GaussDB

FAQs

Issue	01
Date	2025-02-18

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

11 Log Management..38

12 Network Security...40

1 About GaussDB

1.1 What Should I Pay Attention to When Using GaussDB?

1. DB instance OSs are invisible to you. Your applications can access a database only through an IP address and a port.

2. The backup files stored in OBS and the ECS used by GaussDB are invisible to you. They are visible only in the GaussDB instance management system.

3. When viewing the instance list, ensure that the current region is the same as the region where the instance is purchased.

4. After purchasing GaussDB instances, you do not need to perform basic database O&M operations, such as applying HA and security patches. However, you must still pay attention to whether:

   a. The CPU, IOPS, and storage space are abundant for the GaussDB instances. If any of these becomes insufficient, you will have to change the CPU/memory specifications or scale up the storage.

   b. The performance of the GaussDB instance is adequate, a large number of slow SQL queries exist, SQL statements need to be optimized, or if any indexes are redundant or missing.

1.2 What Is the Availability of GaussDB Instances?

Formula for the GaussDB instance availability:

DB instance availability = (1 – Failure duration/Total service duration) × 100%

1.3 Will My GaussDB Instances Be Affected by Other Users' Instances?

No. Your GaussDB instances and resources are isolated from other users' GaussDB instances.

1.4 What Can I Do About Slow Responses of Websites When They Use GaussDB?

To solve this problem:

- You can check the performance of GaussDB on the management console.
- Compare the database connection status of local databases and GaussDB instances. This problem depends on web applications.

1.5 Can Multiple ECSs Connect to the Same GaussDB Instance?

Within the limits of database capability, multiple ECSs can connect to the same GaussDB instance.

1.6 Can I Encrypt the Disk After Purchasing a GaussDB Instance?

You can enable disk encryption when purchasing a GaussDB instance. For details, see **Buying a DB Instance**.

After an instance is created, the disk encryption status and the key cannot be changed.

If you want to configure disk encryption for a created instance, perform the following operations:

- Restore instance backup to a new instance and enable disk encryption for the new instance.
- Use the Data Replication Service (DRS) migration function to migrate instance data to other encrypted instances.

1.7 Will Backups Be Encrypted After Disk Encryption Is Enabled for My GaussDB Instance?

If you enable disk encryption during instance creation, the disk encryption status and the key cannot be changed later. Disk encryption will not encrypt backup data stored in OBS. To enable backup data encryption, contact customer service.

> **NOTICE**
>
> If disk encryption or backup data encryption is enabled, keep the key properly. Once the key is disabled, deleted, or frozen, the database will be unavailable and data may not be restored.
>
> - If disk encryption is enabled but backup data encryption is not enabled, you can **restore instance backup data to a new instance**
> - If both disk encryption and backup data encryption are enabled, data cannot be restored.

1.8 Will Different GaussDB Instances Share CPU and Memory Resources?

No, GaussDB instances are independent from each other and will not share CPU and memory resources. These resources are the instance specifications selected when you buy an instance.

1.9 How Many Databases Can Run on a GaussDB Instance?

The maximum number of databases that can run on a GaussDB instance depends on the DB engine settings.

If there are enough CPU, memory, and storage resources, there are no limitations to the number of databases running on a DB instance.

1.10 What Is the Relationship Between GaussDB and PostgreSQL?

Based on the Postgres-XC architecture of open-source PostgreSQL 9.2, the earliest GaussDB kernel developed a multi-CN architecture, on which some major features, such as distributed execution framework (stream operators) and vectorized engine, are provided. Currently, GaussDB only uses standard APIs and common functions of PostgreSQL. It has developed its own ecosystem, architecture, and key technologies. Its centralized instances are open-source, and the storage engine and optimizer have been rearchitected. The differences between GaussDB and PostgreSQL are as follows:

- PostgreSQL uses a process model, and GaussDB uses a thread pool model.
- PostgreSQL supports only row-store. GaussDB supports row-store, column-store, and Ustore.
- PostgreSQL supports only centralized deployment. GaussDB supports both centralized and distributed deployment.
- Compared with PostgreSQL, GaussDB has many unique features, such as dynamic data masking, fully-encrypted processing, anti-tampering, GTM-Lite mode, NUMA-aware architecture, and geo-redundant and intra-city dual-cluster DR solutions.

1.11 Can Multiple Users Log In to a GaussDB Instance Through DAS at the Same Time? Will the Accounts Be Locked If I Enter Wrong Passwords Several Times in a Row?

Multiple users can log in to DAS at the same time. The passwords will not be locked after multiple failed login attempts.

If you forget the password of your database account when using GaussDB, you can reset the password. On the **Instances** page, click **More** in the **Operation** column of the target instance and choose **Reset Password**.

1.12 What Are the Differences Between the Independent Deployment Model and Combined Deployment Model of GaussDB Distributed Instances?

- Distributed (independent deployment)
  - This model is suitable for scenarios where there is a need to process a large amount of data, improve the storage capacity and concurrency of instances through scale-out, and the performance requirements cannot be met by deploying a distributed instance in combined deployment.
  - Each component is deployed on its own ECS, guaranteeing uninterrupted performance with exclusive resources. For example, in the default configuration (3 CNs, 3 replicas, and 3 shards), a total of 18 ECSs are required, including three CNs, nine DNs, three CMSs, and three GTMs.
- Distributed (combined deployment)
  - This model is suitable for scenarios with large data volumes, potential future data volume growth, moderate performance requirements, and a need for a distributed database.
  - All components are deployed on the same ECS. For example, in a combined deployment with one primary node and two standby nodes, only three ECSs are required. In complex service scenarios, performance contention may arise as all components share the same ECS. Consequently, the overall performance of combined deployment is slightly lower than that of independent deployment with the same specifications.

1.13 What Are the Differences Between the General-Purpose, Dedicated, and Kunpeng Dedicated Specification Types Provided GaussDB?

GaussDB instances support the x86 or Arm architecture, and the specification types supported by different architectures are as follows. For details about the differences between them, see **Table 1-1**.

