

Data Warehouse Service

User Guide

Issue 01
Date 2024-07-31



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

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

Trademarks and Permissions



HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

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

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

Contents

1 Process for Using GaussDB(DWS)	1
2 Preparations	5
3 Creating or Deleting a Cluster	6
3.1 Accessing the GaussDB(DWS) Management Console.....	6
3.2 Creating a Dedicated Cluster.....	6
3.2.1 Creating a DWS 2.0 Cluster.....	6
3.2.2 Creating a DWS 3.0 Cluster.....	21
3.3 Purchasing a Discount Package.....	32
3.4 Deleting a Cluster.....	33
4 Cluster Connection	35
4.1 Methods of Connecting to a Cluster.....	35
4.2 Obtaining the Cluster Connection Address.....	36
4.3 Using the Data Studio GUI Client to Connect to a Cluster.....	38
4.4 Using the gsql CLI Client to Connect to a Cluster.....	43
4.4.1 Downloading the Client.....	43
4.4.2 Using the Linux gsql Client to Connect to a Cluster.....	46
4.4.3 Using the Windows gsql Client to Connect to a Cluster.....	50
4.4.4 Establishing Secure TCP/IP Connections in SSL Mode.....	53
4.5 Using the JDBC and ODBC Drivers to Connect to a Cluster.....	60
4.5.1 Development Specifications.....	60
4.5.2 JDBC Version Description.....	61
4.5.3 Downloading the JDBC or ODBC Driver.....	63
4.5.4 Using JDBC to Connect to a Cluster.....	66
4.5.5 Configuring JDBC to Connect to a Cluster (Load Balancing Mode).....	78
4.5.6 Configuring JDBC to Connect to a Cluster (IAM Authentication Mode).....	80
4.5.7 Using ODBC to Connect to a Cluster.....	84
4.6 Using the Third-Party Function Library psycopg2 of Python to Connect to a Cluster.....	90
4.7 Using the Python Library PyGreSQL to Connect to a Cluster.....	100
4.8 Managing Database Connections.....	116
5 Monitoring and Alarms	120
5.1 Dashboard.....	120
5.2 Databases Monitoring (DMS).....	123

5.2.1 Database Monitoring Overview.....	123
5.2.2 Monitoring Metrics.....	124
5.2.3 Cluster Overview.....	161
5.2.4 Monitoring.....	164
5.2.4.1 Node Monitoring.....	164
5.2.4.2 Performance Monitoring.....	167
5.2.4.3 Database Monitoring.....	168
5.2.4.4 Real-Time Queries.....	169
5.2.4.5 Historical Queries.....	173
5.2.4.6 Instance Monitoring.....	174
5.2.4.7 Resource Pool Monitoring.....	175
5.2.5 Utilities.....	177
5.2.5.1 SQL Diagnosis.....	177
5.2.5.2 SQL Probes.....	180
5.2.5.3 Table Diagnosis.....	182
5.2.5.4 Intelligent Distribution Column Recommendation.....	186
5.2.6 Workload Analysis.....	188
5.2.6.1 Workload Analysis Overview.....	188
5.2.6.2 Workload Snapshots.....	188
5.2.6.3 Workload Reports.....	191
5.2.7 Settings.....	194
5.2.8 Checking Task Details.....	195
5.2.9 Typical Scenarios.....	196
5.2.9.1 SQL Diagnosis.....	196
5.3 Monitoring Clusters Using Cloud Eye.....	196
5.4 Alarms.....	205
5.4.1 Alarm Management.....	205
5.4.2 Alarm Rules.....	209
5.4.3 Alarm Subscriptions.....	211
5.4.4 Alarm Handling.....	213
5.4.4.1 DWS_200000001 Node CPU Usage Exceeds the Threshold.....	213
5.4.4.2 DWS_200000006 Node Data Disk Usage Exceeds the Threshold.....	216
5.4.4.3 DWS_200000009 Node Data Disk I/O Usage Exceeds the Threshold.....	219
5.4.4.4 DWS_200000012 Node Data Disk Latency Exceeds the Threshold.....	221
5.4.4.5 DWS_200000016 Data Spilled to Disks of the Query Statement Exceeds the Threshold.....	223
5.4.4.6 DWS_200000017 Number of Queuing Query Statements Exceeds the Threshold.....	225
5.4.4.7 DWS_200000018 Queue Congestion in the Default Cluster Resource Pool.....	227
5.4.4.8 DWS_200000020 Long SQL Probe Execution Duration in a Cluster.....	230
5.4.4.9 DWS_200000023 A Vacuum Full Operation That Holds a Lock for A Long Time Exists in the Cluster.....	231
5.5 Event Notifications.....	233
5.5.1 Event Notifications Overview.....	233
5.5.2 Subscribing to Event Notifications.....	237

5.5.3 Viewing Events.....	238
6 Specifications Change and Scaling.....	239
6.1 Viewing Inspection Results.....	239
6.2 Managing Nodes.....	247
6.3 Scaling Nodes.....	249
6.3.1 Scaling Out a Cluster.....	249
6.3.2 Cluster Redistribution.....	254
6.3.2.1 Redistributing Data.....	254
6.3.2.2 Viewing Redistribution Details.....	256
6.3.3 Scaling In a Cluster.....	258
6.4 Changing Specifications.....	261
6.4.1 Using the Elastic Specification Change.....	261
6.4.2 Changing All Specifications.....	267
6.4.3 Disk Capacity Expansion of an EVS Cluster.....	270
7 Backup and Disaster Recovery.....	272
7.1 Snapshots.....	272
7.1.1 Overview.....	272
7.1.2 Manual Snapshots.....	273
7.1.2.1 Creating a Manual Snapshot.....	273
7.1.2.2 Deleting a Manual Snapshot.....	274
7.1.3 Automated Snapshots.....	275
7.1.3.1 Automatic Snapshot Overview.....	275
7.1.3.2 Configuring an Automated Snapshot Policy.....	276
7.1.3.3 Copying Automated Snapshots.....	280
7.1.3.4 Deleting an Automated Snapshot.....	281
7.1.4 Viewing Snapshot Information.....	281
7.1.5 Restoration Using a Snapshot.....	283
7.1.5.1 Constraints on Restoring a Snapshot.....	283
7.1.5.2 Restoring a Snapshot to a New Cluster.....	283
7.1.5.3 Restoring a Snapshot to the Original Cluster.....	286
7.1.6 Configuring a Snapshot.....	286
7.1.7 Stopping Snapshot Creation.....	289
7.2 Cluster DR.....	290
7.2.1 DR Overview.....	290
7.2.2 Creating a DR Task.....	292
7.2.3 Viewing DR Information.....	293
7.2.4 DR Management.....	293
7.2.5 Mutually Exclusive DR Cases.....	297
8 Intelligent O&M.....	298
8.1 Overview.....	298
8.2 Managing O&M Plans.....	299

