and choose **Databases > GaussDB**.
  - d. On the **Instances** page, check whether the target instance is available.

If the instance is in an abnormal state, [submit a service ticket](#) to contact technical support, and the engineers will help you analyze and handle the abnormal instance.


2. Check monitoring metrics.
    - a. Click  in the upper left corner of the page, and choose **Management & Governance > Cloud Eye**.
    - b. In the navigation pane, choose **Cloud Service Monitoring > GaussDB**.
    - c. On the **Cloud Service Monitoring** page, click the target instance to go to the metric monitoring page.
      - On the **DB Instance** tab, view the value of **Instance Disk Usage** to check whether the disk usage is insufficient.
      - On the **Node** tab, view the value of **CPU Usage** to check whether the CPU usage remains high for a long time.
      - On the **Node** tab, view the value of **Memory Usage** to check whether the memory usage increases sharply.
- If any metric is abnormal, [submit a service ticket](#) to contact technical support, and the engineers will help you analyze and handle the abnormal metrics.

## Step 2: Perform the Upgrade


You can select auto-commit after the upgrade or perform a rolling upgrade for gray upgrade as required.

### Upgrade Auto-commit

#### [Method 1: upgrading a single instance]

1. Click  in the upper left corner of the page and choose **Databases > GaussDB**.
2. On the **Instances** page, click **More** in the **Operation** column of the target instance and choose **Upgrade**.
3. In the **Upgrade Instance** dialog box, select **Gray upgrade** for **Upgrade Method**.
4. Select **Auto-commit** for **Action**.
5. Select the target version, enter **confirm**, and click **OK**.
6. Check the upgrade result on the **Instances** page.
  - During the upgrade, the instance status is **Upgrading version**.
  - After the upgrade is complete, the instance status changes to **Available**. Go to the **Basic Information** page of the instance and check whether **DB Engine Version** is the target version in the **DB Information** area.


#### [Method 2: upgrading instances in batches]

1. Click  in the upper left corner of the page and choose **Databases > GaussDB**.
2. On the **Instances** page, select the target instances and click **Batch Upgrade**.


3. In the **Batch Upgrade** dialog box, select **Gray upgrade** for **Upgrade Method**.
4. Select **Auto-commit** for **Action**.
5. Select the target version, enter **confirm**, and click **OK**.
6. Check the upgrade result on the **Instances** page.
  - During the upgrade, the instance status is **Upgrading version**.
  - After the upgrade is complete, the instance status changes to **Available**. Go to the **Basic Information** page of the instance and check whether **DB Engine Version** is the target version in the **DB Information** area.

## Rolling Upgrade

### [Method 1: upgrading a single instance]

- Upgrading a distributed instance
  - a. Click  in the upper left corner of the page and choose **Databases > GaussDB**.
  - b. On the **Instances** page, click **More** in the **Operation** column of the target instance and choose **Upgrade**.

Alternatively, click the target instance name to go to the **Basic Information** page. In the **DB Information** area, click **Upgrade Instance** in the **DB Engine Version** field.
  - c. In the **Upgrade Instance** dialog box, select **Gray upgrade** for **Upgrade Method**.
  - d. Select **Rolling upgrade** for **Action**.
  - e. Set **Shards to Upgrade**, select a target version, enter **confirm**, and click **OK**.
  - f. Check the upgrade result on the **Instances** page.
    - i. During the upgrade, the instance status is **Upgrading version**.
    - ii. After the upgrade is complete, the instance status changes to **Observing version upgrade**.
  - g. Check that all shards are upgraded and services are running properly before committing the upgrade.

In the **Upgrade Instance** dialog box, select **Commit** for **Action**, select a target version, enter **confirm**, and click **OK**.
- Upgrading a centralized instance
  - a. Click  in the upper left corner of the page and choose **Databases > GaussDB**.
  - b. On the **Instances** page, click **More** in the **Operation** column of the target instance and choose **Upgrade**.

Alternatively, click the target instance name to go to the **Basic Information** page. In the **DB Information** area, click **Upgrade Instance** in the **DB Engine Version** field.
  - c. In the **Upgrade Instance** dialog box, select **Gray upgrade** for **Upgrade Method**.

- d. Select **Rolling upgrade** for **Action**.
- e. Set **AZs to Upgrade**, select a target version, enter **confirm**, and click **OK**.

 **NOTE**


You can upgrade a single AZ or multiple AZs at a time as needed.

- f. Check the upgrade result on the **Instances** page.
  - i. During the upgrade, the instance status is **Upgrading version**.
  - ii. After the upgrade is complete, the instance status changes to **Observing version upgrade**.
- g. Check that all AZs are upgraded and services are running properly before committing the upgrade.

In the **Upgrade Instance** dialog box, select **Commit** for **Action**, select a target version, enter **confirm**, and click **OK**.

If you choose to upgrade AZs one by one, repeat **b** to **f** until all AZs are upgraded, and then commit the upgrade.

**[Method 2: upgrading instances in batches]**

1. Click  in the upper left corner of the page and choose **Databases > GaussDB**.
2. On the **Instances** page, select the target instances and click **Batch Upgrade**.
3. In the **Batch Upgrade** dialog box, select **Gray upgrade** for **Upgrade Method**.
4. Select **Rolling upgrade** for **Action**.
5. Select the target version, enter **confirm**, and click **OK**.

 **NOTE**

During a batch upgrade in rolling upgrade mode, all AZs or shards of the selected instances are upgraded by default.

6. Check the upgrade result on the **Instances** page.
  - During the upgrade, the instance status is **Upgrading version**.
  - After the upgrade is complete, the instance status changes to **Observing version upgrade**.
7. Check that all shards or AZs are upgraded and services are running properly before committing the upgrade.

In the **Batch Upgrade** dialog box, select **Commit** for **Action**, select a target version, enter **confirm**, and click **OK**.

### Step 3: Verify the Upgrade

After the upgrade is complete, check the instance status, backup creation status, and instance connectivity, and whether you can add, delete, update, and query data in the instance.

1. On the **Instances** page, check whether **Status** of the target instance is **Available**.

2. On the **Instances** page, click the name of the target instance. On the **Basic Information** page that is displayed, check whether the value of **DB Engine Version** in the **DB Information** area is the target version.
3. Check that the automated backup triggered after the upgrade is successfully created.
  - a. On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.
  - b. In the navigation pane, choose **Backups**. Check that a backup has been created and the backup status is **Completed**.
4. Check that the instance is properly connected and you can add, delete, update, and query data in the instance.
  - a. Log in to the database. For details, see [Connecting to an Instance Through DAS](#).
  - b. Go to the **SQL Query** page.
  - c. Create a database.

```
CREATE DATABASE database_name;
```

In this example, run the following command to create a database named **db\_tpcds**:

```
CREATE DATABASE db_tpcds;
```

Switch to the newly created database in the upper left corner.
  - d. Create a table and add, delete, update, and query data in the table.
    - i. Create a schema.

```
CREATE SCHEMA myschema;
```
    - ii. Create a table named **mytable** that has only one column. The column name is **firstcol** and the column type is integer.

```
CREATE TABLE myschema.mytable (firstcol int);
```
    - iii. Insert data into the table.

```
INSERT INTO myschema.mytable VALUES (100);
```
    - iv. View data in the table.

```
SELECT * FROM myschema.mytable;
```

|   | firstcol |
|---|----------|
| 1 | 100      |
    - v. Update data in the table.

```
UPDATE myschema.mytable SET firstcol = 200;
```
    - vi. View the data in the table again.

```
SELECT * FROM myschema.mytable;
```

|   | firstcol |
|---|----------|
| 1 | 200      |
    - vii. Drop the table.

```
DROP TABLE myschema.mytable;
```

## Rollback Operations

During upgrade observation, if a rollback is required due to service reasons or the upgrade using the rolling upgrade mode fails, you can manually roll back the upgrade by performing the steps in this section.

### NOTE

- If the rollback is successful, you can perform the upgrade again.
- If the rollback fails, you can perform the rollback again.

If the problem persists, [submit a service ticket](#) to contact technical support, and the engineers will help you upgrade the instance if necessary.

### [Method 1: Rolling Back a Single Instance]

- Step 1** In the **Upgrade Instance** dialog box, select **Rollback** for **Action**, select a target version, enter **confirm**, and click **OK**.
- Step 2** On the **Instances** page, check the rollback status. After the rollback is complete, the instance status changes to **Available**.
- Step 3** On the **Instances** page, click the name of the target instance. On the **Basic Information** page that is displayed, check that the value of **DB Engine Version** in the **DB Information** area is the source version, that is, the version before upgrade.
- Step 4** Check that the instance is properly connected and you can add, delete, update, and query data in the instance.
1. Log in to the database. For details, see [Connecting to an Instance Through DAS](#).
  2. Go to the **SQL Query** page.
  3. Create a database.  
**CREATE DATABASE** *database\_name*;  
In this example, run the following command to create a database named **db\_tpcds**:  
**CREATE DATABASE db\_tpcds**;  
Switch to the newly created database in the upper left corner.
  4. Create a table and add, delete, update, and query data in the table.
    - a. Create a schema.  
**CREATE SCHEMA** *myschema*;
    - b. Create a table named **mytable** that has only one column. The column name is **firstcol** and the column type is integer.  
**CREATE TABLE** *myschema.mytable (firstcol int)*;
    - c. Insert data into the table.  
**INSERT INTO** *myschema.mytable* values (100);
    - d. View data in the table.  
**SELECT \* FROM** *myschema.mytable*;

```
| firstcol |
-----+
1 | 100 |
```

- e. Update data in the table.  
`UPDATE myschema.mytable SET firstcol = 200;`
- f. View the data in the table again.  
`SELECT * FROM myschema.mytable;`  

|   | firstcol |
|---|----------|
| 1 | 200      |
- g. Drop the table.  
`DROP TABLE myschema.mytable;`

----End

## [Method 2: Rolling Back Instances in Batches]

- Step 1** On the **Instances** page, select the target instances and click **Batch Upgrade**.
- Step 2** In the **Batch Upgrade** dialog box, select **Rollback** for **Action**, select a target version, enter **confirm**, and click **OK**.
- Step 3** On the **Instances** page, check the rollback status. After the rollback is complete, the instance status changes to **Available**.
- Step 4** On the **Instances** page, click the name of the target instance. On the **Basic Information** page that is displayed, check that the value of **DB Engine Version** in the **DB Information** area is the source version, that is, the version before upgrade.
- Step 5** Check that the instance is properly connected and you can add, delete, update, and query data in the instance.
1. Log in to the database. For details, see [Connecting to an Instance Through DAS](#).
  2. Go to the **SQL Query** page.
  3. Create a database.  
`CREATE DATABASE database_name;`  
In this example, run the following command to create a database named **db\_tpcds**:  
`CREATE DATABASE db_tpcds;`  
Switch to the newly created database in the upper left corner.
  4. Create a table and add, delete, update, and query data in the table.
    - a. Create a schema.  
`CREATE SCHEMA myschema;`
    - b. Create a table named **mytable** that has only one column. The column name is **firstcol** and the column type is integer.  
`CREATE TABLE myschema.mytable (firstcol int);`
    - c. Insert data into the table.  
`INSERT INTO myschema.mytable values (100);`
    - d. View data in the table.  
`SELECT * FROM myschema.mytable;`

```
| firstcol |
-----+
1 | 100 |
```

- e. Update data in the table.  
UPDATE myschema.mytable SET firstcol = 200;
- f. View the data in the table again.  
SELECT \* FROM myschema.mytable;

```
| firstcol |
-----+
1 | 200 |
```

- g. Drop the table.  
DROP TABLE myschema.mytable;

----End

# 6 Data Backup

---

## 6.1 Overview

You can back up your GaussDB instances to ensure data reliability. Currently, backups are stored in an unencrypted form.

### Functions

Although GaussDB supports high availability, if a database or table is maliciously or mistakenly deleted, data on the standby nodes is also deleted. In this case, you can only restore the deleted data from backups.

### Full Backup

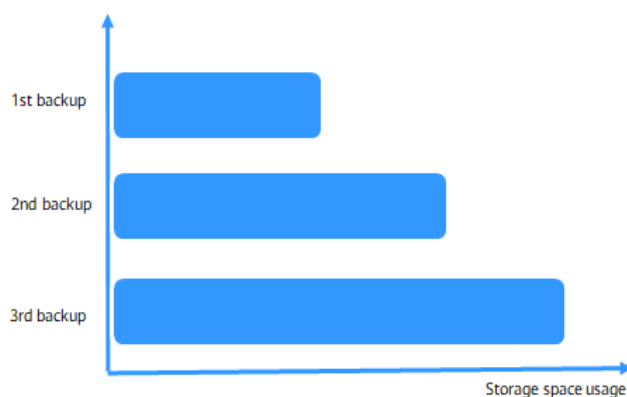
A full backup involves all data of a database at the backup point in time. The time required for full backup is long (in direct proportion to the total data volume of the database). You can use a full backup to restore data of a complete database. A full backup backs up all data even if the data has not changed since the last backup.