- x86: Dedicated (1:4), Dedicated (1:8), and General-purpose (1:4). 1:4 and 1:8 refer to the vCPU-to-memory ratio.
- Arm: Kunpeng dedicated (1:4) and Kunpeng dedicated (1:8)

Table 1-1 Differences between the general-purpose, dedicated, and Kunpeng dedicated specification types

Instance Specifications	Supported Instance Type	Description
General-purpose	Basic edition	This type of specifications uses general-computing ECSs, which prioritize resource sharing and offer better cost-effectiveness. However, if the database is heavily loaded, it can lead to resource contention and fluctuations in computing performance. Consequently, this specification type is not suitable for large applications, and it is not offered in the enterprise edition.
Dedicated	• Basic edition • Enterprise edition	This type of specifications uses general computing-plus x86 ECSs. Such ECSs ensure that there is no resource contention between instances and provide stable performance.
Kunpeng dedicated	• Basic edition • Enterprise edition	This type of specifications uses general computing-plus Kunpeng ECSs. Such ECSs ensure that there is no resource contention between instances and provide stable performance. This specification type is mainly used in IT application innovation scenarios.

1.14 Can I Switch Between the GaussDB Enterprise Edition and Basic Edition?

Switching between the enterprise edition and basic edition of GaussDB is not supported. This is because there are parameter differences caused by unsubscription, re-subscription, and advanced feature changes, which could potentially lead to data loss. You can create another instance with the required edition type and migrate data to the new instance using DRS. For details, see **Real-Time Synchronization**.

2 GaussDB Resource Freezing, Unfreezing, Release, Deletion, and Unsubscription

Will Pay-per-Use Instances Be Charged When They Are Not Used?

Yes. The pay-per-use billing mode is a postpaid mode in which your DB instance will be billed based on the usage duration. The billing starts from the time when your instance is purchased to the time when it is deleted.

Why Are My Resources Released?

If your subscriptions have expired but not been renewed, or you are in arrears due to insufficient balance, your resources enter a grace period. If you still do not complete the payment or renewal after the grace period expires, you will enter a retention period. During the retention period, the resources are not available. If the renewal is still not completed or the outstanding amount is still not paid off when the retention period ends, the stored data will be deleted and the cloud service resources will be released. For details, see **Service Suspension and Resource Release**.

Why Are My Resources Frozen?

Your resources may be frozen for a variety of reasons. The most common reason is that you are in arrears.

Can I Still Back Up Data If My DB Instance Is Frozen?

No. If your instance is frozen due to arrears, you need to unfreeze the instance first.

How Do I Unfreeze My Resources?

Frozen due to arrears: You can renew your resources or top up your account. Instances frozen due to arrears can be renewed, released, or deleted. Yearly/Monthly instances that have expired cannot be unsubscribed from, while those that have not expired can be unsubscribed from.

What Happens When My Resources Are Frozen, Unfrozen, or Released?

- After your resources are frozen:
  - They cannot be accessed, causing downtime. For example, if your instance is frozen, it cannot be connected.
  - If they are yearly/monthly resources, no changes can be made to them.
  - They can be unsubscribed from or deleted manually.
- After your resources are unfrozen, you can connect to them again.
- If your resources are released, your instances will be deleted. Before the deletion, GaussDB determines whether to **move the instance to the recycle bin** based on the recycling policy you specified.

How Do I Renew My Resources?

After a yearly/monthly instance expires, you can renew it on the **Renewals** page. For details, see **Renewal Management**.

Can My Resources Be Recovered After Being Released? /Can I Retrieve an Incorrect Unsubscription?

If your instance is moved to the recycle bin after being deleted and is within the retention period, you can **rebuild it from the recycle bin**. Otherwise, data cannot be restored.

When you unsubscribe from an instance, confirm the instance information carefully. If you have unsubscribed from a DB instance by mistake, you are advised to purchase a new one.

How Do I Delete My Instance?

- For pay-per-use instances, see **Deleting a DB Instance**.
- For yearly/monthly instances, see **Unsubscribing a Yearly/Monthly DB Instance**.

3 Resource and Disk Management

3.1 Can I Scale Down the Storage Space of My GaussDB Instances?

No. You can create another instance with a smaller storage space and migrate data to the new instance using DRS. For details, see **Real-Time Synchronization**.

3.2 Which Items Occupy the Storage Space of My GaussDB Instances?

Both your common data (excluding backup data) and the data required for the operation of your DB instances (such as system database data, rollback logs, redo logs, and indexes) occupy the storage space on your purchased GaussDB instances. The storage space includes the file system overhead required for inodes, reserved blocks, and database operations. It also includes the log files generated by the GaussDB database.

3.3 How Much Storage Space Is Required for DDL Operations?

Data Definition Language (DDL) operations may increase storage usage sharply. To ensure that services are running properly, do not perform DDL operations during peak hours. If DDL operations are required, ensure that the storage space is at least twice the tablespace size plus 10 GB. For example, if your tablespace is 500 GB, ensure that the storage space is at least 1,010 GB (500 GB x 2 + 10 GB).

4 Database Connections

4.1 What Should I Do If I Can't Connect to My GaussDB Instance?

Problem Analysis

Locate the fault by considering the following aspects:

1. **Check whether the DB instance is available.**

    For example, the system is faulty, the DB instance is abnormal, or the DB instance or a table is locked.

2. **(Common) Check whether the client connection is correct.**
    –   If you connect to an instance over a private network, ensure that the instance and ECS are in the same region and VPC.
    –   If you connect to an instance over a public network, bind an EIP to the instance and then connect to the instance through the EIP. If no EIP is available, buy one first.

3. **Check whether the parameters in the connection command are correct.**

    Check whether the following parameters are configured correctly: connection address, port number, username, and password.

4. **(Common) Check whether the network connectivity is normal.**

    For a private network connection:

    a. Ensure that the ECS and GaussDB instance are in the same region and VPC.

    b. Check security group rules.

        To access GaussDB instances in a security group from a public network, **add an inbound rule**.

    c. On the ECS, check whether the GaussDB instance port can be connected.

    For a public network connection:

    a. Check security group rules.

To access GaussDB instances in a security group from a public network, **add an inbound rule**.

b. Check network ACL rules.

c. Ping the ECS to the instance in the same region.

5. **(Common) Check whether the DB instance is in the Storage full state.**

If the DB instance is in the **Storage full** state, data read and write performance is affected.