8.3 Viewing O&M Tasks.....	303
9 Cluster Management.....	305
9.1 Modifying Database Parameters.....	305
9.2 Checking the Cluster Status.....	307
9.3 Viewing Cluster Details.....	313
9.4 Changing a Cluster Name.....	318
9.5 O&M Account.....	319
9.6 Managing Access Domain Names.....	320
9.7 Cluster Topology.....	323
9.8 Managing Tags.....	328
9.8.1 Overview.....	328
9.8.2 Managing Tags.....	329
9.9 Managing Enterprise Projects.....	331
9.10 Managing Clusters That Fail to Be Created.....	333
9.11 Removing the Read-only Status.....	335
9.12 Performing a Primary/Standby Switchback.....	335
9.13 Starting and Stopping a Cluster.....	336
9.14 Resetting a Password.....	338
9.15 Cluster Upgrade.....	338
9.16 Associating and Disassociating ELB.....	341
9.17 Managing CNs.....	346
10 Data Integration.....	349
10.1 Overview.....	349
10.2 Managing Instances.....	349
10.3 Managing Connections.....	363
10.4 Managing Table Mappings.....	368
10.5 Managing Jobs.....	370
11 Cluster Log Management.....	373
12 Database User Management.....	375
12.1 Managing Users.....	375
12.2 Managing Roles.....	378
13 Audit Logs.....	382
13.1 Audit Log Overview.....	382
13.2 Management Console Audit Logs.....	382
13.3 Database Audit Logs.....	385
13.3.1 Configuring the Database Audit Logs.....	385
13.3.2 Dumping the Database Audit Logs.....	387
13.3.3 Viewing Database Audit Logs.....	394
14 Cluster Security Management.....	397
14.1 Configuring Separation of Duties for the GaussDB(DWS) Cluster.....	397

14.2 Encrypting Databases.....	399
14.2.1 Overview.....	400
14.2.2 Rotating Encryption Keys.....	402
14.2.3 Converting an Ordinary Cluster to an Encrypted Cluster.....	402
14.3 Permissions.....	403
14.3.1 Creating a User and Granting GaussDB(DWS) Permissions.....	404
14.3.2 Creating a GaussDB(DWS) Custom Policy.....	405
14.3.3 Syntax of Fine-Grained Permissions Policies.....	407
14.3.4 RBAC Syntax of RBAC Policies.....	437
14.4 Protection for Mission-Critical Operations.....	439
15 Resource Management.....	442
15.1 Overview.....	442
15.2 Resource Pool.....	444
15.2.1 Feature Description.....	444
15.2.2 Creating a Resource Pool.....	448
15.2.3 Modifying a Resource Pool.....	449
15.2.4 Deleting a Resource Pool.....	454
15.3 Resource Management Plan.....	455
15.3.1 Managing Resource Management Plans.....	455
15.3.2 Managing Resource Management Plan Stages.....	456
15.4 Workspace Management.....	460
15.5 Exception Rules.....	461
16 Data Source Management.....	466
16.1 MRS Data Sources.....	466
16.1.1 MRS Data Source Usage Overview.....	466
16.1.2 Creating an MRS Data Source Connection.....	467
16.1.3 Updating the MRS Data Source Configuration.....	470
16.2 Managing OBS Data Sources.....	472
16.3 Managing LakeFormation Data Sources.....	476
17 Managing Logical Clusters.....	479
17.1 Logical Cluster Overview.....	479
17.2 Adding/Deleting a Logical Cluster.....	486
17.3 Managing Logical Clusters.....	487
17.4 Scheduling GaussDB(DWS) 3.0 Logical Cluster Creation and Deletion.....	490
17.5 Tutorial: Converting a Physical Cluster That Contains Data into a Logical Cluster.....	491
17.6 Tutorial: Dividing a New Physical Cluster into Logical Clusters.....	495
17.7 Tutorial: Setting a Read-Only Logical Cluster and Binding It to a User.....	497

1 Process for Using GaussDB(DWS)

GaussDB(DWS) is an online data processing database that uses the Huawei Cloud infrastructure to provide scalable, fully-managed, and out-of-the-box analytic database service, freeing you from complex database management and monitoring. It is a native cloud service based on the converged data warehouse GaussDB, and is fully compatible with the standard ANSI SQL 99 and SQL 2003, as well as the PostgreSQL and Oracle ecosystems. GaussDB(DWS) provides competitive solutions for PB-level big data analysis in various industries.

GaussDB(DWS) provides an easy-to-use management console, allowing you to quickly create clusters and easily manage data warehouses.