### Differential Backup

A differential backup involves only incremental data modified after a specified time point. It takes less time than a full backup in direct proportion to how much data has changed (The total data volume is irrelevant). However, a differential backup cannot be used to restore all of the data of a database. By default, the system automatically backs up updated data every 30 minutes since the last automated backup. The backup period can be changed from 15 minutes to 1,440 minutes.

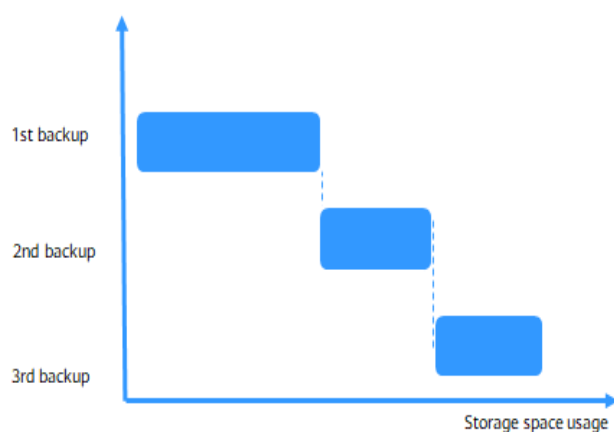
## Backup Principles

**Figure 6-1 Full backup**



Full backup: After the first full backup, all data is backed up in the second and third backups regardless of whether the data is changed.

**Figure 6-2 Differential backup**



Differential backup: After the first full backup, the second backup backs up only the changed data, and the third backup backs up only the data changed after the second backup.

## Automated Backup

Automated backups are created during the backup time window of your GaussDB instances. The system saves automated backups based on a retention period you specify.

## Manual Backup

Manual backups are user-initiated full backups of instances. They are retained until you delete them manually.

## 6.2 Backup Execution

### 6.2.1 Configuring an Automated Backup Policy for GaussDB Instances

#### Scenarios

When you create a GaussDB instance, an instance-level automated backup policy is enabled by default. After your instance is created, you can modify the automated backup policy as needed. GaussDB backs up data based on the automated backup policy you specified.

#### Constraints


- The instance-level automated backup policy cannot be configured for GaussDB single-replica instances of versions earlier than V2.0-3.0.
- To ensure that data can be restored to a specific point in time, the most recent full backup that exceeds the backup retention period will not be deleted immediately.

For example, with the automated backup policy set to perform daily backups and retain them for one day, if backup 1 is generated on November 1, this backup will not be deleted on November 2 when backup 2 is generated, but will be deleted on November 3 when backup 3 is generated and backup 2 is retained.

#### Modifying an Automated Backup Policy

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Backups**. On the displayed page, click **Modify Backup Policy**. You can view the configured backup policy. To modify the backup policy, adjust the parameter values as needed.

**Step 6** Configure parameters.

**Table 6-1** Full backup policy

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Retention Period    | <p>Number of days that your automated backups can be retained. Increasing the retention period will improve data reliability.</p> <p>Full backups are retained till the retention period expires. However, even if the retention period has expired, the most recent backup will be retained.</p> <ul style="list-style-type: none"><li>• Extending the retention period improves data reliability. You can extend the retention period as needed.</li><li>• If you shorten the retention period, the new backup policy takes effect for existing backups. Any automated backups (including full and incremental backups) that have expired will be automatically deleted. Manual backups will not be automatically deleted but you can delete them manually.</li></ul> <p><b>Policy for automatically deleting automated full backups:</b></p> <p>To ensure data integrity, even after the retention period expires, the most recent backup will be retained.</p> <p>If <b>Backup Cycle</b> was set to <b>Monday</b> and <b>Tuesday</b> and the <b>Retention Period</b> was set to <b>2</b>:</p> <ul style="list-style-type: none"><li>• The full backup generated on Monday will be automatically deleted on Thursday. The reasons are as follows:<br/>The backup generated on Monday expires on Wednesday, but it was the last backup, so it will be retained until a new backup expires. The next backup will be generated on Tuesday and will expire on Thursday. So the full backup generated on Monday will not be automatically deleted until Thursday.</li><li>• The full backup generated on Tuesday will be automatically deleted on the following Wednesday. The reasons are as follows:<br/>The backup generated on Tuesday will expire on Thursday, but as it is the last backup, so it will be retained until a new backup expires. The next backup will be generated the next Monday and will expire on the next Wednesday. So the full backup generated on Tuesday will not be automatically deleted until the next Wednesday.</li></ul> |
| Standby Node Backup | If this policy is enabled, full and differential backups of the instance are performed on the hosts where standby DN's reside.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Time Window         | A one-hour period the backup will be scheduled within 24 hours, such as 01:00-02:00 or 12:00-13:00. The backup time is in UTC format. If the DST or standard time is switched, the backup time segment changes with the time zone.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

| Parameter    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Backup Cycle | <p>One or more days from Monday to Sunday when a backup task will be triggered. By default, all options are selected. Select at least one day as required.</p> <p>Backups can be retained for 1 to 732 days.</p> <p>A time window is one hour. A total of 24 time windows are available. You are advised to select an off-peak time window for full backups. By default, each day of the week is selected for <b>Backup Cycle</b>. You can change the backup cycle. At least one day must be selected.</p> <p>A full backup is immediately triggered after a DB instance is created. Then, a full backup or differential backup is performed based on the time window and backup cycle you specified. We recommend that you set the full backup time window to an off-peak hour.</p> |

**Table 6-2** Differential backup policy

| Parameter    | Description                                                                                                |
|--------------|------------------------------------------------------------------------------------------------------------|
| Backup Cycle | You need to select the backup cycle for performing a differential backup. The default value is 30 minutes. |

**Step 7** Click **OK**.

**Step 8** Check the result.

After the task is submitted, click **Modify Backup Policy** to check whether the modification is successful.

----End

## 6.2.2 Configuring an Automated Backup Policy for GaussDB Tables

### Scenarios

When you create a GaussDB instance, the table-level automated backup policy is disabled by default. After your instance is created, no table-level full backup task will be triggered. You can modify the automated backup policy for tables as needed. If such a policy is configured, GaussDB backs up data of specified tables based on the automated backup policy you specified.


### Precautions


- After an instance-level restoration is complete, the databases or tables configured in the table-level automated backup policy may no longer exist. In this case, you need to configure the table-level automated backup policy again.

- To ensure that data can be restored to a point in time, the latest full backup that exceeds the backup retention period will not be deleted immediately. For example, if **Backup Cycle** is set to **All** and **Retention Period** to one day and backup 1 is generated on November 1, this table backup will not be deleted on November 2 when backup 2 is generated, but will be deleted on November 3 when backup 3 is generated.
- Tables in the **postgres**, **template0**, **template1**, **templatem**, **templatea**, and **templatepdb** system databases cannot be backed up using the table-level backup method.
- Tables in system schemas (for example, **public**) cannot be backed up using the table-level backup method.
- Tables in M-compatible databases cannot be backed up using the table-level backup method.
- A maximum of 100 databases or tables can be restored at the same time. If there are more than 100 databases or tables, you are advised to use instance-level backup and restoration.
- Table-level restoration does not support segment-page tables, column-store tables, tables containing user-defined types, synonym tables, temporary tables (including global tables), encrypted tables, TDE-enabled tables, unlogged tables, compressed tables, tables of private users, and tables of ledger databases.
- Table-level backup is unavailable to single-node instances.
- If the name of a table, schema, or database to be backed up is changed, you need to reconfigure the table-level backup policy.
- If there is a new table available for backup, the new table will be backed up during the next scheduled full or differential backup.
- If a table covered by the table-level backup policy is deleted, modify the policy immediately to remove that table. Otherwise, the table-level backup will keep failing.
- To disable the table-level differential backup policy, [submit a service ticket](#).

## Modifying an Automated Backup Policy for Tables

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Backups**. On the displayed page, choose **Full Backups > Table Backup**. Click **Modify Backup Policy**.

**Step 6** Configure parameters.

**Table 6-3** Backup policies

| Backup Policy              | Parameter        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Full Backup Policy         | Retention Period | Number of days that your automated backups can be retained. Increasing the retention period will improve data reliability. The default value is 7.<br><br>If you shorten the retention period, the new backup policy will be applied for all backup files. The backup files that have expired will be deleted.<br><br>Backups can be retained for 1 to 732 days. To extend the retention period, <a href="#">submit a service ticket</a> . Automated backups can be retained for up to 2,562 days. |
|                            | Time Window      | A one-hour period the backup will be scheduled within 24 hours, such as 01:00-02:00 or 12:00-13:00. To minimize the potential impact on services, set the time window to an off-peak hour.                                                                                                                                                                                                                                                                                                         |
|                            | Backup Cycle     | One or more days from Monday to Sunday when a backup task will be triggered. By default, all options are selected. Select at least one day as required.                                                                                                                                                                                                                                                                                                                                            |
| Differential Backup Policy | Backup Cycle     | Interval at which table-level differential backup is performed. The default cycle is 30 minutes.                                                                                                                                                                                                                                                                                                                                                                                                   |

**Step 7** Click **Modify Backup Policy**.

----End

## 6.2.3 Creating a Manual Backup for GaussDB Instances

### Scenarios

GaussDB allows you to create instance-level manual backups for running instances.

### Precautions

- Manual backups are user-initiated full backups of instances. They are retained until you delete them manually.
- You can back up data of instances that are in the **Available** state.
- A user can perform only one instance-level backup operation for a DB instance at a time.
- Instance-level manual backups cannot be created for GaussDB single-replica instances of versions earlier than V2.0-3.0.

## Billing


Backups are saved as packages in OBS buckets.

After a DB instance is deleted, the free backup space of the DB instance is automatically canceled. Manual backups are billed based on the space required.

### Method 1

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

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

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

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

**Step 5** In the displayed dialog box, enter a backup name and enter the description as needed. Then, click **OK**. If you want to cancel the backup creation task, click **Cancel**.

- The backup name is case-sensitive and can contain 4 to 64 characters. It must start with a letter. Only letters, digits, hyphens (-), and underscores (\_) are allowed.
- The description can contain up to 256 characters, but cannot contain carriage returns and the following special characters: > ! < " & ' =
- During the creation process, the instance status is **Backing up**. The time required for creating a manual backup depends on the data volume.

**Step 6** View and manage the created backup on the **Backups** page.


Alternatively, click the instance name. On the **Backups** page, you can view and manage the manual backups.

----End

### Method 2

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Backups**. On the displayed page, click **Create Backup**.

**Step 6** In the displayed dialog box, enter a backup name and description and click **OK**.

- The backup name is case-sensitive and can contain 4 to 64 characters. It must start with a letter. Only letters, digits, hyphens (-), and underscores (\_) are allowed.
- The description can contain up to 256 characters, but cannot contain carriage returns and the following special characters: > ! < " & ' =
- During the creation process, the manual backup status is **Creating**. The time required for creating a manual backup depends on the data volume.

**Step 7** View and manage the created backup on the current page.

Alternatively, go back to the instance list page, and click **Backups** to view and manage the backup.

----End

## 6.2.4 Creating a Manual Backup for GaussDB Tables

### Scenarios

GaussDB allows you to create table-level manual backups for available instances.

### Precautions

- You can back up data of instances that are in the **Available** state.
- A user can perform only one table-level backup operation for a DB instance at a time.
- Tables in the **postgres**, **template0**, **template1**, **templatea**, and **templatepdb** system databases cannot be backed up using the table-level backup method.
- Tables in M-compatible databases cannot be backed up using the table-level backup method.
- Tables in system schemas (for example, **public**) cannot be backed up using the table-level backup method.

By default, **public** is a system schema. However, you can run commands to change the owner of **public**. When the owner is changed to a non-system user, you can back up tables in the **public** schema.

- When a DB instance is deleted, the automated backups for tables are also deleted, but manual table backups are retained.

### Billing


Backups are saved as packages in OBS buckets.

After a DB instance is deleted, the free backup space of the DB instance is automatically canceled. Manual backups are billed based on the space required.

### Procedure

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

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

- Step 3** Click  in the upper left corner of the page and choose **Databases > GaussDB**.
- Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.
- Step 5** In the navigation pane on the left, choose **Backups**. On the displayed page, click **Table Backup** under **Full Backups** and click **Create Backup**.
- Step 6** Enter the backup name and description, select the required data and click **Create Backup**.
- The backup name is case-sensitive and can contain 4 to 64 characters. It must start with a letter. Only letters, digits, hyphens (-), and underscores (\_) are allowed.
  - The description can contain up to 256 characters, but cannot contain carriage returns and the following special characters: >!<"&'=
  - During the creation process, the status is **Creating**. The time required for creating a manual backup depends on the data volume.
- Step 7** View and manage the created backup on the current page.
- Alternatively, go back to the instance list page, and click **Backups** to view and manage the backup.
- End

## 6.3 Backup Management

### 6.3.1 Exporting Backup Information About GaussDB Instances



#### Scenarios


You can export backup information of instances to a CSV file for further analysis. The exported information includes the backup ID, backup name, instance name, instance ID, DB engine, backup type, backup start time, backup end time, backup status, backup size, and backup description.