6. **View common connection error messages.**

Find corresponding solutions based on connection error messages.

Troubleshooting Process

Figure 4-1 Locating instance connection failures
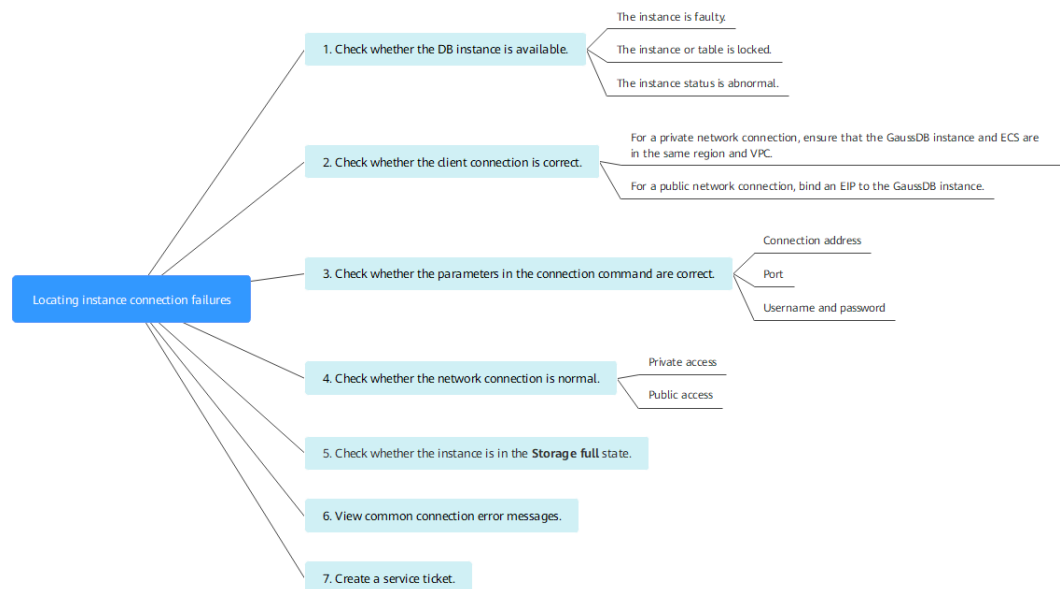
1. **Check whether the DB instance is available.**

Check method: Check whether the DB instance is in the **Available** state.

Possible cause: If the system is faulty, the DB instance is abnormal, or the DB instance or a table is locked.

Solution: If the DB instance is abnormal, reboot it.

Figure 4-2 Checking DB instance status
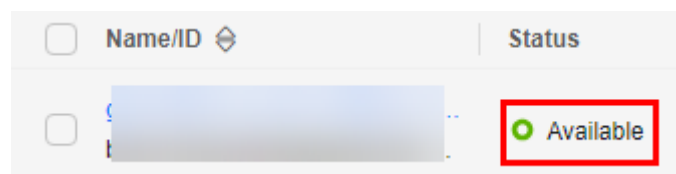
2. **Check whether the client connection is correct.**

For details about how to connect to a DB instance over a private or public network, see **Can an External Server Access GaussDB Instances?**

Table 4-1 Connection methods

Connect Through	Description	Example
Private network	A private IP address is provided by default. If your applications are deployed on an ECS that is in the same region and VPC as your DB instance, use a private IP address to connect to the ECS and the DB instance.	For example, to connect to the **postgres** database, run the following command: **gsql -d** *postgres* **-h** *10.0.0.0* **-U** *root* **-p** *8000* NOTE *postgres* indicates the database name, *10.0.0.0* indicates the CN IP address of a distributed DB instance or the primary DN IP address of a centralized DB instance, *root* is the username for logging in to the database, and *8000* is the port number of the CN.
Public network	If you cannot access a GaussDB instance over a private IP address, bind an EIP to the GaussDB instance and connect it to the ECS (or a host from the public network) over the EIP.	For example, to connect to the **postgres** database, run the following command: **gsql -d** *postgres* **-h** *10.0.0.0* **-U** *root* **-p** *8000* NOTE **postgres** is the name of the database to be connected; **10.0.0.0** is the EIP bound to the DB instance; **root** is the username for logging in to the database; **8000** is the port number of the DB instance.

3. **Check whether the parameters in the connection command are correct.**

   Ensure that the connection address, port, username, and password are correct, and try to connect to the instance again.

   Connection through a private network

   a. Connection address

      Private IP Address in the **Node List** area of the **Basic Information** page of the target instance

   b. Database port

      Database Port of the target instance

   c. Username and password

      User **root** and its password

   Connection through a public network

    a.   Connection address

       EIP in the **Node List** area of the **Basic Information** page of the target instance

    b.   Database port

       Database Port of the target instance

    c.   Username and password

       User **root** and its password

4. **Check the network connection.**

   Private network connection

   a.   Check whether the ECS and GaussDB instance are in the same region and VPC.

      ■ The ECSs and instances in different regions cannot communicate with each other. To reduce network latency, deploy your DB instance in the region nearest to your workloads.

      ■ Perform the following operations in different VPCs:

          ○ Change the VPC hosting the ECS to the same as that hosting the GaussDB instance.

          ○ Create a VPC peering connection to set up a network connection between the two VPCs.

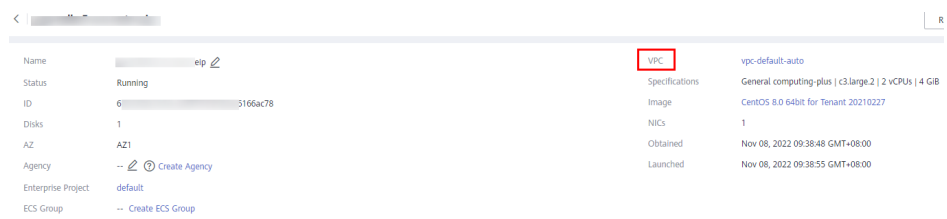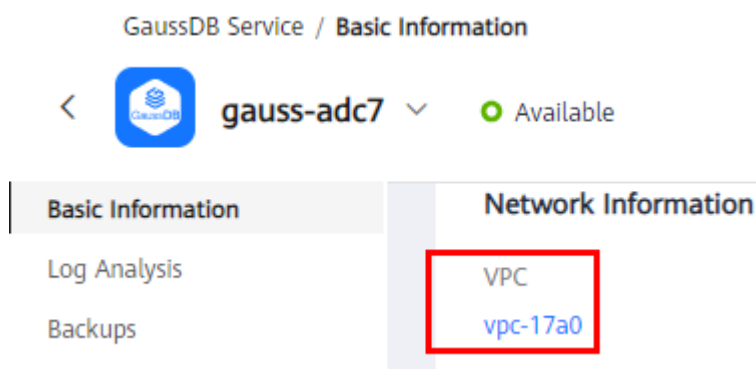**Figure 4-3** Viewing the VPC of an ECS

Figure 4-4 Viewing the VPC of a GaussDB instance

   b.   Check security group rules.