Process Description

Figure 1-1 Process for using GaussDB(DWS)

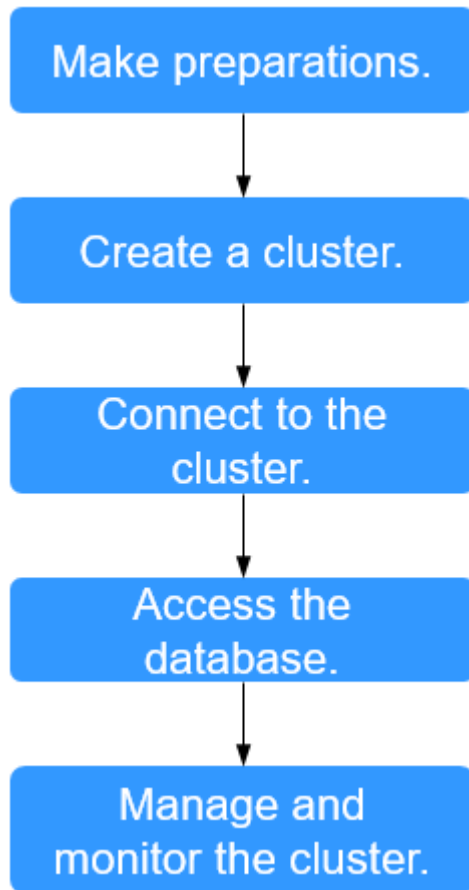


Table 1-1 Process description

Process	Task	Description	Operation Instruction
Make preparations.	-	Before using GaussDB(DWS), select an open port on your firewall as the database port of your data warehouse cluster.	Preparations
Create a cluster.	-	Create a cluster before using GaussDB(DWS) to execute data analysis tasks. A GaussDB(DWS) cluster contains nodes in the same subnet. These nodes jointly provide services. During cluster creation, the system creates a default database.	<ul style="list-style-type: none"> Creating a DWS 2.0 Cluster

Process	Task	Description	Operation Instruction
Connect to the cluster.	-	After the data warehouse cluster is successfully created, use the SQL client tool or a third-party driver such as JDBC or ODBC to connect to the database in the cluster. You can download the SQL client tool and JDBC/ODBC driver on the Client Connections page of the GaussDB(DWS) management console.	Methods of Connecting to a Cluster
Access the database.	-	After connecting to the cluster, you can create and manage databases, manage users and permissions, import and export data, and query and analyze data.	Data Warehouse Service (DWS) Developer Guide
Manage and monitor the cluster.	Cluster management	View the cluster status, modify cluster configurations, add cluster tags, and scale out, restart, and delete the cluster.	Managing clusters
	Snapshot management	Create snapshots to back up and restore the cluster.	Snapshots
	O&M and monitoring	View the running status and performance of the cluster through monitoring, log auditing, event notification, and resource load management.	<ul style="list-style-type: none"> • Monitoring Clusters Using Cloud Eye • Event Notifications Overview • Audit Logs • Resource Management
	Scaling and specification change	<ul style="list-style-type: none"> • Expand the capacity of an existing cluster on the console if your service requires additional compute or storage resources. • Change the specifications of created clusters on the console. 	<ul style="list-style-type: none"> • Scaling Out a Cluster • Changing Specifications

Process	Task	Description	Operation Instruction
	Cluster upgrade	Cluster 8.1.1 and later versions allow users to deliver cluster upgrade operations on the console.	Cluster Upgrade
	Resource load management	GaussDB(DWS) provides the resource management function. You can put resources (CPU, memory, I/O, and storage space) into different resource pools, which are isolated from each other.	Resource Management

2 Preparations

Before using GaussDB(DWS), make the following preparations:

- [Registering a Public Cloud Account](#)
- [Determining the Cluster Ports](#)

Registering a Public Cloud Account

If you do not have a Huawei Cloud account, register an account.

1. Open the official public cloud website (<https://www.huaweicloud.com/eu/>) and click **Register** in the upper right corner. The registration page is displayed.
2. Enter registration information as prompted..
3. After the registration is successful, you can be automatically logged in to Huawei Cloud.

Determining the Cluster Ports

- When creating a GaussDB(DWS) cluster, you need to specify a port for SQL clients or applications to access the cluster.
- If your client is behind a firewall, you need an available port so that you can connect to the cluster and perform query and analysis from the SQL client tool.
- If you do not know an available port, contact the network administrator to specify an open port on your firewall. GaussDB(DWS) supports port numbers between **8000** and **30000**. By default, a data warehouse cluster is created with port number **8000**.
- After a cluster is created, its port number cannot be changed. Ensure that the port specified is available.

3 Creating or Deleting a Cluster


3.1 Accessing the GaussDB(DWS) Management Console

Scenario

This section describes how to log in to the GaussDB(DWS) management console and use GaussDB(DWS).

Procedure

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner of the console homepage to expand the **Service List**, and choose **Analytics > GaussDB(DWS)**.

----End

3.2 Creating a Dedicated Cluster

3.2.1 Creating a DWS 2.0 Cluster

To use Huawei Cloud GaussDB(DWS), create a data warehouse cluster first.

This section describes how to create a data warehouse cluster on the GaussDB(DWS) console.

 NOTE

- To balance loads, achieve high availability, and avoid single-node faults, if no ELB is bound during cluster creation, you can bind an ELB on the cluster details page after the cluster is created. For details, see [Associating and Disassociating ELB](#).
- The GaussDB(DWS) clusters under the same account are physically isolated and cannot share data. You can import data from a remote GaussDB(DWS) cluster to a local one by using a foreign table. For details, see [Tutorial: Importing Remote GaussDB\(DWS\) Data Sources](#).

Preparations Before Creating a Cluster

- You have evaluated the flavor of cluster nodes.
You can select the number of nodes by data volume, service load, and performance. More nodes bring you stronger storage and compute capabilities.
When first using GaussDB(DWS), you can create a cluster with a smaller flavor. Then, you can adjust the cluster scale and node flavor based on the data volume and service load changes without interrupting services. For details, see [Scaling Out a Cluster](#).
- Determine the number of nodes that can be used by users.
The number of nodes that can be used by users must meet the following requirements. Otherwise, the system displays a message indicating that the cluster cannot be created.
The number of nodes that can be used by a user depends on the product type you select. A hybrid data warehouse cluster (standalone mode) has only one node. For other types of clusters, the number of nodes can be greater than or equal to 3. You can view the number of available nodes on the **Clusters > Dedicated Clusters** page.