#### Precautions

- Backup information cannot be exported for GaussDB single-replica instances of versions earlier than V2.0-3.0.
- You can only export the backup data shown on the current page; data on other pages cannot be exported.

#### Procedure

- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Click  in the upper left corner of the page and choose **Databases > GaussDB**.

**Step 4** In the navigation pane on the left, choose **Backups**. Select the backups you want to export and click  to export the backup information.

Alternatively, click the instance name. On the **Backups** page, select the backups you want to export and click **Export** above the backup list to export the backup information.

The exported backup information is in a CSV file which facilitates your analysis.

**Step 5** View the exported backup information.

----End

## 6.3.2 Deleting a Manual Backup of a GaussDB Instance

### Scenarios

You can delete manual backups for GaussDB instances to release storage space.


### Precautions

- Deleted manual backups cannot be recovered. Exercise caution when performing this operation.
- Automated backups cannot be manually deleted.
- Backups that are being restored cannot be deleted.
- To delete a backup, you must log in to the account that the backup belongs to.
- Manual backups cannot be deleted for GaussDB single-replica instances of versions earlier than V2.0-3.0.

### Procedure

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

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

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

**Step 4** In the navigation pane on the left, choose **Backups**. On the displayed page, locate the manual backup you want to delete and click **Delete** in the **Operation** column.

Alternatively, on the **Instances** page, click the instance name to go to the **Basic Information** page. In the navigation pane on the left, choose **Backups**. On the displayed page, locate the manual backup you want to delete and click **Delete** in the **Operation** column.

**Step 5** Click **OK**.

After the backup is deleted, it will not be displayed on the **Backups** page.

----End

# 7 Data Restoration

---

## 7.1 Restoring a Backup File to a GaussDB Instance

### Scenarios

You can use an instance-level automated or manual backup to restore data to the point in time when the backup was created. The restoration is at the DB instance level.

Data can be restored to a new DB instance, an existing DB instance, or the original DB instance.


### Precautions

- Restoration will fail if the instance is in the **Abnormal**, or **Storage full** state.
- GaussDB currently only supports restoration between DB instances running the same major version. For example, backup data can only be restored from version V2.0-1.4.x to version V2.0-1.4.y.
- In addition to full backups and incremental backups, the system backs up incremental log files to ensure data consistency. It takes some time to back up and upload incremental log files (The time depends on the network and OBS traffic control). Note that the backup completion time does not represent the data consistency point that can be specified when this backup set is used to restore data. (Generally, the data consistency point is within several minutes before the backup completion time.) If you have strict requirements on data consistency after restoration, restore data to a specified point in time.

### Restoring Data to a New Instance

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

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

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

**Step 4** In the navigation pane, choose **Backups**. On the **Backups** page, locate the backup to be restored and click **Restore** in the **Operation** column.

Alternatively, click the name of the target instance on the **Instances** page. In the navigation pane, choose **Backups**. On the **Full Backups** page, click the **Instance Backup** tab, and click **Restore** in the **Operation** column of the backup to be used for restoration.

**Step 5** Set **Restoration Method** to **Create New Instance** and click **OK**.

- The new DB instance must have the same major version as the original instance when it was backed up. For example, backup data can only be restored from version V2.0-1.4.x to version V2.0-1.4.y.
- The storage space of the new instance is the same as that of the original DB instance by default and the new instance must be at least as large as the original DB instance. The storage space for a single shard starts from 40 GB and can be increased at a step of 4 GB.
- The administrator password needs to be reset.
- By default, the instance specifications of the new instance are the same as those of the original instance. To change the instance specifications, ensure that the instance specifications of the new instance are at least those of the original instance.
- The backup media, CPU architecture, OS type and version, instance type (distributed or centralized), deployment model, replica consistency protocol, and transaction consistency settings of the new instance must be the same as those of the original instance to be restored.
- If there are M-compatible databases in the instance to be restored, the **Tables Names in M-compatible Databases** parameter of the new instance must be set to the same value as that of the original instance. The default value is **Case-sensitive**.

**Step 6** On the displayed page, configure parameters of the new instance and click **Create Now**.

**Step 7** View the restoration results.

A new instance that contains the same data as the backup is created. When the instance status changes from **Creating** to **Available**, the restoration is complete.


The new instance is independent from the original one.

----End

## Restoring Data to the Original Instance

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

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

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

**Step 4** In the navigation pane, choose **Backups**. On the **Backups** page, locate the backup to be restored and click **Restore** in the **Operation** column.

Alternatively, click the name of the target instance on the **Instances** page. In the navigation pane, choose **Backups**. On the **Full Backups** page, click the **Instance Backup** tab, and click **Restore** in the **Operation** column of the backup to be used for restoration.

**Step 5** Set **Restoration Method** to **Restore to Original**, select the confirmation check box, and click **Next**.

- The instance version and node configuration must be the same as those of the original instance when the backup was created.
- Restoring to the original DB instance will overwrite all data on it and cause the DB instance to be unavailable during the restoration.
- You are advised to manually back up data before the restoration.
- If you use a backup created before advanced compression is enabled to restore data to the original instance, you must enable this feature for the instance again.

**Step 6** Confirm the backup task information and click **OK**.

**Step 7** View the restoration results.


On the **Instances** page, when the instance status changes from **Restoring** to **Available**, the restoration is complete. After the restoration is complete, an instance-level full backup will be automatically triggered.


After the restoration is complete, check whether the restored data is consistent with the time point to which the data is restored. If they are consistent, click **Confirm Data Integrity** on the **Backups** page. Before confirming data integrity, you can restore data for multiple times. Once data integrity has been confirmed, any logs archived after the point in time data was restored from will be lost, but normal log archiving will be restored.

----End

## Restoring Data to an Existing Instance

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

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

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

**Step 4** In the navigation pane, choose **Backups**. On the **Backups** page, locate the backup to be restored and click **Restore** in the **Operation** column.

Alternatively, click the name of the target instance on the **Instances** page. In the navigation pane, choose **Backups**. On the **Full Backups** page, click the **Instance Backup** tab, and click **Restore** in the **Operation** column of the backup to be used for restoration.

**Step 5** Set **Restoration Method** to **Restore to Existing**, select the confirmation check box, select an existing instance, and click **Next**.

- Restoring data to an existing instance will overwrite all data on it and cause the instance to be unavailable during the restoration.

- You are advised to manually back up data of the selected instance before the restoration.
- If there are M-compatible databases in the instance to be restored, the **Tables Names in M-compatible Databases** parameter of the new instance must be set to the same value as that of the original instance. The default value is **Case-sensitive**.
- The backup media, CPU architecture, OS type and version, instance type (distributed or centralized), deployment model, replica consistency protocol, transaction consistency, resource specifications, and failover priority settings of the selected instance must be the same as those of the original instance to be restored.
- An existing instance cannot be selected as the restoration target if it is currently undergoing a full or differential backup. If the backup is not required, you can stop it before proceeding with the restoration.

**Step 6** Confirm the backup task information and click **OK**.

**Step 7** View the restoration results.

On the **Instances** page, when the instance status changes from **Restoring** to **Available**, the restoration is complete. After the restoration is complete, an instance-level full backup will be automatically triggered.

----End

## 7.2 Restoring a GaussDB Database or Table Using a Backup File

### Scenarios

You can use an instance-level automated or manual backup to restore data of specified databases or tables to the point in time when the backup was created. This operation restores data in specified databases or tables.

You can use a table-level automated or manual backup to restore data in tables to the point in time when the backup was created. This operation restores data only in specified tables.

Data can be restored to a new DB instance, an existing DB instance, or the original DB instance.


### Precautions


- Only databases and tables in instances of version V2.0-3.200 or later can be restored.
- Restoration will fail if the instance is in the **Abnormal**, or **Storage full** state.
- Backup data cannot be restored across major versions. For example, backup data of an instance of version V2.0-3.200.x can only be restored to version V2.0-3.200.y.
- A maximum of 100 databases or tables can be restored at the same time. If there are more than 100 databases or tables, you are advised to use instance-level restoration.

- The table names in a given database and schema as well as the database names must be different before and after the restoration.
- Table-level restoration does not support column-store tables, user-defined tables, synonym tables, temporary tables (including global tables), unlogged tables, tables of private users, and tables of ledger databases.
- Currently, only heap data, index data, and auto-increment column data related to tables can be restored. Other data related to tables, such as foreign key information, triggers, and SQL throttling rules, cannot be restored.
- After table-level restoration, row-level access control and dynamic masking information will be lost.
- System databases (**postgres**, **template0**, **template1**, **templatem**, **templatea**, and **templatepdb**) and their tables cannot be selected for database- and table-level restoration.
- System schemas (for example, **public**) and their tables cannot be selected for database- and table-level restoration.
- M-compatible databases and their tables cannot be selected for database- and table-level restoration.
- Single-replica instances do not support database- and table-level backup and restoration.
- Snapshot-based backups cannot be used to restore databases or tables.
- If transparent data encryption (TDE) is enabled for an instance and the instance version is earlier than V2.0-8.200, database- and table-level restoration is not supported.
- Restoring a database or table using a backup file is only available to whitelisted users. To request this feature, [submit a service ticket](#).
- In addition to table-level full backups and incremental backups, the system backs up incremental log files to ensure data consistency. It takes some time to back up and upload incremental log files (The time depends on the network and OBS traffic control). Note that the backup completion time does not represent the data consistency point that can be specified when this backup set is used to restore data. (Generally, the data consistency point is within several minutes before the backup completion time.) If you have strict requirements on data consistency after restoration, restore data to a specified point in time.

## Restoring Data to a New Instance

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

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

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

**Step 4** In the navigation pane, choose **Backups**. On the **Backups** page, locate the backup to be restored and click **Restore** in the **Operation** column.

- If the selected backup is an instance-level backup, you can restore data to specified databases or tables.
- If the selected backup is a table-level backup, you can only restore data to specified tables.

Alternatively, on the **Instances** page, click the instance name to go to the **Basic Information** page. In the navigation pane on the left, choose **Backups**. Click **Instance Backup** or **Table Backup** under **Full Backups**. On the displayed page, locate the backup to be restored and click **Restore** in the **Operation** column.

**Step 5** Set **Restoration Method** to **Create New Instance** and click **OK**.

When you attempt to restore data to tables, a maximum of 200 tables can be displayed for **Backup Tables** by default. If the required table is not displayed, click the + icon in the upper right corner to add a custom table. (You can also use this method to add required tables if the instance is faulty or deleted.)

- The new DB instance must have the same major version as the original instance when it was backed up. For example, backup data of an instance of version V2.0-3.200.x can only be restored to version V2.0-3.200.y.
- The storage space of the new instance is the same as that of the original DB instance by default and the new instance must be at least as large as the original DB instance.
- The administrator password needs to be reset.
- By default, the specifications of the new instance are the same as those of the original instance. You can configure higher specifications for the new instance. For instances of version V2.0-8.200 or later, an instance backup can be restored to an instance with lower specifications than the original instance. (Minimum specifications for the destination instance: 8 vCPUs | 32 GB)
- The backup media, CPU architecture, OS type and version, instance type (distributed or centralized), deployment model, replica consistency protocol, and transaction consistency settings of the new instance must be the same as those of the original instance to be restored.
- If there are M-compatible databases in the instance to be restored, the **Tables Names in M-compatible Databases** parameter of the new instance must be set to the same value as that of the original instance. The default value is **Case-sensitive**.
- Parameters of the original instance will not be automatically restored to the new instance. To use the original parameter settings, select the required parameter template for **Parameter Template** when creating an instance for restoring data to a new instance. After the new instance is created, click the instance name and manually change the parameter values on the **Parameters** page.

**Step 6** On the displayed page, configure parameters of the new instance, select the table or database data to be restored, and click **Create Now**.

**Step 7** View the restoration results.


A new instance that contains the same databases or tables as the backup is created. When the instance status changes from **Creating** to **Available**, the restoration is complete.


The new instance is independent from the original one.

----End

## Restoring Data to the Original Instance

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

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

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

**Step 4** In the navigation pane, choose **Backups**. On the **Backups** page, locate the backup to be restored and click **Restore** in the **Operation** column.

- If the selected backup is an instance-level backup, you can restore data to specified databases or tables.
- If the selected backup is a table-level backup, you can only restore data to specified tables.

Alternatively, on the **Instances** page, click the instance name to go to the **Basic Information** page. In the navigation pane on the left, choose **Backups**. Click **Instance Backup** or **Table Backup** under **Full Backups**. On the displayed page, locate the backup to be restored and click **Restore** in the **Operation** column.