      ■ If **Destination** is not **0.0.0.0/0** and **Protocol & Port** is not **All** on the **Outbound Rules** page of the ECS, add the private IP address and port of the GaussDB instance to the outbound rules.

- If **Source** is not **0.0.0.0/0** and **Protocol & Port** is not **All** on the **Inbound Rules** page of the GaussDB instance, add the IP address and port of the ECS to the **inbound rules**.

c. On the ECS, check whether the GaussDB instance port can be connected.

   telnet *<IP address> <port number>*

   - If the ECS can connect to the DB instance, the network between them is normal.

   - If the ECS cannot connect to the DB instance, choose **Service Tickets > Create Service Ticket** in the upper right corner of the management console to submit a service ticket for help.

Public network connection

a. Check security group rules.

   - If **Destination** is not **0.0.0.0/0** and **Protocol & Port** is not **All** on the **Outbound Rules** page of the ECS, add the public IP address and port of the GaussDB instance to the outbound rules.

   - If **Source** is not **0.0.0.0/0** and **Protocol & Port** is not **All** on the **Inbound Rules** page of the GaussDB instance, add the IP address and port of the ECS to the **inbound rules**.

b. Check network ACL rules.

   i.   Go to the **Network ACLs** page.

   ii. Check whether the NIC bound to the EIP is in the subnet associated with the network ACL.

   iii. Check whether the network ACL is enabled.

        If it is enabled, **add an ICMP rule to allow traffic**.

   The default network ACL rule denies all inbound and outbound packets. This default rule is still applied even if the network ACL is disabled.

c. Ping the ECS to the instance in the same region.

   If you cannot ping the EIP on the original ECS, select another ECS in the same region and ping the EIP again. If the ping is successful, the virtual network is functional. In this case, choose **Service Tickets > Create Service Ticket** in the upper right corner of the management console to submit a service ticket for help.

5. **Check whether the DB instance is in the Storage full state.**

    Check method: View the storage space usage on the console or Cloud Eye.

    –   On the management console

        Locate the target instance and click its name to go to the **Basic Information** page. In the **Storage & Backup** area, view the storage space usage.

    –   On Cloud Eye

        Locate the target instance and click **View Metric** in the **Operation** column. On the displayed page, view the storage space usage.

        Possible cause: If the kernel system detects that the disk usage exceeds the specified threshold, the instance is set to read-only and no data can be written to the DB instance.

Solution: See **Scaling Up Storage Space**.

6. **View common connection error messages.**

When you run commands to connect to a DB instance, the possible error and solution are provided in the following:

gsql: ERROR: dn_6001_6002_6003: Invalid username/password,login denied.

If this error message is displayed when you attempt to connect to a GaussDB instance, check whether the username or password is correct.

7. If the problem persists, **create a service ticket** to contact customer service for help.

4.2 Can an External Server Access GaussDB Instances?

DB Instance Bound with an EIP

If the GaussDB instance is publicly accessible, you can access the DB instance over public networks.

DB Instance Not Bound with an EIP

- Enable a VPN in a VPC and use the VPN to connect to the GaussDB instance.
- Create a GaussDB instance and an ECS in the same VPC and access the GaussDB instance through the ECS.

For details, see **Using gsql to Connect to an Instance**.

4.3 Do Applications Need to Support Automatic Reconnections to GaussDB Databases?

It is recommended that your applications support automatic reconnections to the database. After a database reboot, your applications will automatically reconnect to the database to increase service availability and continuity.

To reduce resource consumption and improve performance, configure your applications to connect to the database using a persistent connection.

4.4 Why Can't I Ping My EIP After It Is Bound to a GaussDB Instance?

Troubleshooting Process

1. Check security group rules.
2. Check network ACL rules.
3. Ping the ECS to the instance in the same region.

Solution

1. Check security group rules.

   a. **Log in to the management console**.

   b. Click ⊙ in the upper left corner and select a region and project.
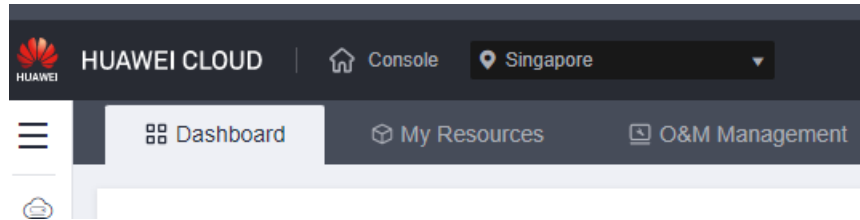
   c. Click ☰ in the upper left corner of the page and choose **Databases** > **GaussDB**.

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

   e. In the **Network Information** area, click the security group name in the **Security Group** field.

   f. Check whether the ECS NIC security group allows the inbound ICMP traffic.

   Table 4-2 Security group rules

Direction	Type	Protocol/Port Range	Source IP Address
Inbound	IPv4	Any/Any	0.0.0.0/0 **0.0.0.0/0** indicates all IP addresses.
Inbound	IPv4	ICMP/Any	0.0.0.0/0 **0.0.0.0/0** indicates all IP addresses.

2. Check network ACL rules.

   a. Check whether the network ACL is enabled or disabled.

   b. Check whether the NIC to which the EIP bound belongs to the subnet associated with the network ACL.

   c. If the network ACL is enabled, add an ICMP rule to allow traffic.

   📖 NOTE

   > The default network ACL rule denies all inbound and outbound packets. This default rule is still applied even if the network ACL is disabled.