Creating a Cluster

Step 1 Go to the [page for creating a GaussDB\(DWS\) cluster](#).

Step 2 Choose **Region** and select the actual working region of the cluster node.

Step 3 Select an AZ. You can select **Single AZ** or **Multi-AZ** as required.

For more information, see [Regions and AZs](#).

 NOTE

- Multi-AZ clusters are supported only by clusters of version 8.2.0.100 or later.
- The **Multi-AZ** option is displayed only if the number of AZs in the selected region is greater than or equal to 3. If this condition is not met, only a single-AZ cluster can be created.
- For a multi-AZ cluster, only three AZs can be selected at a time so far. Server nodes are evenly distributed among the three AZs.
- The multi-AZ cluster supports only DWS 2.0 standard data warehouses.
- The numbers of nodes in a multi-AZ cluster must be a multiple of 3.
- In a multi-AZ cluster, the number of DNs must be less than or equal to 2.

Step 4 Configure **Resource**, **CPU Architecture**, and **Node Flavor**.

 **NOTE**

- The number of nodes in a new cluster cannot exceed the quota that can be used by a user or 256. If the node quota is insufficient, click **Increase quota** to submit a service ticket and apply for higher node quota.
- After a cluster is created, its type cannot be changed. For details about the differences between product types, see [Data Warehouse Types](#).

Table 3-1 Node configuration parameters

Parameter	Description	Example Value
Resource	<p>Product type. It can be:</p> <ul style="list-style-type: none">• Standard data warehouse: It can analyze hot and cold data and is highly cost-effective. Its storage and computing resources are not limited, and can be elastically scaled and billed per use. It is suitable for the converged analysis that requires integrated databases, warehouses, marts, and lakes. It is most suitable for OLAP workloads.• Stream data warehouse: It provides efficient time series computing and IoT analysis capabilities based on the standard data warehouse and supports correlation between real-time and historical data. The compression ratio can reach 40:1. It can be used for IoT real-time analysis.• Hybrid data warehouse: It provides high-concurrency, high-performance, and low-latency transaction processing capabilities at low costs based on large-scale data query and analysis capabilities. The data warehouse can be used to process HTAP hybrid loads, and can be deployed in standalone or cluster mode. <p>NOTE</p> <p>A hybrid data warehouse can be deployed in cluster or standalone mode.</p> <ul style="list-style-type: none">• Cluster deployment: If the name of the selected node flavor contains h (for example, dwsx2.h.4xlarge.4.c6), the hybrid data warehouse can be deployed in cluster mode. You can deploy multiple nodes, scale nodes, and manage resource pools.• Standalone deployment: If the name of the selected node flavor contains h1 (for example, dwsx2.h1.xlarge.2.c6), the hybrid data warehouse only supports standalone deployment, which does not provide HA capabilities. The storage cost can be reduced by half. A standalone data warehouse can be restored by the automatic reconstruction of ECS, and its data reliability is ensured by the EVS multi-copy mechanism. It is less expensive than other specifications. It is a good choice for lightweight services.	Standard


Parameter	Description	Example Value
Compute Resource	<p>It can be:</p> <ul style="list-style-type: none">• ECS: ECS is a scalable virtual block storage service built on a distributed architecture. It boasts high data reliability and robust I/O throughput, ensuring that data can be swiftly migrated and restored in the event of any backup failure. This helps prevent data loss resulting from hardware failures. Backup and restoration can be performed on ECSs and EVS disks. You can configure automatic backup policies for them.	-
Storage Type	<p>It can be:</p> <ul style="list-style-type: none">• Cloud SSD• Extreme SSD: suitable for workloads that demand super-high bandwidth and super-low latency.• Extreme SSD V2: ultra-high-performance SSD EVS disks dedicated for latency-sensitive mission-critical applications.• Local SSD <p>NOTE Local SSD disks do not support disk scale-out. For more information, see Disk Types and Performance.</p>	-
CPU Architecture	<p>The CPU architecture includes:</p> <ul style="list-style-type: none">• x86• Kunpeng <p>NOTE The only difference between the x86 and Kunpeng architectures lies in the underlying architecture, of which the application layer is unaware. The same SQL syntax is used. If x86 servers are sold out when you create a cluster, select the Kunpeng architecture.</p>	-


Parameter	Description	Example Value
Node Flavor	<p>Select the desired node flavor based on service requirements. Each node flavor displays the vCPU, memory, and recommended application scenario.</p> <p>For more information about the node flavors supported by GaussDB(DWS) and their prices, see the GaussDB(DWS) pricing details.</p> <p>For details about the node flavors supported by GaussDB(DWS), see Data Warehouse Specifications.</p>	dws2.m6.4xlarge. 8
Hot storage	<p>Available storage capacity of each node.</p> <p>NOTE</p> <ul style="list-style-type: none">• The storage capacity you apply for has the necessary file system overhead, which includes index nodes and the space required for database running. The storage space must be an integer multiple of 100.• 200 GB per node is the actual storage capacity for service data. For example, if the number of nodes is set to 3, the total resource capacity is 600 GB.• By default, tablespaces are automatically created when you configure cold and hot data storage. You do not need to manually create tablespaces. This feature is supported only in clusters of 8.1.3 and later versions.	-
Nodes	<p>Specify the number of nodes in the cluster.</p> <p>The number of nodes ranges from 3 to 256.</p>	3
Total	<p>Displays the total capacity of a cluster.</p> <p>The storage capacity of each flavor is the actual database space used for storing data. The displayed storage capacity has deducted the disk space consumed by backups and RAIDs.</p>	-

Step 5 Click **Next: Configure Network**.