**Step 5** Set **Restoration Method** to **Restore to Original**, select the confirmation check box, and click **Next**.

When you attempt to restore data to tables, a maximum of 200 tables can be displayed for **Backup Tables** by default. If the required table is not displayed, click the + icon in the upper right corner to add a custom table. (You can also use this method to add required tables if the instance is faulty or deleted.)

- The instance version and node configuration must be the same as those of the original instance when the backup was created.
- Databases and tables will be created on the target DB instance. During the restoration, the source database can be used properly.
- You are advised to manually back up data before the restoration.
- After a database (for example, **db1**) is restored to the original instance, you need to wait for the DB instance to automatically perform a full or differential backup before restoring the data of **db1** using other restoration processes. The time to wait depends on the backup policy. During point-in-time restoration, if you select a time point that is later than when the current database-level restoration is complete but earlier than when the next backup is performed, **db1** cannot be restored.

**Step 6** On the **Restore DB Instance** page, select the data to be restored and click **Submit**.


**Step 7** View the restoration results.


On the **Instances** page, when the instance status changes from **Restoring** to **Available**, the restoration is complete.

----End

## Restoring Data to an Existing Instance

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

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

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

**Step 4** In the navigation pane, choose **Backups**. On the **Backups** page, locate the backup to be restored and click **Restore** in the **Operation** column.

- If the selected backup is an instance-level backup, you can restore data to specified databases or tables.
- If the selected backup is a table-level backup, you can only restore data to specified tables.

Alternatively, on the **Instances** page, click the instance name to go to the **Basic Information** page. In the navigation pane on the left, choose **Backups**. Click **Instance Backup** or **Table Backup** under **Full Backups**. On the displayed page, locate the backup to be restored and click **Restore** in the **Operation** column.

**Step 5** Set **Restoration Method** to **Restore to Existing**, select the confirmation check box, select an existing instance, and click **Next**.

When you attempt to restore data to tables, a maximum of 200 tables can be displayed for **Backup Tables** by default. If the required table is not displayed, click the + icon in the upper right corner to add a custom table. (You can also use this method to add required tables if the instance is faulty or deleted.)

- Databases and tables will be created on the target instance. During the restoration, the databases on the instance can be used properly.
- You are advised to manually back up data of the selected instance before the restoration.
- The backup media, CPU architecture, OS type and version, instance type (distributed or centralized), deployment model, replica consistency protocol, transaction consistency, resource specifications, and failover priority settings of the selected instance must be the same as those of the original instance to be restored.
- An existing instance cannot be selected as the restoration target if it is currently undergoing a full or differential backup. If the backup is not required, you can stop it before proceeding with the restoration.

**Step 6** Confirm the backup task information and click **OK**.

**Step 7** View the restoration results.

On the **Instances** page, when the instance status changes from **Restoring** to **Available**, the restoration is complete.

----End

## 7.3 Restoring a GaussDB Instance to a Specific Point in Time

### Scenarios

You can use an instance-level automated backup to restore a GaussDB instance to a specified point in time.


You can restore backup data to the original GaussDB instance, an existing instance, or a new one.


## Precautions

- Only instances of version V2.0-2.1 or later can be restored to any point in time. Single-replica instances are not supported.
- Data from the time period of the following operations cannot be restored: nodes are being added to an instance, the instance version is being upgraded, or the instance data is being restored.
- If an instance is faulty or a CN is removed, archive logs cannot be generated, and data from the affected time period cannot be restored.
- If you restore backup data to a new DB instance:
  - The DB engine and major version are the same as those of the original DB instance and cannot be changed.
  - The administrator password needs to be reset.
- If you restore backup data to the original DB instance, data on the original instance will be overwritten and the original DB instance will be unavailable during the restoration. Additionally, log archiving stops. After the restoration is complete, the **Confirm Data Integrity** button is displayed. Before clicking **Confirm Data Integrity**, you can restore data for multiple times. Once data integrity has been confirmed, any logs archived after the point in time data was restored from will be lost, but normal log archiving will be restored.
- When a DB instance is deleted, all archive logs are deleted by default and cannot be retained. After an instance is deleted, it cannot be rebuilt or restored to any point in time.

## Restoring Data to a New Instance

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

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

**Step 6** Set **Restoration Method** to **Create New Instance** and click **OK**.

- The new DB instance must have the same major version as the original instance when it was backed up. For example, backup data can only be restored from version V2.0-1.4.x to version V2.0-1.4.y.
- The storage space of the new instance is the same as that of the original DB instance by default and the new instance must be at least as large as the original DB instance.
- The administrator password needs to be reset.

- The backup media, CPU architecture, OS type and version, instance type (distributed or centralized), deployment model, replica consistency protocol, and transaction consistency settings of the new instance must be the same as those of the original instance to be restored.
- By default, the instance specifications of the new instance are the same as those of the original instance. To change the instance specifications, ensure that the instance specifications of the new instance are at least those of the original instance.
- If there are M-compatible databases in the instance to be restored, the **Tables Names in M-compatible Databases** parameter of the new instance must be set to the same value as that of the original instance. The default value is **Case-sensitive**.

**Step 7** On the displayed page, configure parameters of the new instance and click **Create Now**.

**Step 8** View the restoration results.

A new instance that contains the same data as the backup is created. When the instance status changes from **Creating** to **Available**, the restoration is complete.


The new instance is independent from the original one.

----End

## Restoring Data to the Original Instance

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

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

**Step 6** Set **Restoration Method** to **Restore to Original**, select the confirmation check box, and click **Next**.

- The instance version and node configuration must be the same as those of the original instance when the backup was created.
- Restoring to the original DB instance will overwrite all data on it and cause the DB instance to be unavailable during the restoration.
- You are advised to manually back up data before the restoration.
- If you use a backup created before advanced compression is enabled to restore data to the original instance, you must enable this feature for the instance again.

**Step 7** Confirm the backup task information and click **OK**.

**Step 8** View the restoration results.

On the **Instances** page, when the instance status changes from **Restoring** to **Available**, the restoration is complete. After the restoration is complete, a full backup will be automatically triggered.


After the restoration is complete, check whether the restored data is consistent with the time point to which the data is restored. If they are consistent, click **Confirm Data Integrity** on the **Backups** page. Before confirming data integrity, you can restore data for multiple times. Once data integrity has been confirmed, any logs archived after the point in time data was restored from will be lost, but normal log archiving will be restored.

----End

## Restoring Data to an Existing Instance

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

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

**Step 6** Set **Restoration Method** to **Restore to Existing**, select the confirmation check box, select an existing instance, and click **Next**.

- Restoring data to an existing instance will overwrite all data on it and cause the instance to be unavailable during the restoration.
- You are advised to manually back up data of the selected instance before the restoration.
- If there are M-compatible databases in the instance to be restored, the **Tables Names in M-compatible Databases** parameter of the target instance must be set to the same value as that of the original instance. The default value is **Case-sensitive**.
- The backup media, CPU architecture, OS type and version, instance type (distributed or centralized), deployment model, replica consistency protocol, transaction consistency, resource specifications, and failover priority settings of the selected instance must be the same as those of the original instance to be restored.
- An existing instance cannot be selected as the restoration target if it is currently undergoing a full or differential backup. If the backup is not required, you can stop it before proceeding with the restoration.

**Step 7** Confirm the backup task information and click **OK**.

**Step 8** View the restoration results.

On the **Instances** page, when the instance status changes from **Restoring** to **Available**, the restoration is complete. After the restoration is complete, a full backup will be automatically triggered.

----End

## 7.4 Restoring a GaussDB Database or Table to a Specific Point in Time

### Scenarios

You can use an instance-level automated backup to restore data of specified databases or tables to a specified point in time.

You can use a table-level automated backup to restore data in tables to a specified point in time.

You can restore backup data to the original GaussDB instance, an existing instance, or a new one.

### Precautions


- To use this function, your DB instance cannot be a single-replica instance and its version must be V2.0-3.200 or later.
- Data from the time period of the following operations cannot be restored: nodes are being added to an instance, the instance version is being upgraded, or the instance data is being restored.
- If an instance is faulty or a CN is removed, archive logs cannot be generated, and data from the affected time period cannot be restored.
- If you restore backup data to a new DB instance:
  - The DB engine and major version are the same as those of the original DB instance and cannot be changed.
  - The administrator password needs to be reset.
- If you restore backup data to the original DB instance, new databases or tables are created on the original instance and the original databases are still available during the restoration. The archive logs are normal.
- When a DB instance is deleted, all archive logs are deleted by default and cannot be retained. After an instance is deleted, it cannot be rebuilt or restored to any point in time.
- The table names in a given database and schema as well as the database names must be different before and after the restoration. If they are same, you need to rename the tables and databases that the data is restored to.
- If **ALTER DATABASE SET TABLESPACE** and **ALTER TABLE SET TABLESPACE** are executed in the original instance, table-level data can be restored only after a differential backup or full backup is triggered in the original instance.
- A maximum of 100 databases or tables can be restored at the same time. If there are more than 100 databases or tables, you are advised to use instance-level restoration.
- Table-level restoration does not support column-store tables, user-defined tables, synonym tables, temporary tables (including global tables), unlogged tables, tables of private users, and tables of ledger databases.
- Currently, only heap data, index data, and auto-increment column data related to tables can be restored. Other data related to tables, such as foreign key information, triggers, and SQL throttling rules, cannot be restored.

- After table-level restoration, row-level access control and dynamic masking information will be lost.
- System databases (**postgres**, **template0**, **template1**, **templatem**, **templatea**, and **templatepdb**) and their tables cannot be selected for database- and table-level restoration.
- System schemas (for example, **public**) and their tables cannot be selected for database- and table-level restoration.
- M-compatible databases and their tables cannot be selected for database- and table-level restoration.
- Single-replica instances do not support database- and table-level backup and restoration.
- PITR based on snapshots does not support database- or table-level restoration.
- If transparent data encryption (TDE) is enabled for an instance and the instance version is earlier than V2.0-8.200, database- and table-level restoration is not supported.

## Restoring Data to a New Instance

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the navigation pane, choose **Backups**. On the displayed page, click the **Instance Backup** or **Table Backup** tab.

- On the **Instance Backup** tab, click **Restore to Point in Time** to restore data of specified databases or tables.
- On the **Table Backup** tab, click **Restore to Point in Time** to restore data of specified tables.

**Step 6** Set **Restoration Method** to **Create New Instance** and click **OK**.

- The new DB instance must have the same major version as the original instance when it was backed up. For example, backup data of an instance of version V2.0-3.200.x can only be restored to version V2.0-3.200.y.
- The storage space of the new instance is the same as that of the original DB instance by default and the new instance must be at least as large as the original DB instance.
- The administrator password needs to be reset.
- The backup media, CPU architecture, OS type and version, instance type (distributed or centralized), deployment model, replica consistency protocol, and transaction consistency settings of the new instance must be the same as those of the original instance to be restored.
- If there are M-compatible databases in the instance to be restored, the **Tables Names in M-compatible Databases** parameter of the new instance must be

set to the same value as that of the original instance. The default value is **Case-sensitive**.

- By default, the instance specifications of the new instance are the same as those of the original instance. To change the instance specifications, ensure that the instance specifications of the new instance are at least those of the original instance.
- Parameters of the original instance will not be automatically restored to the new instance. To use the original parameter settings, select the required parameter template for **Parameter Template** when creating an instance for restoring data to a new instance. After the new instance is created, click the instance name and manually change the parameter values on the **Parameters** page.

**Step 7** On the displayed page, configure parameters of the new instance, select the table or database data to be restored, and click **Create Now**.

**Step 8** View the restoration results.


A new instance that contains the same data as the backup is created. When the instance status changes from **Creating** to **Available**, the restoration is complete.


The new instance is independent from the original one.

----End

## Restoring Data to the Original Instance

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the navigation pane, choose **Backups**. On the displayed page, click the **Instance Backup** or **Table Backup** tab.

- On the **Instance Backup** tab, click **Restore to Point in Time** to restore data of specified databases or tables.
- On the **Table Backup** tab, click **Restore to Point in Time** to restore data of specified tables.

**Step 6** Set **Restoration Method** to **Restore to Original**, select the confirmation check box, and click **Next**.

- The instance version and node configuration must be the same as those of the original instance when the backup was created.
- New databases or tables will be created on the original instance. The databases to be restored are available during the restoration.
- After a database (for example, **db1**) is restored to the original instance, you need to wait for the DB instance to automatically perform a full or differential backup before restoring the data of **db1** using other restoration processes. The time to wait depends on the backup policy. During point-in-time

restoration, if you select a time point that is later than when the current database-level restoration is complete but earlier than when the next backup is performed, **db1** cannot be restored.

- You are advised to manually back up data before the restoration.

**Step 7** On the page that is displayed, select the data to be restored and click **Submit**.


**Step 8** View the restoration results.


On the **Instances** page, when the instance status changes from **Restoring** to **Available**, the restoration is complete.