3. Ping the ECS to the instance in the same region.

   If the affected EIP can be pinged from another ECS in the same region, the virtual network is functional. In such a case, contact technical support.

4.5 Can I Access a GaussDB Instance Over an Intranet Connection Across Regions?

By default, instances cannot be accessed over an intranet across regions. Cloud services in different regions cannot communicate with each other over an intranet. You can use EIP, Cloud Connect, or Virtual Private Network (VPN) to connect to instances across regions.

- You can access GaussDB instances across regions using EIP. For details, see **Using gsql to Connect to an Instance**.

- Cloud Connect allows you to connect two VPCs in the same account or in different accounts even if they are in different regions. For details, see **Connecting VPCs in the Same Account**.

- VPN uses an encrypted tunnel to connect VPCs in different regions and sends traffic over the Internet. It is inexpensive, easy to configure, and easy to use. However, the quality of VPN connections depends on the quality of Internet connections. For details, see **Connecting to a VPC Through a VPN**.

5 Database Storage

5.1 What Should I Do If My Data Exceeds the Available Storage of a GaussDB Instance?

Symptom

There is not enough storage available for a GaussDB instance and the instance becomes read-only, so applications cannot write any data to the instance.

Possible Causes

1. Increased workload data
2. Too much data being stored
3. Too many logs generated due to a large number of transactions and write operations
4. Too many temporary files generated due to a large number of sorting queries executed by applications

Solution

1. For insufficient storage caused by increased workload data, **scale up storage space**.

   If the original storage has reached the maximum, **upgrade the instance specifications** first.

2. If too much data is stored, delete unnecessary historical data.

   a. If the instance becomes read-only, you need to contact customer service to cancel the read-only status first.

   b. To clear up space, you can optimize tables with a high fragmentation rate during off-peak hours.

   To delete data of an entire table, use the **DROP** or **TRUNCATE** statement. To delete part of table data, use the **DELETE** statement.

3. If log files occupy too much storage, clear log files to release storage space.

4.   If temporary files generated by sorting queries occupy too much storage space, optimize your SQL statements.

5.2 How Do I View the Storage Usage of My GaussDB 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 On the **Basic Information** page, view the storage usage in the **Storage & Backup** area.

----End

6 Database Usage

6.1 How Do I Use DAS to Query SQL Statements Executed in GaussDB?

DAS is a professional database management tool with a graphical user interface. You can enable SQL Explorer to query related SQL statements.

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** > **Data Admin Service** The DAS console is displayed.

Step 4 In the navigation pane, choose **Intelligent O&M** > **Instance List** to go to the **Instance Overview** page.

Step 5 On the **Instance Overview** page, select **GaussDB Instances** in the filter area.

Figure 6-1 Instance overview
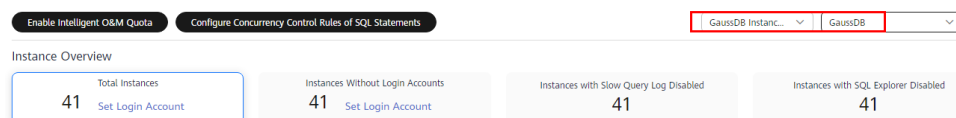
Step 6 Locate the instance you want to view and click **Details**.

Step 7 Choose **SQL** > **SQL Explorer** to view full SQL query details of the instance.

Step 8 On the **SQL Statements** tab, toggle on **Enable DAS SQL Explorer**. Query the SQL statements executed by the current instance by time range, user, keyword, operation type, database, or other filters.

Step 9 Filter operation types by referring to **Table 6-1** and click **Export** to export the corresponding SQL statements.

Table 6-1 Common SQL statement types

Type	Keyword
DDL	CREATE, DROP, ALTER
DML	INSERT, UPDATE, DELETE, SELECT
DCL	GRANT, REVOKE

☐ **NOTE**

A maximum of 10,000 SQL statements can be displayed. If you need to view more, click **Export**.

A maximum of 100,000 records can be exported.

----End

6.2 What Should I Do If Garbled Characters Are Displayed After SQL Query Results Are Exported to an Excel File for My GaussDB Instance?

The default encoding of exported data UTF-8. You need to convert the default encoding to Unicode in the exported Excel file.

6.3 What Do I Do If the root Account of My GaussDB Instance Is Locked?

Step 1 On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.

Step 2 In the **Basic Information** area, click **Reset Password** next to the **Administrator** field.

Step 3 Enter a new password and confirm the password.

----End

6.4 Why Did the New Password Not Take Effect After I Reset the Administrator Password of a GaussDB Instance?

Possible Causes

You may have restored from a backup before you reset the administrator password.

Locating Method

Check whether the DB instance is restored after the administrator password is reset.

Solution

On the GaussDB console, reset the administrator password again. For details, see **Resetting the Administrator Password**.

6.5 What Do I Do If Replay Speed of Standby DNs Cannot Catch Up with Write Speed of Primary DN?

Symptom

When workloads on a DB instance are heavy, the replay speed of standby DNs cannot catch up with the write speed of the primary DN. After the system runs for a long time, logs are accumulated on the standby DNs. If the primary DN is faulty, data restoration takes a long time and the database is unavailable, severely affecting system availability.

Solution

GaussDB provides ultimate RTO to minimize the data recovery time after a primary DN is faulty and improve availability.

To use ultimate RTO, submit an application by choosing **Service Tickets > Create Service Ticket** in the upper right corner of the console.

Precautions

- Ultimate RTO focuses only on whether the RTO of the standby DN meets the requirements. Ultimate RTO has no inherent flow control and uses the **recovery_time_target** parameter for flow control instead.
- Ultimate RTO uses multi-page redo threads to accelerate the replay progress. When the replay on the standby DN catches up with that on the primary DN and the standby DN is unloaded, the CPU usage of a single page redo thread is about 15% (the actual value depends on the hardware and parameter configuration). Total CPU usage of the replay on the standby DN = CPU usage of a single page redo thread x Number of page redo threads. Because more threads are started, the CPU and memory consumption is higher than that of parallel replay and serial replay.
- Ultimate RTO supports read on standby nodes. Because historical data pages are read, the query performance on the standby DNs is worse than that on the primary DN and worse than that of read on standby nodes during parallel redo. However, query blocking is alleviated.
- The replay speed of DDL logs is much slower than that of page modification logs. Frequent DDL operations may increase the primary/standby latency.
- When the I/O and CPU usage of a node is too high (it is recommended that the I/O and CPU usage be less than or equal to 70%), the performance of replay and read on standby nodes deteriorates significantly.