Step 6 Configure the network.

Table 3-2 Network parameters

Parameter	Description	Example Value
VPC	<p>Specify a virtual private network for nodes in a cluster to isolate networks of different services.</p> <p>If you create a data warehouse cluster for the first time and have not configured the VPC, click View VPC. On the VPC management console that is displayed, create a VPC that satisfies your needs.</p> <p>For details about how to create a VPC, see Creating a VPC in the <i>Virtual Private Cloud User Guide</i>.</p> <p>After selecting a VPC from the drop-down list, click View VPC to enter the VPC management console and view the detailed information about the VPC.</p> <p>You can click  to refresh the options in the VPC drop-down list.</p>	vpc-dws
Subnet	<p>Specify a VPC subnet.</p> <p>A subnet provides dedicated network resources that are isolated from other networks, improving network security.</p> <p>NOTE</p> <p>After a cluster is created, the subnet cannot be modified. If you need to modify the subnet, you can restore the snapshot of the cluster to a new cluster. The data of the new cluster is the same as that of the old cluster, and the subnet can be modified when the new cluster is created.</p>	subnet-dws

Parameter	Description	Example Value
Security Group	<p>Specify a VPC security group.</p> <p>A security group restricts access rules to enhance security when GaussDB(DWS) and other services access each other.</p> <ul style="list-style-type: none">• Automatic creation If Automatic creation is selected, the system automatically creates a default security group. This option is selected by default.<p>The rule of the default security group is as follows: The outbound allows all access requests, while the inbound is open only to the database port that you set to connect to the GaussDB(DWS) cluster.</p><p>The format of the default security group name is <code>dws-<Cluster_name>-<Cluster_database_port></code>, for example, dws-dws-demo-8000.</p><p>NOTE</p><p>If the quotas of the security group and the security group rule are insufficient, an error message will be displayed after you submit the cluster creation application. Select an existing group and retry.</p>• Manual creation You can also log in to the VPC management console to manually create a security group. Then, go back to the page for creating data warehouse clusters, click  next to the Security Group drop-down list to refresh the page, and select the new security group.<p>To enable the GaussDB(DWS) client to connect to the cluster, you need to add an inbound rule to the new security group to grant the access permission to the database port of the GaussDB(DWS) cluster. The following is an example of an inbound rule..</p><ul style="list-style-type: none">- Protocol: TCP- Port: 8000. Use the database port set when creating the GaussDB(DWS) cluster. This port is used for receiving client connections to GaussDB(DWS).- Source: Select IP address and use the host IP address of the client host, for example, 192.168.0.10/32.	Automatic creation

Parameter	Description	Example Value
	<p>After a GaussDB(DWS) cluster is created, you can change the security group. You can also add, delete, or modify security group rules in the current security group. For details, see Modifying a Security Group. Changing the security group of a cluster may cause brief service disruption. Exercise caution when performing this operation. For better network performance, do not select more than five security groups.</p>	

Parameter	Description	Example Value
Public Network Access	<p>Specify whether users can use a client to connect to a cluster's database over the Internet. The following methods are supported:</p> <ul style="list-style-type: none">• Do not use: The EIP is not required. If DWS is used in the production environment, first bind it to ELB, and then bind it to an EIP on the ELB page.• Buy Now: Users specify the bandwidth of the EIP and the system automatically assigns an EIP that exclusively uses bandwidth to each cluster so that users can use the EIP to access the cluster over the Internet. The bandwidth name of an automatically assigned EIP starts with the cluster name.• Specify: A specified EIP is bound to the cluster. If no available EIPs are displayed in the drop-down list, click Create EIP to go to the Elastic IP page and create an EIP that satisfies your needs. You can set the bandwidth as needed. <p>NOTE</p> <ul style="list-style-type: none">• If you use the EIP binding function for the first time in each project of each region, the system prompts you to create the DWSAccessVPC agency to authorize GaussDB(DWS) to access VPC. After the authorization is successful, GaussDB(DWS) can switch to a healthy VM when the VM bound with the EIP becomes faulty.• By default, only Huawei Cloud accounts or users with Security Administrator permissions can query and create agencies. By default, the IAM users in those accounts cannot query or create agencies. When the users use the EIP, the system makes the binding function unavailable. Contact a user with the DWS Administrator permissions to authorize the agency on the current page.• Do not use indicates disabling access to the cluster over the public network. After a cluster is created, if you want to access it over the public network, bind an EIP to the cluster and create a public network domain name. For details, see Creating a Public Network Domain Name.• If GaussDB(DWS) is used for the production environment, the new GaussDB(DWS) cluster needs to be bound to ELB and then to EIP. Select Do not use here.	Buy now

Parameter	Description	Example Value
ELB	<p>Specifies whether ELB is bound. With ELB health checks, CN requests of a cluster can be quickly forwarded to normal CNs. If a CN is faulty, the workload can be immediately shifted to a healthy node, minimizing cluster access faults. Currently, ELBs can be bound in the same VPC or across VPCs.</p> <ul style="list-style-type: none">• Do not use: The load balancer is not used. If GaussDB(DWS) is used in the production environment, first bind it to ELB, and then bind it to an EIP on the ELB page.• Specify: Specify an ELB to be bound to the cluster. If no available load balancers are displayed in the drop-down list, click Create ELB to go to the Elastic Load Balance page and create a load balancer as needed.	Specify
Bandwidth	When EIP is set to Buy now , you need to specify the bandwidth of the EIP, which ranges from 1 Mbit/s to 100 Mbit/s.	50Mbit/s

Step 7 Click **Next: Configure Advanced Settings**.

Step 8 Configure cluster parameters.