----End

## Restoring Data to an Existing Instance

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the navigation pane, choose **Backups**. On the displayed page, click the **Instance Backup** or **Table Backup** tab.

- On the **Instance Backup** tab, click **Restore to Point in Time** to restore data of specified databases or tables.
- On the **Table Backup** tab, click **Restore to Point in Time** to restore data of specified tables.

**Step 6** Set **Restoration Method** to **Restore to Existing**, select the confirmation check box, select an existing instance, and click **Next**.

- Databases and tables will be created on the target instance. During the restoration, the databases on the instance can be used properly.
- You are advised to manually back up data of the selected instance before the restoration.
- The backup media, CPU architecture, OS type and version, instance type (distributed or centralized), deployment model, replica consistency protocol, transaction consistency, resource specifications, and failover priority settings of the selected instance must be the same as those of the original instance to be restored.
- An existing instance cannot be selected as the restoration target if it is currently undergoing a full or differential backup. If the backup is not required, you can stop it before proceeding with the restoration.
- If there are M-compatible databases in the instance to be restored, the **Tables Names in M-compatible Databases** parameter of the new instance must be set to the same value as that of the original instance. The default value is **Case-sensitive**.

**Step 7** On the page that is displayed, select the data to be restored and click **Submit**.

**Step 8** View the restoration results.

On the **Instances** page, when the instance status changes from **Restoring** to **Available**, the restoration is complete. After the restoration is complete, an instance-level full backup will be automatically triggered.

----**End**

# 8 Parameter Management

---

## 8.1 Modifying GaussDB Instance Parameters

You can modify parameters of a GaussDB instance to bring out the best possible performance of the instance. You can also check the parameter values of an instance.

GaussDB provides the following types of parameters:


- **Public parameters:** GaussDB uses a set of default running parameters after it is installed. You can modify the parameters to better fit your application scenarios and data volume.
- **Parameters for data redistribution:** These parameters are used to control the data redistribution policy during database scale-out.


### Precautions

- Parameters for data redistribution can be modified only for distributed instances of version V2.0-2.6 or later.
- Parameters of read replicas can be modified only for centralized (1 primary + 2 standby) instances of version V2.0-2.7.1 or later.
- Before modifying parameters, make sure you understand their meanings and fully verify the changes in a test environment to avoid instance or service exceptions caused by inappropriate parameter settings.

### Modifying Common Parameters of the Current Instance

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Parameters**.


- You can modify and query the parameters applied to the instance on this page. After modifying parameters, you can preview the changes or cancel the modification. After confirming that all changes are correct, click **Save**.  
The modification of some parameters is only applied after the instance is rebooted. After you modify a parameter value, view the value in the **Effective upon Reboot** column.
  - If the value is **Yes** and the instance status on the **Instances** page is **Parameter change. Pending reboot**, you must reboot the instance for the modifications to take effect.
  - If the value is **No**, the modifications take effect immediately for the instance.
- You can click **Replicate** to save the parameters of the instance as a parameter template. You can view the parameter template under the **Custom Templates** tab of the **Parameter Templates** page. For details, see [Managing Parameter Templates for GaussDB Instances](#).
- You can click **Export** to download the parameters of the instance to your local PC.
- You can click **Compare** to compare the parameter template applied to the current instance with an existing parameter template.

----End

## Modifying Common Parameters of Multiple Instances at a Time

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

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

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

**Step 4** In the navigation pane on the left, choose **Parameter Templates**. Click the **Custom Templates** tab, and click the name of the target parameter template.

**Step 5** Modify parameters as needed.

- Click **Save**. In the displayed dialog box, click **Yes** to save the modification.
- To cancel your changes, click **Cancel**.
- To preview your changes, click **Preview**.



**Step 6** After the parameters are modified, click **Change History** to view what changes have been made.

The changes take effect only after you apply the parameter template to instances. For details, see [Applying a Parameter Template](#).

----End

## Modifying Data Redistribution Parameters of the Current Instance

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

- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Click  in the upper left corner of the page and choose **Databases > GaussDB**.
- Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.
- Step 5** In the navigation pane on the left, choose **Parameters**. On the displayed page, click **Change Parameters for Scale-out** or **Change Parameters for Redistribution**.

----End

## 8.2 Viewing Parameter Change History of a GaussDB Instance



### Scenarios

After modifying parameters of a GaussDB instance or those in a custom template, you can check their change history.

### Precautions

- The change history of a newly replicated or created parameter template is initially blank.
- The change history of the last seven days is displayed.
- The parameter change history of read replicas is only available for centralized (1 primary + 2 standby) instances of version V2.0-2.7.1 or later.

### Viewing Change History of DB Instance Parameters

- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Click  in the upper left corner of the page and choose **Databases > GaussDB**.
- Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.
- Step 5** In the navigation pane on the left, choose **Parameters**.
- Step 6** On the displayed page, click **Change History**.


You can view the parameter name, original parameter value, new parameter value, modification status, modification time, application status, and application time.


You can apply the parameter template to instances as required by referring to [Applying a Parameter Template](#).

----End

## Viewing Change History of a Parameter Template

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

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

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

**Step 4** Choose **Parameter Templates** in the navigation pane on the left. On the **Custom Templates** page, click the parameter template name.

**Step 5** On the displayed page, choose **Change History** in the navigation pane on the left.

You can view the parameter name, original parameter value, new parameter value, modification status, and modification time.

----End

## 8.3 Exporting Parameters of a GaussDB Instance

### Scenarios


You can export the parameter template information (parameter names, values, and descriptions) of a DB instance to a CSV file for analysis.


### Precautions

The parameters of read replicas can be exported only for centralized (1 primary + 2 standby) instances of version V2.0-2.7.1 or later.

### Procedure

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Parameters**. On the displayed page, click **Export** above the parameter list.

Exporting to a file: You can export the parameter template information (parameter names, parameter values, and descriptions) of an instance to a CSV file for analysis.

**Step 6** In the displayed dialog box, enter the file name and click **OK**.

The file name must start with a letter and can contain 4 to 81 characters. Only letters, digits, hyphens (-), and underscores (\_) are allowed.

----End

## 8.4 Creating a Custom Parameter Template for GaussDB Instances

### Scenarios

Parameter templates store engine configurations that can be applied to one or more DB instances. In GaussDB, there are two types: default and custom parameter templates.

If you do not specify a custom parameter template when creating an instance or restoring data to a new instance, the system defaults to applying the default parameter template.

Note that its parameters are immutable. For tailored needs, you can create a custom parameter template and select it during instance creation. You can also apply a custom parameter template to existing instances. For details, see [Applying a Parameter Template](#).



If you have already created a parameter template and want to provide most of its custom parameters and values in a new parameter template, you can replicate the template you created following the instructions provided in [Replicating a Parameter Template](#).

### Precautions

- Not all of the DB engine parameters in a custom parameter template can be changed. The parameters related to compute specifications (that is, the parameters whose default values are **default**) in the parameter template cannot be modified.
- A maximum of 100 GaussDB database parameter templates can be created for each project. All GaussDB engines share the parameter template quota.
- Certain parameter changes require a manual restart of the instance to take effect.
- Improper parameter settings may have unintended consequences, including reduced performance and system instability. Exercise caution when changing database parameters; always back up your data beforehand. Do not perform boundary tests within the parameter template to prevent system instability. Before applying parameter template changes to a production DB instance, you should try out these changes on a test DB instance.

### Procedure

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



- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Click  in the upper left corner of the page and choose **Databases > GaussDB**.
- Step 4** In the navigation pane on the left, choose **Parameter Templates**.
- Step 5** On the **Parameter Templates** page, click **Create Parameter Template**.
- Step 6** In the displayed dialog box, configure required information and click **OK**.
- Select a DB engine for the parameter template.
  - The template name is case-sensitive and can contain 1 to 64 characters. Only uppercase letters, lowercase letters, digits, hyphens (-), underscores (\_), and periods (.) are allowed.
  - The template description can contain a maximum of 256 characters and cannot include carriage return characters and the following special characters: `>|<"&'=`
- End

## 8.5 Managing Parameter Templates for GaussDB Instances

You can perform the following operations to manage GaussDB parameter templates:

- **Applying a parameter template:** Modifications to parameters in a parameter template take effect for instances only after you apply this parameter template to target instances. A parameter template can be applied only to instances of the same version.
- **Replicating a parameter template:** You can replicate a custom parameter template that you have created. When you have already created a parameter template and want to include most of the custom parameters and values from that template in a new parameter template, you can replicate that parameter template.
- **Comparing instance parameters with a parameter template:** You can compare instance parameters with a parameter template to see the differences of parameter settings.
- **Comparing parameter templates:** You can compare two default GaussDB parameter templates to see the differences between them. You can also compare two custom parameter templates.
- **Resetting a parameter template:** You can reset all parameters in a parameter template you have created to their default settings.
- **Modifying the description of a parameter template:** You can edit the description of a parameter template you have created.
- **Deleting a parameter template:** You can create up to 100 parameter templates and delete parameter templates that are no longer used.

## Applying a Parameter Template

- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Click  in the upper left corner of the page and choose **Databases > GaussDB**.
- Step 4** On the **Parameter Templates** page, perform the following operations based on the type of the parameter template to be applied:
- If you intend to apply a default parameter template to instances, click the **Default Templates** tab and click **Apply** in the **Operation** column of the target parameter template.
  - If you intend to apply a custom parameter template to instances, click **Custom Templates** and choose **More > Apply** in the **Operation** column of the target parameter template.

A parameter template can be applied to one or more instances.



- Step 5** In the displayed dialog box, select one or more instances to which the parameter template will be applied and click **OK**.
- Step 6** After the parameter template is applied, check its application records.
- If you intend to check the application records of a default parameter template, click the **Default Templates** tab and click **View Application Record** in the **Operation** column of the target parameter template.
  - If you intend to check the application records of a custom parameter template, click **Custom Templates** and choose **More > View Application Record** in the **Operation** column of the target parameter template.

If the application status of an instance is **Applying**, it will not be displayed in the instance list when you apply a parameter template again. If you want to apply the parameter template to the same instance again, ensure that the application status is **Successful**.

After the parameter template is successfully applied, if you modify parameters in the parameter template and the instance status is **Parameter change. Pending reboot**, you must reboot the instance for the modifications to take effect. If no parameter that requires an instance reboot is modified, the instance status remains unchanged.

-----End

## Replicating a Parameter Template

- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Click  in the upper left corner of the page and choose **Databases > GaussDB**.
- Step 4** On the **Parameter Templates** page, click the **Custom Templates** tab. Locate the parameter template to be replicated and click **Replicate** in the **Operation** column.

**Step 5** In the displayed dialog box, configure required details and click **OK**.

- The template name is case-sensitive and can contain 1 to 64 characters. Only letters, digits, hyphens (-), underscores (\_), and periods (.) are allowed.
- The template description can contain up to 256 characters, but cannot contain carriage returns and the following special characters: > ! < " & ' =

After the parameter template is replicated, a new template is generated in the list on the **Custom Templates** tab of the **Parameter Templates** page. The new parameter template may not be displayed immediately. You are advised to wait for at least 5 minutes before using the new template.

Default parameter templates cannot be replicated, but you can create parameter templates based on them.

----End

## Comparing Instance Parameters with a Parameter Template


---


### CAUTION

Parameters of read replicas can be compared only for centralized (1 primary + 2 standby) instances of version V2.0-2.7.1 or later.

---

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

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

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

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

**Step 5** In the navigation pane on the left, choose **Parameters**.

**Step 6** On the displayed page, click **Compare** to compare the parameters of the current instance.

**Step 7** In the displayed dialog box, select a parameter template that you want to compare with parameters of the current instance and click **OK**.


- If their settings are different, the parameter names and values of both parameter templates are displayed.
- If their settings are the same, no data is displayed.

----End

## Comparing Parameter Templates

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

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

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

**Step 4** On the **Parameter Templates** page, click **Default Templates** or **Custom Templates**. Locate a parameter template and click **Compare** in the **Operation** column.


**Step 5** In the displayed dialog box, select a parameter template that uses the same DB engine as the target template and click **OK**.


- If their settings are different, the parameter names and values of both parameter templates are displayed.
- If their settings are the same, no data is displayed.

----End

## Resetting a Parameter Template

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

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

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

**Step 4** On the **Parameter Templates** page, click the **Custom Templates** tab. Locate the parameter template and choose **More > Reset** in the **Operation** column


**Step 5** Click **Yes** to reset all parameters to their default values.


----End



## Modifying the Description of a Parameter Template

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

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

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

**Step 4** On the **Parameter Templates** page, click the **Custom Templates** tab. Locate the parameter template and click  in the **Description** column.

**Step 5** Enter a new description. You can click  to submit or  to cancel the modification.

- After you submit the modification, you can view the new description in the **Description** column.
- The description can contain up to 256 characters, but cannot contain the following special characters: >!<"&'=

### NOTE

You cannot modify the description of any default parameter template.