6.6 Can I Change the VPC to Which My GaussDB Instance Belongs?

No, you cannot directly change the VPC on the GaussDB console.

However, you can change a VPC by restoring the full backup of your instance to the VPC you want to use. For details, see **Data Restoration**.

7 Backup and Restoration

7.1 How Long Does GaussDB Store Backup Data?

Automated backup data is kept based on the backup retention period you specified. For details, see **Configuring an Automated Backup Policy for Instances**.

There is no limit for the manual backup retention period. You can delete manual backups as needed. For details, see **Deleting a Manual Backup**.

7.2 How Do I Clear GaussDB Backup Space?

- **Automated full and incremental backups**

  Automated backups cannot be manually deleted. You need to change the backup retention period by referring to **Configuring an Automated Backup Policy**. Backups that have expired will be automatically deleted.

- **Manual full backups**

  You can manually delete manual backups. For details, see **Deleting a Manual Backup**.

  After the manual backups are deleted, you can go to the **Instances** page and click the name of the target instance to go to the **Basic Information** page. In the **Storage & Backup** area, you can check the backup space usage.

7.3 Can My GaussDB Instance Still Be Used in the Backup Window?

A backup window is a user-specified time segment during which backup of GaussDB instances is performed. With these periodic data backups, GaussDB allows you to restore DB instances to a point in time within the backup retention period.

- During the backup window, you can still use your GaussDB instance. However, you cannot perform operations such as rebooting it on the console.

- When starting a full backup task, GaussDB first tests connectivity to your instance. If either of the following conditions is met, the test fails and a retry is performed. If the retry fails, the backup task fails.
  - DDL operations are being performed on the DB instance.
  - The backup lock fails to be obtained from the DB instance.

7.4 How Is GaussDB Backup Data Billed?

All the GaussDB backups are stored on OBS without occupying the storage of your DB instances. GaussDB provides free backup space of the same size as your purchased storage. If the backup storage usage exceeds 100% of your provisioned database storage, tiered pricing starts. For details, see **Billing**.

The lifecycle of automated backups is the same as that of the instance. If you delete an instance, its automated backups are also deleted, but manual backups are not. For details, see **Deleting a Manual Backup**.

> **NOTICE**
>
> If your storage is frozen, it is no longer billed and the free backup space is also unavailable.
>
> If your instance is frozen, no free backup space is available. As a result, the original automated backups of the instance will be billed.
>
> - If you unfreeze the instance, the free backup space will be restored.
> - If you directly delete the frozen instance, its automated backups will also be deleted and the backup space will not be billed any longer.

7.5 How Can I Back Up a GaussDB Database to an ECS?

You can back up data to an ECS in the same way you **export SQL statements**. The ECS service does not have restrictions on the types of data to be backed up as long as the data complies with local laws and regulations. You can store database backup data on an ECS. However, you are advised not to use an ECS as the database backup space.

You are advised to use the **automated backup** and **manual backup** functions of GaussDB to store backups to a professional storage object, such as the OBS, to ensure high data reliability and service assurance.

7.6 Will Backups Be Retained After My GaussDB Instance Is Deleted?

If your DB instance is deleted, its manual backups are retained by default and will be billed based on the OBS pricing details. If you do not need the backups anymore, **delete them manually**. If your instance is frozen, its backups are not billed.

Automated backups and their related files are automatically deleted.

7.7 Why Does Automated Backup for a GaussDB Instance Fail?

The possible causes and solutions for automatic backup failures are as follows:

- The network environment may be unstable due to problems such as network delay or interruptions.

  If the system detects any of these problems, it triggers another automated backup half an hour later. Alternatively, you can perform a manual backup immediately.

- If multiple tasks are being executed simultaneously, there can be problems such as excessive task wait times or interruptions.

  If the system detects any of these problems, it triggers another automated backup half an hour later. Alternatively, you can perform a manual backup immediately.

- The DB instance is abnormal probably because it is faulty or being modified.

  If the system detects any of these problems, it triggers another automated backup half an hour later. Alternatively, you can perform a manual backup immediately.

- A parameter change was incorrect.

  If your instance becomes faulty after you modify parameters of a parameter template and apply the template to the instance, check whether the modified parameters are set to correct values and whether there are any associated parameters that need to be changed, or reset the parameters to their defaults and reboot the instance.

- An error occurred during data import.

  For example, system catalog records are lost due to inappropriate data import. You can import data again by referring to the migration plan.

- If the problem persists, contact customer service.

7.8 Why Is Data Table Lost or Data Deleted from My GaussDB Instance?

GaussDB does not delete or perform any operations on any user data. If this problem occurs, check whether there have been any misoperations and restore the data from backup files, if necessary.

Restore data using backup files:

- Use the GaussDB restoration function to restore data.
- Import backup data from the ECS to GaussDB.

7.9 Does GaussDB Support Restoring Data for Individual Tables?

Yes. GaussDB supports table-level restoration.

- You can use an automated or manual backup to restore data in tables to the state at the time when the backup was created. For details, see **Restoring a GaussDB Database or Table Using a Backup File**.
- You can use an automated backup to restore data in tables to a specified point in time. For details, see **Restoring a GaussDB Database or Table to a Specific Point in Time**.

7.10 How Can I Delete the GaussDB Backup Policy?

Sorry, you cannot delete the GaussDB backup policy.

When you create a GaussDB instance, the instance-level automated backup policy is enabled by default. You can modify the backup cycle and retention period on the console. For operation details, see **Configuring an Automated Backup Policy for Instances**.

8 Database Monitoring

8.1 Which Monitoring Metrics of GaussDB Instances Do I Need to Pay Most Attention To?

For a GaussDB instance, you need to pay the most attention to the CPU, memory, and storage space (disk) usage.

You can configure the system to report alarms as needed and take measures to clear any reported alarms.

Configuration examples:

- Configure the system to report alarms to Cloud Eye if its average CPU usage reaches or exceeds a specific value (for example, 90%) multiple times (for example, 4 times) within a set period (for example, 5 minutes).