Table 3-3 Cluster parameters

Parameter	Description	Example Value
Cluster Name	<p>Set the name of the data warehouse cluster.</p> <p>The cluster name contains 4 to 64 case-insensitive characters and must start with a letter. Only letters, digits, hyphens (-), and underscores (_) are allowed.</p> <p>NOTE If the cluster name cannot be changed on the console, contact technical support.</p>	dws-demo
Cluster Version	Displays the version of the database instance installed in the cluster. The example version number is for reference only.	-
Default Database	<p>The default database name of the cluster is gaussdb.</p> <p>NOTE This name cannot be changed.</p>	gaussdb

Parameter	Description	Example Value
Administrator Account	<p>Set the database administrator name.</p> <p>The administrator username must:</p> <ul style="list-style-type: none"> • Consist of lowercase letters, digits, or underscores. • Start with a lowercase letter or an underscore. • Contain 6 to 64 characters. • Cannot be a keyword of the GaussDB(DWS) database. For details about the keywords of the GaussDB(DWS) database, see Keyword in <i>Data Warehouse Service (DWS) Developer Guide</i>. 	dbadmin
Administrator Password	<p>Set the password of the database administrator account.</p> <p>The password complexity requirements are as follows:</p> <ul style="list-style-type: none"> • Consists of 12 to 32 characters. • Cannot be the username or the username spelled backwards. • Must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters (~!?,,;_){}[]/<>@#%^&*+ \=-) • Passes the weak password check. <p>NOTE Change the password regularly and keep it secure.</p>	-
Confirm Password	Enter the database administrator password again.	-
Database Port	<p>Specify the port used when the client or application connects to the database in the cluster.</p> <p>The port number ranges from 8000 to 30000.</p> <p>NOTE The database port of a created cluster cannot be changed. You can specify the database port only when creating a cluster.</p>	8000
Time Zone	You can set the time zone for the tenant cluster, including the system OS time zone and cluster data warehouse time zone.	-

- Step 9** Configure the enterprise project to which the cluster belongs. You can configure this parameter only when the Enterprise Project Management service is enabled. The default value is **default**.

An enterprise project facilitates project-level management and grouping of cloud resources and users.

You can select the default enterprise project **default** or other existing enterprise projects. To create an enterprise project, log in to the Enterprise Management console. For details, see the *Enterprise Management User Guide*.

- Step 10** Configure advanced settings. Select **Default** to keep the default values of the advanced parameters. You can also select **Custom** to modify the values.
- After this function is enabled, you can set the retention mode and the backup device used by the current cluster for automated snapshots. For more information, see [Table 3-4](#).

Table 3-4 Automated snapshot parameters

Parameter	Description
Backup Device	Select OBS or NFS from the drop-down list.
NFS backup file system address (NFS parameter)	NFS shared IP address. Enter the IP address of the SFS shared path. After the mounting is successful, a mount directory is created in the /var/chroot/nfsbackup directory of the cluster instance by default.

- CNs**
CNs, or Coordinators, receive access requests from the clients and return the execution results. They also split and distribute tasks to the Datanodes (DNs) for parallel execution.
The value ranges from 3 to the number of cluster nodes. The maximum value is **20** and the default value is **3**. In a large-scale cluster, you are advised to deploy multiple CNs.
- Tag**
A tag is a key-value pair used to identify a cluster. For details about the keys and values, see [Table 3-5](#). By default, no tag is added to the cluster.
If your organization has configured GaussDB(DWS) tag policies, you need to add tags to clusters based on the tag policies. If a tag does not comply with the tag policies, cluster creation may fail. Contact your organization administrator to learn more about tag policies.
For more information about tags, see [Overview](#).

Table 3-5 Tag parameters

Parameter	Description	Example Value
Key	<p>The options are as follows:</p> <ul style="list-style-type: none">Select a predefined tag key or an existing resource tag key from the drop-down list of the text box. <p>NOTE To add a predefined tag, you need to create one on TMS and select it from the drop-down list of Tag key. You can click View predefined tags to enter the Predefined Tags page of TMS. Then, click Create Tag to create a predefined tag. For more information, see section Creating Predefined Tags in the <i>Tag Management Service User Guide</i>.</p> <ul style="list-style-type: none">Enter a tag key in the text box. A tag key can contain a maximum of 36 characters. It cannot be an empty string or start or end with a space. The value cannot contain the following characters: =*<>\\, / <p>NOTE A key must be unique in a given cluster.</p>	key01
Value	<p>You can select:</p> <ul style="list-style-type: none">Select a predefined tag value or resource tag value from the drop-down list of the text box.Enter a tag value in the text box. A tag value can contain a maximum of 43 characters, which can be an empty string. It cannot start or end with a space. The value cannot contain the following characters: =*<>\\, /	value01

- **Encrypt DataStore**

If this function is enabled, Key Management Service (KMS) encrypts the cluster and the cluster's snapshot data.

When you enable database encryption for each project in each region for the first time, the system displays a **Create Agency** dialog box. Click **Yes** to create **DWSAccessKMS** to authorize GaussDB(DWS) to access KMS. If you click **No**, the encryption function is not enabled. Select the created KMS key from the **KMS Key Name** drop-down list.

NOTICE

- Only users with the Tenant Admin permission can view and toggle the **Encrypt DataStore** switch.
- By default, only Huawei Cloud accounts or users with Security Administrator permissions can query and create agencies. IAM users under an account do not have the permission to query or create agencies by default. Contact a user with that permission and complete the authorization on the current page.
- The database encryption function cannot be disabled once it is enabled.
- After **Encrypt DataStore** is enabled, the key cannot be disabled, deleted, or frozen when being used. Otherwise, the cluster becomes abnormal and the database becomes unavailable.
- After database encryption is enabled, you cannot use open APIs to restore created snapshots.

- Method 1: Select a key name.
- Method 2: Enter the key ID. Enter the key ID used for authorizing the current tenant..

When you grant permissions on the [Creating a Grant](#) page, the authorized object must be an account instead of a user. The authorized operations must at least contain **Querying key details**, **Encrypting data**, and **Decrypting data**.