----End

## Deleting a Parameter Template




### CAUTION

- Deleted parameter templates cannot be recovered. Exercise caution when performing this operation.
  - Default parameter templates cannot be deleted.
- 

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

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

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

**Step 4** On the **Parameter Templates** page, click the **Custom Templates** tab. Locate the parameter template to be deleted and choose **More > Delete** in the **Operation** column.

**Step 5** Click **Yes** to delete it.

----End

# 9 Logs and Auditing

---

## 9.1 Downloading Slow Query Logs of a GaussDB Instance

GaussDB allows you to download slow query logs. Slow query logs help you locate slow SQL statement execution problems.


### Precautions

- CNs and DN of the instance are normal.
- The IaaS network is normal.
- Logs generated in the last 12 hours are collected for the analysis of slow query logs.

### Slow Query Logs

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the navigation pane, choose **Log Analysis**.

**Step 6** The system checks whether there has been a slow query log task in the last 5 minutes and, if there is not, generates one. Click the **Slow Query Logs** tab. Click **Download** in the **Operation** column of the record whose status is **Preparation completed**.

After the log is downloaded, you can analyze the log on the local PC.

----End

## 9.2 Downloading Switchover/Failover Logs of a GaussDB Instance

You can download switchover/failover logs of a GaussDB instance. If switchover/failover log collection is enabled for an instance whose **Failover Priority** is **Availability**, GaussDB can collect Xlogs that cannot be replayed on the standby node in time when a switchover or failover occurs and convert the Xlogs into a SQL file. You can download the SQL file and run SQL statements to replay the data in the SQL file as required.


### Precautions

- CNs and DN of the instance are normal.
- The IaaS network is normal.
- Switchover/Failover logs are available only for distributed instances whose **Failover Priority** is **Availability**.

### Switchover/Failover Logs

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Log Analysis**.

**Step 6** On the displayed page, click the **Switchover/Failover Logs** page, enable **Switchover/Failover Log Collection**, and click **Download** in the **Operation** column of the record whose status is **Preparation completed** to download the switchover/failover log file.

----End

## 9.3 Viewing GaussDB Operation Logs on CTS

With Cloud Trace Service (CTS), you can record operations associated with GaussDB for future query, audit, and backtracking.

## GaussDB Operations That Can Be Recorded by CTS

**Table 9-1** Operations supported by CTS

| Operation                                                     | Resource Type | Trace Name             |
|---------------------------------------------------------------|---------------|------------------------|
| Creating a DB instance or restoring data to a new DB instance | instance      | createInstance         |
| Deleting an instance                                          | instance      | deleteInstance         |
| Changing instance specifications                              | instance      | resizeFlavor           |
| Upgrading the instance version                                | instance      | upgradeVersion         |
| Resetting a password                                          | instance      | resetPassword          |
| Rebooting a DB instance                                       | instance      | instanceRestart        |
| Modifying resource tags                                       | instance      | modifyTag              |
| Deleting resource tags                                        | instance      | deleteTag              |
| Adding resource tags                                          | instance      | createTag              |
| Renaming a DB instance                                        | instance      | renameInstance         |
| Adding nodes                                                  | instance      | instanceAction         |
| Deleting task records                                         | workflowTask  | deleteTaskRecord       |
| Reducing the number of replicas                               | instance      | reduceReplica          |
| Deleting coordinator nodes                                    | instance      | reduceCoordinator-Node |
| Modifying the recycling policy                                | backup        | setRecyclePolicy       |
| Creating a manual backup                                      | backup        | createManualSnapshot   |
| Deleting a manual backup                                      | backup        | deleteManualSnapshot   |
| Modifying the backup policy                                   | backup        | setBackupPolicy        |

| Operation                                          | Resource Type  | Trace Name            |
|----------------------------------------------------|----------------|-----------------------|
| Restoring a DB instance                            | backup         | restoreInstance       |
| Restoring data of an instance using a backup       | instance       | restoreInstance       |
| Changing the retention period of automated backups | instance       | setBackupPolicy       |
| Creating a parameter group                         | parameterGroup | createParameterGroup  |
| Applying a parameter group                         | parameterGroup | applyParameterGroup   |
| Replicating a parameter group                      | parameterGroup | copyParameterGroup    |
| Deleting a parameter group                         | parameterGroup | deleteParameterGroup  |
| Resetting a parameter group                        | parameterGroup | resetParameterGroup   |
| Updating a parameter group                         | parameterGroup | updateParameterGroup  |
| Changing the port                                  | instance       | modifyPort            |
| Creating slow query log download tasks             | instance       | createSlowLogDownload |
| Enabling or disabling switchover/failover logs     | instance       | switchErrorLog        |
| Scaling up storage for shards                      | instance       | resizeVolume          |
| Modifying storage autoscaling policies             | instance       | autoEnlargeVolume     |
| Deleting shards                                    | instance       | reduceShard           |
| Changing standby data nodes to log nodes           | instance       | switchReplica         |
| Performing a primary/standby switchover            | instance       | switchShard           |
| Changing the disk type                             | instance       | changeVolumeType      |

| Operation                                                                        | Resource Type | Trace Name                |
|----------------------------------------------------------------------------------|---------------|---------------------------|
| Starting an instance or node                                                     | instance      | startInstance             |
| Stopping an instance or node                                                     | instance      | stopInstance              |
| Changing a single-replica instance to an instance with primary and standby nodes | instance      | changeDeployment-Solution |
| Enabling advanced features                                                       | instance      | updateFeatures            |

## Querying Audit Logs

You can query GaussDB traces (audit logs) on the CTS console. For details, see [Querying Real-Time Traces](#).

## 9.4 Interconnecting with LTS and Querying Database Audit Logs

### Scenarios

Audit logs capture database start, stop, and connection events, along with DDL, DML, and DCL operations. Each log includes the event time, type, result, username, database, connection details, affected object, instance name, and port number. Security administrators can use these logs to reconstruct events leading to the current database state and identify unauthorized operations, including who performed them, when, and what was done, strengthening traceability and accountability.

If **Upload Audit Logs to LTS** is enabled for a GaussDB instance, new audit logs of it are automatically uploaded to Log Tank Service (LTS) and you can search for logs, monitor logs, download logs, and view real-time logs on the LTS console.


### Precautions

- To apply for the permission needed for enabling **Upload Audit Logs to LTS**, [submit a service ticket](#).
- Currently, this function is available only for centralized instances. The DB engine version must be V2.0-2.1.0 or later.
- Audit logs record all requests sent to your DB instance and are stored in LTS.
- Toggling on or off this function will not be applied immediately. There is a delay of about 10 minutes.
- After this function is enabled, audit policies you configured are reported to LTS by default.

## Enabling Upload Audit Logs to LTS

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

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

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

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

**Step 5** Click the instance name to go to the **Basic Information** page.

**Step 6** In the navigation pane on the left, click **Audit Logs**.

**Step 7** Click  next to **Upload Audit Logs to LTS**.

**Step 8** In the displayed dialog box, configure **Log Group** and **Log Stream**.

**Step 9** Click **OK**.


After this function is enabled, audit logs will not be uploaded immediately to LTS. There is a delay of about 10 minutes. For details, see [Viewing Real-Time Logs](#).

----End

## Disabling Upload Audit Logs to LTS

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


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

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

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

**Step 5** Click the instance name to go to the **Basic Information** page.

**Step 6** In the navigation pane on the left, click **Audit Logs**.

**Step 7** Click  next to **Upload Audit Logs to LTS**.

**Step 8** In the displayed dialog box, confirm the information.

**Step 9** In the displayed dialog box, click **OK**.

----End

# 10 Quota Adjustment

---

## 10.1 Adjusting GaussDB Resource Quotas of an Enterprise Project


The GaussDB management console provides quota management for enterprise projects of tenants.

Quota management is available only for enterprise accounts configured in the whitelist. To apply for the permissions required, .

### Managing 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** Click  in the upper left corner of the page and choose **Databases > GaussDB**.

**Step 4** In the navigation pane on the left, click **Quotas**.

On this page, you can view the usage of instances, vCPUs, memory, and storage under each project.

**Step 5** Locate the enterprise project to be managed, and click **Edit** in the **Operation** column.

**Table 10-1** Parameter description

| Category     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Instances    | <ul style="list-style-type: none"><li>• The first number indicates the number of existing instances in the enterprise project.</li><li>• The second number indicates the maximum number of instances that can be created in the enterprise project. The minimum value must be greater than or equal to the number of existing instances. The maximum value is <b>100000</b>. If this parameter is set to <b>-1</b>, the number is not limited.</li></ul>                                       |
| vCPUs        | <ul style="list-style-type: none"><li>• The first number indicates the number of vCPUs used by existing instances in the enterprise project.</li><li>• The second number indicates the maximum number of vCPUs that can be used by instances in the enterprise project. The minimum value must be greater than or equal to the number of vCPUs used by existing instances. The maximum value is <b>2147483646</b>. If this parameter is set to <b>-1</b>, the number is not limited.</li></ul> |
| Memory (GB)  | <ul style="list-style-type: none"><li>• The first number indicates the size of memory used by existing instances in the enterprise project.</li><li>• The second number indicates the maximum size of memory that can be used by instances in the enterprise project. The minimum value must be greater than or equal to the size of memory used by existing instances. The maximum value is <b>2147483646</b>. If this parameter is set to <b>-1</b>, the number is not limited.</li></ul>    |
| Storage (GB) | <ul style="list-style-type: none"><li>• The first number indicates the storage space used by existing instances in the enterprise project.</li><li>• The second number indicates the maximum storage space that can be used by instances in the enterprise project. The minimum value must be greater than or equal to the storage space used by existing instances. The maximum value is <b>2147483646</b>. If this parameter is set to <b>-1</b>, the number is not limited.</li></ul>       |

 **NOTE**

When you access the **Quotas** page for the first time, **Settings** is displayed.

**Step 6** In the displayed dialog box, enter a new quota. Click **OK**.

**Step 7** On the **Quotas** page, check whether the quotas are changed.

-----**End**

# 11 Disaster Recovery Management

---

## 11.1 Constraints

GaussDB provides the cross-region disaster recovery (DR) capability to improve your data security. If a disaster occurs in the primary data center, the DR data center takes over services quickly. The cross-region DR is suitable for data centers with a distance greater than 200 km.

### Constraints

Constraints on DR relationship establishment:

- Currently, a primary instance can only establish a DR relationship with one DR instance. If you want to re-establish a new DR relationship, you must stop the DR relationship between the primary and original DR instances and then select an available instance to be the DR instance.
- To establish a DR relationship for distributed instances, the number of shards of the primary instance must be the same as that of the DR instance.
- The instances for which a DR relationship is established must use the same OS. DR relationships cannot be established between instances running different OSs, for example, they run on EulerOS and HCE separately.

Constraints during DR relationship establishment:

- The DR instance is only used to synchronize data of the primary instance. It cannot read or write data.
- During DR relationship creation, storage of the primary and DR instances cannot be scaled up.
- For the primary instance, if the instance is forcibly started when a majority of AZs are faulty, data cannot be synchronized to the DR instance. In this case, you need to disconnect the DR relationship.
- If a minority of nodes for the DR instance are faulty, the DR instance still works properly.
- If a majority of nodes for the DR instance are faulty, the DR instance cannot be promoted to primary.

- When you restore the primary instance to a new instance, the DR user account cannot be automatically deleted. You need to delete the account manually.
- The DR instance in a DR relationship does not support adding replicas. After the DR instance is promoted to primary, replicas in the instance can be added.
- The resource types (virtualization types) of the primary and DR instances must be the same.
- It is recommended that the specifications of the DR instance be the same as those of the primary instance. Smaller specifications of the DR instance will slow down Xlog replay on the DR instance.
- Before setting up cross-cloud DR, you need to delegate the DR operation permissions to a system user and connect the two instances properly.
  - No action is required if the two instances are in the same VPC.
  - If the two instances are deployed in different VPCs, create a VPC peering connection on the **Virtual Private Cloud > VPC Peering Connections** page of the VPC Endpoint console and configure routes between subnets.
  - In cross-cloud scenarios, submit a request to Direct Connect for creating a direct connection between the VPCs in the two regions. Specifically, [submit a service ticket](#).
  - Ensure that the security group allows IPv4 inbound traffic from port 12016.
- Compatibility constraints:
  - In the cross-cloud DR scenario, the peer DR instance name is not displayed in the DR relationship list.
- If the pre-verification for a cross-cloud DR operation fails, you can deliver the same operation again on the page after the fault is rectified.

## 11.2 Creating a DR Relationship

This section describes how to create a disaster recovery (DR) relationship.