- Configure the system to report alarms to Cloud Eye if its average memory usage reaches or exceeds a specific value (for example, 90%) multiple times (for example, 3 times) within a set period (for example, 5 minutes).

- Configure the system to report alarms to Cloud Eye if its maximum storage usage reaches or exceeds a specific value (for example, 85%) multiple times (for example, 2 times) within a set period (for example, 5 minutes).

> 📖 **NOTE**
>
> For details on Cloud Eye alarm configuration, see section "Creating an Alarm Rule" in the *Cloud Eye User Guide*.

Measures:

- If a CPU or memory alarm is reported, you can scale up the CPU or memory by changing the DB instance specifications.

  For details, see **Changing CPU and Memory Specifications**.

- If a storage space usage alarm is reported, you can scale up the storage space.

  For details, see **Scaling Up Storage Space**.

8.2 How Can I Calculate the Memory Usage of a GaussDB Instance?

The formula for calculating the memory usage is as follows:

Memory usage = (Total memory – (Available memory + Buffer memory + Cache memory))/Total memory

9 Scaling and Specification Change

9.1 Are My GaussDB Instances Still Available During Storage Scale-up and Instance Specification Change?

Currently, you can scale up storage space and change the CPU and memory specifications of a GaussDB instance.

- When storage space scale-up, GaussDB instances are available and services are not interrupted. However, you cannot delete or reboot the instances that are being scaled.

- During CPU and memory specification change, the network is intermittently disconnected for one or two times in seconds. A failover may occur during this period and services may be briefly interrupted.

  For DB instances earlier than V2.0-3.100, the instances will be rebooted after their specifications are changed. During the reboot, services are unavailable. For instances of version V2.0-3.100 or later, no reboot is required, which greatly reduces the impact of intermittent disconnections. To prevent service interruption, perform the operation during off-peak hours. Rebooting a DB instance will clear the cached memory in it. You are advised to reboot it during off-peak hours.

10 Database Parameter Modification

10.1 How Can I Change the Time Zone of a GaussDB Instance?

You can set the time zone only on the GaussDB console. To change the time zone for an instance, perform the following steps:

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 Search for the **timezone** parameter in the search box.

Step 7 Select a time zone, and click **Save**.

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

For example, to change the time zone to UTC+08:00, select **Asia/Shanghai** from the drop-down list.

----End

10.2 How Can I Configure a Password Expiration Policy?

You can set the global variable **password_effect_time** to control the default validity period of a user password.

You can change the value of the **password_effect_time** parameter on the GaussDB console. For operation details, see **Modifying Parameters in a Parameter Template**.

The value of **password_effect_time** indicates how many days until a password expires. The default value is **0**, indicating that the created user password will never expire.

10.3 Can I Use SQL Commands in GaussDB to Modify Global Parameters?

Sorry, you cannot use SQL commands to modify global parameters, but you can modify specific parameters on the GaussDB console.

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 Change the values of target parameters and click **Save**.

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

----End

10.4 Can I Enable SELECT Operation Auditing in GaussDB?

Yes, you can enable SELECT operation auditing on the GaussDB console. Specifically, perform the following operations:

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 Search for the **audit_dml_state_select** parameter in the search box.

Step 7 Select a value from the drop-down list and click **Save**.

- **0**: Disable the function.

- **1**: Enable the function.

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

----End

10.5 How Do I Set the Disk Usage Threshold for Putting a Database Node into Read-Only Mode in GaussDB?

You can set the disk usage threshold for putting a database node into read-only mode on the GaussDB console. Specifically, perform the following operations:

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 Search for the **cms:datastorage_threshold_value_check** parameter in the search box.

Step 7 Enter a value ranging from 1 to 99 and click **Save**.

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

----End

10.6 How Do I Set the Date and Time Display Format in GaussDB?

You can set the date and time display format on the GaussDB console. Specifically, perform the following operations:

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 Search for the **datestyle** parameter in the search box.

Step 7 Select a format from the drop-down list and click **Save**.

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

----End

10.7 How Do I Set the Time Zone Used for Displaying and Interpreting Timestamps in GaussDB?

You can set the time zone for displaying and interpreting timestamps on the GaussDB console. Specifically, perform the following operations:

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 Search for the **timezone** parameter in the search box.

Step 7 Select a timezone from the drop-down list and click **Save**.

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

----End

10.8 How Do I Change the Maximum Number of DN Connections Allowed in GaussDB?

You can change the maximum number of concurrent connections to DNs in a centralized instance on the GaussDB console. Specifically, perform the following operations:

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 Search for the **dn:max_connections** parameter in the search box.

Step 7 Change the parameter value and click **Save**.

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

----End

10.9 How Do I Enable the Audit of INSERT, UPDATE, and DELETE Operations in GaussDB?

You can enable the audit of INSERT, UPDATE, and DELETE operations on the GaussDB console. Specifically, perform the following operations:

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 Search for the **audit_dml_state** parameter in the search box.

Step 7 Select a value from the drop-down list and click **Save**.

- **0**: Disable the function.
- **1**: Enable the function.

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

----End

10.10 How Do I Set the Maximum Number of WAL Sender Threads That Can Be Created in GaussDB?

WAL sender threads are occupied when standby DNs connect to primary DNs to obtain physical logs and when logical replication tools connect to primary DNs to obtain logical logs. On the GaussDB console, you can set the maximum number of WAL sender threads that can be created. Specifically, perform the following operations:

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 Search for the **max_wal_senders** parameter in the search box.

Step 7 Change the parameter value and click **Save**. The value ranges from 8 to 100.

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

----End

> ⚠️ **CAUTION**
>
> ● If this parameter is set to a value smaller than 20, scale-out may fail.
> ● The value of this parameter must be smaller than that of **max_connections**.

10.11 How Do I Set the Maximum Number of Bytes to Be Logged for Each SQL Statement in GaussDB?

The **track_activity_query_size** parameter specifies the maximum number of bytes that can be logged for each SQL statement. If the number of bytes of a SQL statement exceeds the specified parameter value, the SQL statement will be truncated. Setting this parameter to a large value can lead to excessive memory usage, potentially causing the system to run out of memory. You are advised to set this parameter to a value no more than 4096.