Step 11 Enable IPv6 networking. If this function is enabled, a client or application can connect to the database using an IPv6 address.

 **NOTE**

To enable IPv6 networking, the following conditions must be met:

- The subnet configured in [Step 6](#) is an IPv6 dual-stack subnet.
- The cluster supports IPv6 addresses and a maximum of three NICs.
- The cluster version must be 8.2.1.210 or later.

Step 12 Click **Next: Confirm**.

 **NOTE**

If the number of requested nodes, vCPU (cores), or memory (GB) exceed the user's remaining quota, a warning dialog box is displayed, indicating that the quota is insufficient and displaying the detailed remaining quota and the current quota application. You can click **Increase quota** in the warning dialog box to submit a service ticket and apply for higher node quota. Once approved, we will update your resource quota accordingly and send you a notification. For details about quota operations, see [Quotas](#).

Step 13 Click **Buy Now**. If the billing mode is yearly/monthly billing, click **Buy Now**. The payment page is displayed.

After the submission is successful, the creation starts. Click **Back to Cluster List**. The cluster management page is displayed. The initial status of the cluster is **Creating**. Cluster creation takes some time. Wait for a while. Clusters in the **Available** state are ready for use.

 NOTE

- For load balancing and high availability purposes, and to prevent single CN failures, a cluster must be bound to ELB. For details, see [Associating and Disassociating ELB](#).

----End

3.2.2 Creating a DWS 3.0 Cluster

DWS 3.0 uses the cloud-native, cost-effective architecture with decoupled storage and computing. It supports hot and cold data analysis, elastic scaling of storage and computing, unlimited computing power and capacity, and pay-per-use pricing. It is applicable to OLAP analysis scenarios.

This section describes how to create a DWS 3.0 cluster on the GaussDB(DWS) console.

 NOTE

To balance loads, achieve high availability, and avoid single-node faults, if no ELB is bound during cluster creation, you can bind an ELB on the cluster details page after the cluster is created. For details, see [Associating and Disassociating ELB](#).

Preparations Before Creating a Cluster

- You have evaluated the flavor of cluster nodes.
You can select the number of nodes by data volume, service load, and performance. More nodes bring you stronger storage and compute capabilities.
When first using GaussDB(DWS), you can create a cluster with a smaller flavor. Then, you can adjust the cluster scale and node flavor based on the data volume and service load changes without interrupting services. For details, see [Scaling Out a Cluster](#).
- Determine the number of nodes that can be used by users.
Ensure that the number of available nodes is greater than or equal to 3. Otherwise, the system displays a message indicating that the cluster cannot be created. You can choose **Clusters > Dedicated Clusters** to view the number of available nodes.

Creating a Cluster

Step 1 Go to the [page for creating a GaussDB\(DWS\) cluster](#).

Step 2 Select a region and an AZ.

Step 3 Configure **Resource**, **CPU Architecture**, and **Node Flavor**.

 NOTE

- The number of nodes in a new cluster cannot exceed the quota that can be used by a user or 256. If the node quota is insufficient, click **Increase quota** to submit a service ticket and apply for higher node quota.
- If you have yearly/monthly nodes that meet service requirements, you are advised to use these nodes first to save costs. You can select **Yearly/Monthly** for **Billing Mode**.

Table 3-6 Node configuration parameters


Parameter	Description	Example Value
Resource	The options are as follows: Standard data warehouse: It can analyze hot and cold data and is highly cost-effective. Its storage and computing resources are not limited, and can be elastically scaled and billed per use. It is suitable for the converged analysis that requires integrated databases, warehouses, marts, and lakes. It is most suitable for OLAP workloads.	Standard
Version	DWS 3.0	DWS 3.0
Compute Resource	It can be: ESC: Scalable, reliable, and high-throughput virtual block storage is provided in a distributed architecture. This ensures that data can be quickly migrated and restored if any data replica is unavailable, preventing data from being lost because of a single hardware fault. Backup and restoration can be performed on ECSs and EVS disks. You can configure automatic backup policies for them.	-
Storage Type	It can be: <ul style="list-style-type: none">• Cloud SSD• Extreme SSD: suitable for workloads that demand super-high bandwidth and super-low latency.• Extreme SSD V2: ultra-high-performance SSD EVS disks dedicated for latency-sensitive mission-critical applications.	-
CPU Architecture	The following CPU architectures can be selected: <ul style="list-style-type: none">• x86• Kunpeng NOTE The only difference between the x86 and Kunpeng architectures lies in the underlying architecture, of which the application layer is unaware. The same SQL syntax is used. If x86 servers are sold out when you create a cluster, select the Kunpeng architecture.	-


Parameter	Description	Example Value
Node Flavor	Select a node flavor. Each node flavor shows the vCPU, memory, and recommended application scenario. For more information about the node flavors supported by GaussDB(DWS) and their prices, see the GaussDB(DWS) pricing details . For details about the node flavors supported by GaussDB(DWS), see Data Warehouse Specifications .	-
Hot Storage (with Cache)	Available storage capacity of each node. NOTE <ul style="list-style-type: none">The storage capacity you apply for has the necessary file system overhead, which includes index nodes and the space required for database running. If the idle space is less than 100 GB, scale-out may fail due to insufficient running space.The displayed 200GB/node includes the storage for cache. For example, if you create 3 nodes, each having 200 GB capacity, the total resource capacity is 600 GB, and the actual storage space available to you is 300 GB.	-
Cold Data	Store data in separate OBS buckets, which are billed on a pay-per-use basis.	-
Nodes	Specify the number of nodes in the cluster. The number of nodes ranges from 3 to 256.	3
Total	Display the cluster's total capacity. The storage capacity of each flavor includes the storage for cache. The displayed storage capacity includes the disk space consumed by backups and RAIDs.	-

Step 4 Click **Next: Configure Network**.

Step 5 Configure the network.