### Precautions

- Before setting up cross-cloud DR, you need to delegate the DR operation permissions to a system user and connect the two instances properly.
  - No action is required if the two instances are in the same VPC.
  - If the two instances are deployed in different VPCs, create a VPC peering connection on the **Virtual Private Cloud > VPC Peering Connections** page of the VPC Endpoint console and configure routes between subnets.
  - In cross-cloud scenarios, submit a request to Direct Connect for creating a direct connection between the VPCs in the two regions. Specifically, [submit a service ticket](#).
  - Ensure that the security group allows IPv4 inbound traffic from port 12016.
- Before establishing the DR relationship, log in to the console and select the cloud where the primary or DR instance is deployed separately, go to the **Basic Information** page of the instance from the **Instances** page and click **Reset Configuration** for the **DR IP Address** field in the **Configuration** area.

- A DR relationship cannot be established between instances with different CPU architectures.
- The replica consistency protocol and DB engine kernel version of the primary instance must match those of the DR instance.
- After a DR task is created, all databases on the selected DR instance will be cleared. The backup policy will be disabled for the selected DR instance, but historical backup data will be retained. If you want to retain data of the selected DR instance, create a manual backup before the DR task is created.
- After a DR relationship is created, DBS will be automatically authorized to access VPC resource information and query IaaS APIs.
- After a DR relationship is created, the DR network settings and other configurations of the instance will be reset.
- The following table lists the supported DR modes.


**Table 11-1** Dual-instance streaming DR

| Deployment  | Primary Instance                      | DR Instance                                             |
|-------------|---------------------------------------|---------------------------------------------------------|
| Distributed | Combined (Basic edition)              | Combined (Basic edition)                                |
| Centralized | 1 primary + 2 standby                 | 1 primary + 2 standby or Single                         |
|             | 1 primary + 1 standby + 1 log (Paxos) | 1 primary + 1 standby + 1 log (Paxos) or Single (Paxos) |

## Procedure

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

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

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

**Step 4** In the navigation pane on the left, choose **Disaster Recovery**. On the displayed page, click **Create DR Task**.

**Step 5** Configure required parameters described in [Table 11-2](#) and select the confirmation information.

**Table 11-2** Parameter description

| Parameter | Description                                                  |
|-----------|--------------------------------------------------------------|
| DR Type   | Type of the disaster recovery.<br>Streaming DR is supported. |

| Parameter                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Primary Instance                   | Primary instance in the DR relationship. <ul style="list-style-type: none"><li>Only available instances that run V2.0-2.0 or later versions and have more than two replicas are displayed.</li><li>When <b>Replica Consistency Protocol</b> is set to <b>Paxos</b>, only 1 primary + 1 standby + 1 log instances can create DR relationships.</li></ul>                                                                                                                                                                                                                                                                      |
| DR Instance IP Address             | To obtain the IP address of the DR instance, log in to the cloud where the DR instance is deployed, go to the <b>Basic Information</b> page of the instance from the <b>Instances</b> page and view the value of <b>DR IP Address</b> in the <b>Configuration</b> area.<br><br>Before establishing the DR relationship, log in to the console and select the cloud where the primary or DR instance is deployed separately, go to the <b>Basic Information</b> page of the instance from the <b>Instances</b> page and click <b>Reset Configuration</b> for the <b>DR IP Address</b> field in the <b>Configuration</b> area. |
| DR Instance Administrator          | Administrator of the DR instance, which is used to authenticate the primary and DR instances during DR. The default account name is user <b>root</b> of the DR instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| DR Instance Administrator Password | Administrator password of the DR instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

**Step 6** Click **OK** to create a DR relationship. You can view the task status on the **Disaster Recovery** page.

----End


## 11.3 Checking DR Task Statuses

After the DR relationship is established between the primary and DR instances, you can view the data synchronization status on the DR task details page.

### Procedure

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

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

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

**Step 4** Choose **Disaster Recovery** in the navigation pane on the left.

**Step 5** Click the name of the target DR task to go to the task details page.

View the statuses of the primary and DR instances.

- Statuses of the primary instance
  - **Full synchronizing**: The primary instance is fully synchronizing data to the DR instance after the DR relationship is established.
  - **Full synchronization failed**: The primary instance fails to fully synchronize data to the DR instance after the DR relationship is established.
  - **Demoting to standby**: The roles of the primary and DR instances are being switched. In this case, the status of the primary instance is **Demoting to standby**.
  - **Incremental synchronization in progress**: In the DR relationship, the primary instance synchronizes incremental data to the DR instance.
  - **Incremental synchronization failed**: In the DR relationship, the primary instance fails to synchronize incremental data to the DR instance.
  - **RPO** indicates the time, in seconds, from when a transaction (generally the latest transaction) was submitted to when the transaction was sent to the DR instance.  
Description of special RPO values:
    - If the RPO is 0, all transactions have been sent to the DR instance.
    - If the RPO is -1, there are transactions that were not synchronized to the DR instance database or the DR instance is stopped.
- Statuses of the DR instance
  - **Full restoration**: The DR instance is in the full restoration status after DR relationship is established.
  - **Full restoration failed**: The DR instance fails to be fully restored.
  - **Incremental restoration**: The DR instance is synchronizing incremental data from the primary instance.
  - **Incremental restoration failed**: The DR instance fails to synchronize incremental data from the primary instance.
  - **Promoting to primary**: The roles of the primary and DR instances are being switched, or the system is failing over to the DR instance in case the primary instance becomes unavailable. In this case, the status of the DR instance is **Promoting to primary**.
  - **Promotion failed**: The DR instance fails to be promoted to primary.
  - **RTO** measures the amount of data being transmitted.  
Description of special RTO values:
    - If the RTO is 0, all transactions have been transmitted to and successfully executed on the DR instance.
    - If the RTO is -1, there are transactions that were not restored to the DR instance or the DR instance is stopped.

----End

## 11.4 Promoting the DR Instance to Primary

If the primary instance is unavailable, you can manually promote the DR instance to primary.


### Precautions

- If the DR instance is available and in the incremental restoration status, you can promote the DR instance to the primary instance.
- If the DR instance is promoted to the primary instance, it can process read and write requests.
- A DR instance can be promoted to primary regardless of whether the primary instance is available. In the streaming DR scenario, if the primary instance is available, you can stop the DR task on the primary instance after the DR instance is promoted to primary.
- If a DR instance node is faulty before the promotion, repair or replace the faulty node after the promotion is complete.
- In the streaming DR scenario, you can determine whether to enable **Re-establish DR relationship** when promoting the DR instance to primary. For details, see [Re-creating a DR Task After the Primary Instance Is Faulty](#).
- After a DR task is complete, the whitelist configuration set in the **Reset Configuration** dialog box will be automatically cleared.
- After the promotion is complete, the DR task cannot be retried. You need to create a new DR task. After the DR instance is promoted to primary, you need to manually stop the DR task on the original primary instance to clear the DR data.

### Procedure

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

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

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

**Step 4** In the navigation pane on the left, choose **Disaster Recovery**. Locate a DR task, and click **Promote DR Instance to Primary** in the **Operation** column.

**Step 5** Enable **Re-establish DR relationship**.

- If you enable **Re-establish DR relationship**, the system saves the DR relationship record. After the DR instance is promoted to primary, you can re-establish the DR relationship by one click.
- If you disable this option, the DR relationship will be disconnected and cannot be re-established by one click.

**Step 6** In the displayed dialog box, select **Promote DR instance to primary** and click **Yes**.  
The DR instance becomes the primary instance.

----End

## 11.5 Stopping a DR Task


This section describes how to stop the DR relationship between the primary and DR instances.


### Precautions

- After the DR task is stopped, you need to create a new DR task if needed.
- After the DR task is complete, the network configuration will be automatically cleared. If you need to set up a new DR relationship, go to the instance details page to reset the configuration.
- In the streaming DR scenario, if instance B is promoted to the primary instance and removes the DR relationship with instance A, instance A does not know that the relationship has been deleted and can successfully remove the DR relationship with instance B only when instance B is available. If instance B is abnormal, you need to manually skip the removal step related to instance B when deleting the relationship on the instance A side.

### Procedure

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

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

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

**Step 4** In the navigation pane on the left, choose **Disaster Recovery**. Locate a DR task and click **Stop** in the **Operation** column.

**Step 5** In the displayed dialog box, select **Stop task** and click **Yes**.


----End


## 11.6 Deleting a DR Task

After a DR task is stopped or fails to be created, you can delete the DR task record on the console.

### Procedure

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

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

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

**Step 4** In the navigation pane on the left, choose **Disaster Recovery**. Locate a stopped DR task and click **Delete** in the **Operation** column.

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

----End

## 11.7 Re-creating a DR Task After the Primary Instance Is Faulty

### Scenarios

A DR relationship has been established with instance A as the primary instance and instance B as the DR instance. If primary instance A is faulty, you need to manually promote instance B to primary.

After instance A is recovered, re-establish the DR relationship between instance B and instance A, and then perform a DR switchover to promote instance A to the primary instance.

### Procedure

**Step 1** If the primary instance A is faulty, log in to the console, select the region where the DR instance B is deployed, promote instance B to primary by referring to [Promoting the DR Instance to Primary](#), and enable **Re-establish DR Relationship**.

In this case, instance B takes over the role of instance A to provide services.

**Step 2** Log in to the console, select the region where instance B is located, and re-establish a DR relationship by referring to [Re-establishing a DR Relationship](#).

**Step 3** After the fault of instance A is rectified, switch roles of instance A and instance B by referring to [Switching Roles of Primary and DR Instances](#).

In this case, the services are taken over by instance A and the DR relationship is restored to the status before the fault occurs.

After the switchover, instance A provides services. You need to switch services to instance A.

----End

## 11.8 Switching Roles of Primary and DR Instances

### Scenarios

If the primary and DR regions and primary and DR instances are normal, you can log in to either the primary instance or DR instance to switch their roles. This section describes how to switch roles of primary and DR instances.

### Procedure

**Step 1** Log in to the console and select the region where the primary or DR instance is located.

**Step 2** In the navigation pane on the left, choose **Disaster Recovery**. Locate a DR task, and click **Switch Roles** in the **Operation** column.

**Step 3** In the dialog box that is displayed, confirm related information.

**Step 4** Select **Switch roles** and click **Yes**.

----End

## 11.9 Re-establishing a DR Relationship

### Scenarios

After the DR instance is promoted to primary and the original primary instance restores, you can re-establish the DR relationship between the two instances. If the DR relationship is re-established at the DR instance side, the original primary instance becomes the DR instance, and the original DR instance becomes primary.

### Precautions

- Streaming DR is available only for instances whose DB engine version is V2.0-3.200 or later.

### Procedure

**Step 1** Log in to the console and select the region where the new primary instance is located.

**Step 2** In the navigation pane on the left, choose **Disaster Recovery**. Locate the DR task and click **Re-establish DR Relationship** in the **Operation** column.

**Step 3** In the displayed dialog box, select the confirmation information, and click **Yes**.

If the DR relationship is re-established at the DR instance side, the original primary instance becomes the DR instance, and the original DR instance becomes primary.

----End

## 11.10 Performing a DR Drill

### Scenarios

If the primary and DR regions and primary and DR instances are normal, you can log in to the DR instance and perform a DR drill. This section describes the precautions and how to perform a DR drill.

### Precautions

- This function is supported only when streaming DR is used and the DB engine version is V2.0-3.200 or later.
- After a DR drill starts, data synchronization between the primary and DR instances is suspended. Automated backups for the DR instance are created.

- When a DR drill is being performed on the DR instance, do not remove the DR relationship from the primary instance side.
- DR drills are supported for dual-instance streaming DR only.

## Procedure

**Step 1** Log in to the console and select the region where the DR instance is located.

**Step 2** In the navigation pane on the left, choose **Disaster Recovery**. Click **More** in the **Operation** column of the target DR task, and choose **Enable DR Drill**.

**Log Cache Space/Primary Instance Space:** During the DR drill, the DR instance provides services and the incremental data is temporarily cached in the primary instance. After the drill is complete, the DR instance can perform incremental replay based on the temporarily stored incremental data. If the cached incremental data exceeds the value of this parameter, the DR instance will replay all data after the drill is complete.

**Step 3** After the DR drill is enabled, perform DR drill operations on the DR instance.

**Step 4** After the drill is complete, log in to the console and select the region where the DR instance is located. On the **Disaster Recovery** page, click **More** in the **Operation** column of the target DR task, and choose **Disable DR Drill**.

**Step 5** After the DR drill is complete, view that the DR task is restored to the status before the DR drill is performed. The data generated by the primary instance during the DR drill is synchronized to the DR instance, and the data generated by the DR instance during the DR drill is cleared.

----End

## 11.11 Enabling or Disabling Log Cache

### Scenarios

If the primary and standby regions are disconnected for a long time due to network problems, you can enable log cache to temporarily store incremental data in the primary instance storage space. When the network recovers, the incremental data, instead of full data, is synchronized to the DR instance.

### Precautions

- This function is supported only when streaming DR is used and the DB engine version is V2.0-3.200 or later.
- During the log caching period in the cross-cloud DR scenario, the DR relationship cannot be removed.