You can set the maximum number of bytes to be logged for each SQL statement in a centralized instance on the GaussDB console. Specifically, perform the following operations:

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 Search for the **track_activity_query_size** parameter in the search box.

Step 7 Change the parameter value and click **Save**. The value ranges from 100 to 102,400.

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

----End

10.12 How Do I Set the Connection Timeout Interval for GaussDB?

The connection timeout interval indicates the maximum time (in seconds) allowed during which no operation is performed after the connection to the server is established. You can set the connection timeout interval on the GaussDB console. For example, if you want to change the timeout interval to 5 minutes, change the

value of the related parameter to **300**. The value **0** indicates that the timeout setting is disabled. Specifically, perform the following operations:

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 Search for the **session_timeout** parameter in the search box.

Step 7 Change the parameter value and click **Save**. The value ranges from 0 to 86,400.

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

----End

10.13 How Do I Set the Automatic Unlocking Time After an Account Is Locked in GaussDB?

The **password_lock_time** parameter specifies the duration before a locked account is automatically unlocked, in days. The value is a floating-point number ranging from 0 to 365. The value **0** indicates that the account is unlocked immediately. The integral part of the parameter value indicates the number of days and its decimal part can be converted into hours, minutes, and seconds. For example, if **password_lock_time** is set to **1.5**, the duration is one day and 12 hours.

You can set the automatic account unlocking time on the GaussDB console. Specifically, perform the following operations:

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 Search for the **password_lock_time** parameter in the search box.

Step 7 Change the parameter value and click **Save**. The value ranges from 0 to 365.

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

----End

> ⚠️ **CAUTION**
>
> The locking and unlocking functions take effect only when the values of both **password_lock_time** and **failed_login_attempts** are positive numbers. **failed_login_attempts** specifies the maximum number of incorrect password attempts before an account is locked. For details, see **How Do I Set the Maximum Number of Failed Login Attempts for Accounts in GaussDB?**

10.14 How Do I Set the Maximum Number of Failed Login Attempts for Accounts in GaussDB?

The **failed_login_attempts** parameter specifies the maximum number of incorrect password attempts before an account is locked. The account will be automatically unlocked after the time (in seconds) specified in **password_lock_time** elapses. You can set the maximum number of failed login attempts on the GaussDB console. Specifically, perform the following operations:

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 Search for the **failed_login_attempts** parameter in the search box.

Step 7 Change the parameter value and click **Save**. The value ranges from 0 to 1000. The value **0** indicates that the number of incorrect password attempts is not limited.

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

----End

> ⚠️ **CAUTION**
>
> The locking and unlocking functions take effect only when the values of both **password_lock_time** and **failed_login_attempts** are positive numbers. For details about how to configure **password_lock_time**, see **How Do I Set the Automatic Unlocking Time After an Account Is Locked in GaussDB?**

11 Log Management

11.1 How Do I View All SQL Logs Executed by GaussDB?

You can use the visualized database management service Data Admin Service (DAS) to quickly search for target SQL execution records.
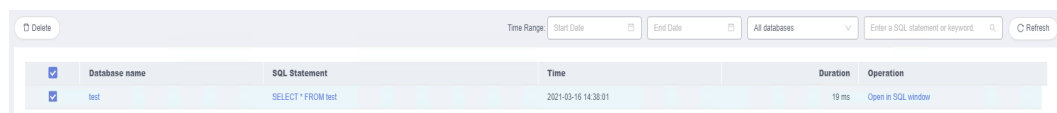
Procedure

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 target DB instance and click **Log In** in the **Operation** column.

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

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

Step 7 On the displayed page, search for execution information about the target SQL statement by time range, database name, or keyword.

Figure 11-1 SQL History

- To access the **Database Management** page, click a database name.
- To copy and use a SQL statement, click it in the **SQL Statement** column.

- To directly execute a SQL statement, click **Open in SQL window** in the **Operation** column.

  ----End

11.2 How Do I View Deadlock Logs of GaussDB?

Database deadlock logs are recorded in error logs. If **Error Log Collection** is enabled, you can download error logs and quickly locate deadlock-related records by searching for keywords **Lock wait timeout** in the logs.

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 **Log Analysis**.

Step 6 On the **Error Logs** tab, select a specified time period and click **Download** in the **Operation** column to download the error log file.

Step 7 Open the error log file, and search for **Lock wait timeout** to quickly locate deadlock errors.

----End

12 Network Security

12.1 How Can I Prevent Untrusted Source IP Addresses from Accessing GaussDB Instances?

- If you enable public accessibility, your EIP DNS and database port may be vulnerable to hacking. To protect information such as your EIP, DNS, database port, database account, and password, you are advised to set the range of source IP addresses in the GaussDB **security group** to ensure that only trusted source IP addresses can access your DB instances.

- To prevent your database password from being cracked, set a strong password and periodically change it.

12.2 How Can I Import the Root Certificate to a Windows or Linux Server?

Importing the Root Certificate to a Windows Server

1. Press **Win+R** to open the **Run** dialog box, enter **MMC**, and press **Enter**.
2. On the displayed console, choose **File** > **Add/Remove Snap-in**.
3. In the **Available snap-ins** pane to the left of the displayed **Add or Remove Snap-ins** dialog box, select **Certificates** and click **Add**.
4. In the displayed **Certificates snap-in** dialog box, select **Computer account** and click **Next**.
5. In the displayed **Select Computer** dialog box, click **Finish**.
6. In the **Add or Remove Snap-ins** dialog box, click **OK**.
7. On the console, double-click **Certificates**.
8. Right-click **Trusted Root Certification Authorities** and choose **All Tasks** > **Import**.
9. Click **Next**.
10. Click **Browse** to change the file type to **All Files (*.*)**.

11. Locate the downloaded root certificate (a ca.pem file) and click **Open**. Then, click **Next**.

NOTICE

You must change the file type to **All Files (*.*)** because **.pem** is not a standard certificate extension name.

---

1. Click **Next**.
2. Click **Finish**.
3. Click **OK** to complete the import of the root certificate.

Importing the Root Certificate to a Linux Server

You can use a connection tool (such as WinSCP or PuTTY) to upload the certificate to any directory on a Linux server.