## Procedure

**Step 1** Log in to the console and select the region where the DR instance is located.

**Step 2** In the navigation pane on the left, choose **Disaster Recovery**. Click **More** in the **Operation** column of the target DR task, and choose **Enable Log Cache**.

- Step 3** After the log cache is enabled, the task status of the primary instance changes to **Caching logs**. The incremental data temporarily stores in the primary instance storage space.
- Step 4** After the DR relationship is normal, log in to the console, select the region where the primary instance is located, and click **Disable Log Cache** in the **Operation** column.
- Step 5** After log cache is disabled, check that the DR task is restored to the status before the log cache is enabled. The data generated during the log cache on the primary instance is synchronized to the DR instance.
- End

## 11.12 Restrictions on the DR Instance

After a DB instance serves as a DR instance, some basic functions are unavailable and some monitoring metrics cannot be reported. This section describes the restrictions on the DR instance.

### Restrictions on the Primary Instance

After a DR relationship is established, the DR instance only allows for storage scaling, specifications change, rebooting, and parameter modification. The basic functions of the primary instance have the following restrictions:

| Function                                                          | Supported by the Primary Instance During DR |
|-------------------------------------------------------------------|---------------------------------------------|
| Rebooting an instance                                             | Yes                                         |
| Modifying parameters                                              | Yes                                         |
| Applying parameters                                               | Yes                                         |
| Resetting a password                                              | Yes                                         |
| Performing a full backup                                          | Yes                                         |
| Creating a differential backup                                    | Yes                                         |
| Deleting a backup                                                 | Yes                                         |
| Modifying the backup policy                                       | Yes                                         |
| Restoring data to the original instance                           | No                                          |
| Restoring data to a new instance                                  | Yes                                         |
| Scaling up storage space                                          | Yes                                         |
| Adding nodes (only for distributed instances)                     | No                                          |
| Changing instance specifications (only for centralized instances) | Yes                                         |

| Function                                           | Supported by the Primary Instance During DR |
|----------------------------------------------------|---------------------------------------------|
| Upgrading a minor version                          | Yes                                         |
| Viewing slow query logs                            | Yes                                         |
| Viewing metrics                                    | Yes                                         |
| Deleting an instance                               | No                                          |
| Deleting an automated backup                       | N/A                                         |
| Rebuilding a deleted instance from the recycle bin | N/A                                         |
| Backing up and restoring data in a single table    | Yes                                         |
| Querying the disk usage                            | Yes                                         |
| Creating a database                                | Yes                                         |
| Querying a database                                | Yes                                         |
| Creating a schema and user                         | Yes                                         |
| Deleting a schema and user                         | Yes                                         |
| Performing database operations                     | Yes                                         |
| Repairing a node                                   | Yes                                         |
| Replacing a node                                   | Yes                                         |
| Forcibly starting a minority of AZs                | No                                          |
| Resuming AZs                                       | No                                          |
| Switching AZs                                      | Yes                                         |
| Resetting configurations (for DR)                  | Yes                                         |
| Adding replicas                                    | No                                          |
| Managing tags                                      | Yes                                         |
| Binding an EIP                                     | Yes                                         |

## Restrictions on Monitoring Metrics

After the DR relationship is established, some monitoring metrics of the DR instance are unavailable. The following table lists the monitoring metrics of the DR instance.

 **CAUTION**

If the deployment model of an instance is **1 primary + 1 standby + 1 log**, the log node does not have a data disk, and the following node-level metrics related to data disks cannot be displayed for the node:

- Data Disk IOPS
- Data Disk Write Throughput
- Data Disk Read Throughput
- Time Required for per Data Disk Write
- Time Required for per Data Disk Read

| Category       | Sub-category | Metric                                | Displayed for Distributed DR Instance | Displayed for Centralized DR Instance |
|----------------|--------------|---------------------------------------|---------------------------------------|---------------------------------------|
| Instance level | -            | Used Instance Disk Size               | Yes                                   | No                                    |
|                |              | Total Instance Disk Size              | Yes                                   | No                                    |
|                |              | Instance Disk Usage                   | Yes                                   | No                                    |
|                |              | Deadlocks                             | No                                    | No                                    |
|                |              | Response Time of 80% SQL Statements   | No                                    | No                                    |
|                |              | Response Time of 95% SQL Statements   | No                                    | No                                    |
|                |              | Buffer Hit Rate                       | No                                    | No                                    |
| Node level     | -            | CPU Usage                             | Yes                                   | Yes                                   |
|                |              | Memory Usage                          | Yes                                   | Yes                                   |
|                |              | Data Write Volume                     | Yes                                   | Yes                                   |
|                |              | Outgoing Data Volume                  | Yes                                   | Yes                                   |
|                |              | Disk IOPS                             | Yes                                   | Yes                                   |
|                |              | Disk Write Throughput                 | Yes                                   | Yes                                   |
|                |              | Disk Read Throughput                  | Yes                                   | Yes                                   |
|                |              | Time Required for per Data Disk Write | Yes                                   | Yes                                   |
|                |              | Time Required for per Data Disk Read  | Yes                                   | Yes                                   |
|                |              | Swap Memory Usage                     | Yes                                   | Yes                                   |

| Category | Sub-category | Metric                                                                           | Displayed for Distributed DR Instance | Displayed for Centralized DR Instance |
|----------|--------------|----------------------------------------------------------------------------------|---------------------------------------|---------------------------------------|
|          |              | Total Swap Memory                                                                | Yes                                   | Yes                                   |
|          |              | IOPS Usage (This metric applies only to nodes that use EVS disks.)               | Yes                                   | Yes                                   |
|          |              | Disk I/O Bandwidth Usage (This metric applies only to nodes that use EVS disks.) | Yes                                   | Yes                                   |

| Category        | Sub-category | Metric                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Displayed for Distributed DR Instance | Displayed for Centralized DR Instance |
|-----------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|---------------------------------------|
| Component level | CN           | User Logins, User Logouts, Lock Waiting Session Rate, Active Session Rate, CN Connections, User Committed Transactions, User Rollback Transactions, Background Committed Transactions, Background Rollback Transactions, Average Response Time of User Transactions, User Transaction Rollback Rate, Background Transaction Rollback Rate, Data Definition Language, Data Manipulation Language, Data Control Language, DDL and DCL Rate, Data Volume to Be Flushed to Disks, Physical Reads per Second, Physical Writes per Second, Online Sessions, Active Sessions, Online Session Rate, Used Dynamic Memory, Dynamic Memory Usage, Maximum Execution Duration of Database Transactions, Idle Transactions, Slow SQL Statements in the System Database, Slow SQL Statements in the User Database, and Xlog Rate | No                                    | N/A                                   |
|                 | Primary DN   | Used Disk Size, Total Disk Size, Disk Usage, Data Volume to Be Flushed to Disks, Physical Reads per Second, and Physical Writes per Second                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | N/A                                   | N/A                                   |

| Category | Sub-category | Metric                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Displayed for Distributed DR Instance | Displayed for Centralized DR Instance |
|----------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|---------------------------------------|
|          |              | Standby Redo Progress, User Logins, User Logouts, Lock Waiting Session Rate, Active Session Rate, User Committed Transactions, User Rollback Transactions, Background Committed Transactions, Background Rollback Transactions, Average Response Time of User Transactions, User Transaction Rollback Rate, Background Transaction Rollback Rate, Data Definition Language, Data Manipulation Language, Data Control Language, DDL and DCL Rate, Online Sessions, Active Sessions, Online Session Rate, Used Dynamic Memory, Dynamic Memory Usage, Size of WALs Retained in the Replication Slot, and Xlog Rate | N/A                                   | N/A                                   |
|          | Standby DN   | Used Disk Size                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Yes                                   | Yes                                   |
|          |              | Total Disk Size                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Yes                                   | Yes                                   |
|          |              | Disk Usage                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Yes                                   | Yes                                   |
|          |              | Primary Node Flow Control Duration                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | No                                    | No                                    |
|          |              | Standby Node RTO                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | No                                    | No                                    |
|          |              | standby_xlog_replay_speed                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | No                                    | No                                    |
|          |              | Difference Between Redo and Receipt Positions on Standby Node                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | No                                    | No                                    |

| Category | Sub - category | Metric                     | Displayed for Distributed DR Instance | Displayed for Centralized DR Instance |
|----------|----------------|----------------------------|---------------------------------------|---------------------------------------|
|          |                | Standby Node Redo Progress | No                                    | N/A                                   |
|          |                | Lock-Waiting Session Rate  | No                                    | No                                    |
|          |                | Active Session Rate        | No                                    | No                                    |
|          |                | Online Session Rate        | No                                    | No                                    |
|          |                | Online Sessions            | No                                    | No                                    |
|          |                | Dynamic Memory Usage       | No                                    | No                                    |

# 12 Managing GaussDB Tasks

You can view the progresses and results of tasks on the **Task Center** page.

## Precautions

Deleted task records cannot be recovered. Exercise caution when performing this operation.

## Tasks That Can Be Viewed and Managed


**Table 12-1** Supported task types

| Category               | Task                                                                                                                                                                           |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Instance creation      | Creating a GaussDB instance                                                                                                                                                    |
| Instance upgrade       | <ul style="list-style-type: none"><li>• Rolling upgrade</li><li>• Upgrade commit</li><li>• Upgrade auto-commit</li><li>• In-place upgrade</li><li>• Upgrade rollback</li></ul> |
| Backup and restoration | <ul style="list-style-type: none"><li>• Creating a manual backup</li><li>• Restoring data to a new instance</li></ul>                                                          |

## Viewing a Task

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

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

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

**Step 4** Choose **Task Center** in the navigation pane on the left. On the displayed page, view the task details.

- To identify a task, you can use the task name/ID or instance name/ID, or simply select a task name in the search box displayed in the upper pane of the page.
- You can view the progress and status of tasks in a specific period. The default period is seven days.  
The task list can only show up to 30 days of past tasks.
- You can view tasks in the following statuses:
  - Running
  - Completed
  - Failed
- You can view the task creation and completion time.


----End

## Deleting a Task Record

You can delete the task records that no longer need to be displayed. The deletion only deletes the task records, and does not delete the DB instances or terminate the tasks that are being executed.

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

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

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

**Step 4** Choose **Task Center** in the navigation pane on the left. On the displayed page, locate the task record to be deleted and click **Delete** in the **Operation** column. In the displayed dialog box, click **OK**.

You can delete tasks in the following statuses:

- Completed
- Failed

----End

# 13 Managing GaussDB 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.


You can use unified tags to manage various resources and use custom tags to define different DB instance categories. This section describes how to edit and delete tags.


## Precautions

- You are advised to configure predefined tags on the TMS console.
- A tag consists of a key and value. You can add only one value for each key.
- A maximum of 20 tags can be added for an instance.

## Adding Tags

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the navigation pane, choose **Tags**. On the displayed page, click **Add Tag**. In the displayed dialog box, enter a tag key and value, and click **OK**.

- A tag key can contain up to 128 characters. It cannot start with **\_sys\_** or a space, and cannot end with a space. Only letters, digits, spaces, and the following special characters are allowed: `_./=+-@`
- A tag value can contain up to 255 characters. Only letters, digits, spaces, and the following special characters are allowed: `_./=+-@`


**Step 6** View and manage the tag on the **Tags** page.

----End

## Editing a Tag

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, click **Tags**, locate the tag to be edited, and click **Edit** in the **Operation** column.

**Step 6** In the displayed dialog box, enter a tag value and click **OK**.

- Only the tag value can be edited.
- The tag value can contain 0 to 43 characters. Only letters, digits, at signs (@), hyphens (-), underscores (\_), and periods (.) are allowed.


**Step 7** View and manage the tag on the **Tags** page.

----End

## Deleting a Tag

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

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

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

**Step 4** On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Tags**. On the displayed page, locate the tag to be deleted, and click **Delete**.

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

**Figure 13-1** Deleting a tag



**Step 7** Check that the tag is no longer displayed on the **Tags** page.

----End

# 14

## Resetting the Administrator Password of a GaussDB Instance

---

### Scenarios

If you forget the password of your **root** account when using GaussDB, 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 the DB instance is abnormal, the administrator password cannot be reset.
- The volume of data being processed by the instance determines how long it takes for the new password to take effect.
- To prevent brute force cracking and ensure system security, change your password periodically.
- If you log in to the database as the **root** user, resetting the password may interrupt services. Exercise caution when performing this operation.

### Procedure

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

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

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

**Step 4** On the **Instances** page, locate the instance that you want to reset password for and click **More > Reset Password** in the **Operation** column.

Alternatively, click the instance name to go to the **Basic Information** page. In the **DB Information** area, click **Reset Password** next to the **Administrator** field.

**Step 5** Enter a new password and confirm the password.

The new password must:

- Contain 8 to 32 characters.
- Contain at least three types of the following: uppercase letters, lowercase letters, digits, and special characters. Supported special characters: ~!@#%^\*\_-=+?,
- Be different from the old password or the old password written backwards.

----End