Table 3-7 Network parameters

Parameter	Description	Example Value
VPC	<p>Specify a VPC to isolate the cluster's network.</p> <p>If you create a data warehouse cluster for the first time and have not configured the VPC, click View VPC. On the VPC management console that is displayed, create a VPC as needed.</p> <p>For details about how to create a VPC, see Creating a VPC in the <i>Virtual Private Cloud User Guide</i>.</p> <p>After selecting a VPC from the drop-down list, click View VPC to enter the VPC management console and view the detailed information about the VPC.</p> <p>You can click  to refresh the options in the VPC drop-down list.</p>	vpc-dws
Subnet	<p>Specify a VPC subnet.</p> <p>A subnet provides dedicated network resources that are isolated from other networks, improving network security.</p>	subnet-dws

Parameter	Description	Example Value
Security Group	<p>Specify a VPC security group.</p> <p>A security group restricts access rules to enhance security when GaussDB(DWS) and other services access each other.</p> <ul style="list-style-type: none">• Automatic creation If Automatic creation is selected, the system automatically creates a default security group. This option is selected by default.<p>The rule of the default security group is as follows: The outbound allows all access requests, while the inbound is open only to the database port that you set to connect to the GaussDB(DWS) cluster.</p><p>The format of the default security group's name is <i>dws-<cluster name>-<database port of the GaussDB(DWS) cluster></i>, for example, dws-dws-demo-8000.</p><p>NOTE</p><p>If the quotas of the security group and the security group rule are insufficient, an error message will be displayed after you submit the cluster creation application. You can select an existing group and retry.</p>• Manual creation You can also log in to the VPC management console to manually create a security group. Then, go back to the page for creating data warehouse clusters, click  next to the Security Group drop-down list to refresh the page, and select the new security group.<p>To enable the GaussDB(DWS) client to connect to the cluster, you need to add an inbound rule to the new security group to grant the access permission to the database port of the GaussDB(DWS) cluster. The following is an example of an inbound rule..</p><ul style="list-style-type: none">- Protocol: TCP.- Port: 8000. Use the database port set when creating the GaussDB(DWS) cluster. This port is used for receiving client connections to GaussDB(DWS).- Source: Select IP address and use the host IP address of the client host, for example, 192.168.0.10/32.	Automatic creation

Parameter	Description	Example Value
	After a GaussDB(DWS) cluster is created, you can change the security group. You can also add, delete, or modify security group rules in the current security group. For details, see Modifying a Security Group .	

Parameter	Description	Example Value
EIP	<p>Specify whether users can use a client to connect to a cluster's database over the Internet. The following methods are supported:</p> <ul style="list-style-type: none">• Do not use: Do not specify any EIPs here. If GaussDB(DWS) is used in the production environment, first bind it to ELB, and then bind it to an EIP on the ELB page.• Buy Now: Users specify the bandwidth of the EIP and the system automatically assigns an EIP that exclusively uses bandwidth to each cluster so that users can use the EIP to access the cluster over the Internet. The bandwidth name of an automatically assigned EIP starts with the cluster name.• Specify: Specify an EIP to be bound to the cluster. If no available EIPs are displayed in the drop-down list, click Create EIP to go to the Elastic IP page and create an EIP as needed. The bandwidth can be customized. <p>NOTE</p> <ul style="list-style-type: none">• If you use the EIP binding function for the first time in each project of each region, the system prompts you to create the DWSAccessVPC agency to authorize GaussDB(DWS) to access VPC. After the authorization is successful, GaussDB(DWS) can switch to a healthy VM when the VM bound with the EIP becomes faulty.• By default, only cloud accounts or users with Security Administrator permissions can query and create agencies. By default, the IAM users in those accounts cannot query or create agencies. When the users use the EIP, the system makes the binding function unavailable. Contact a user with the DWS Administrator permissions to authorize the agency on the current page.• Do not use indicates disabling access to the cluster over the public network. After a cluster is created, if you want to access it over the public network, bind an EIP to the cluster and create a public network domain name. For details, see Creating a Public Network Domain Name.• If GaussDB(DWS) is used for the production environment, the new GaussDB(DWS) cluster needs to be bound to ELB and then to EIP. Select Do not use here.	Buy now

Parameter	Description	Example Value
ELB	<p>Specifies whether ELB is bound. With ELB health checks, CN requests of a cluster can be quickly forwarded to normal CNs. If a CN is faulty, the workload can be immediately shifted to a healthy node, minimizing cluster access faults. Currently, ELBs can be bound in the same VPC or across VPCs.</p> <ul style="list-style-type: none"> • Do not use: The load balancer is not used. If GaussDB(DWS) is used in the production environment, first bind it to ELB, and then bind it to an EIP on the ELB page. • Specify: Specify an ELB to be bound to the cluster. If no available load balancers are displayed in the drop-down list, click Create ELB to go to the Elastic Load Balance page and create a load balancer as needed. 	Specify
Bandwidth	Specifies the EIP bandwidth. The value ranges from 1 Mbit/s to 100 Mbit/s. This parameter is mandatory if EIP is set to Buy now .	50 Mbit/s

Step 6 Click **Next: Configure Advanced Settings**.

Step 7 Configure cluster parameters.

Table 3-8 Cluster parameters

Parameter	Description	Example Value
Cluster Name	<p>Set the name of the data warehouse cluster. The cluster name contains 4 to 64 case-insensitive characters and must start with a letter. Only letters, digits, hyphens (-), and underscores (_) are allowed.</p> <p>NOTE If the cluster name cannot be changed on the console, contact technical support.</p>	dws-demo
Cluster Version	Version of the database instance installed in the cluster. The example version number is for reference only.	9.0.0
Default Database	<p>The default database name of the cluster is gaussdb.</p> <p>NOTE This name cannot be changed.</p>	gaussdb

Parameter	Description	Example Value
Administrator or Account	<p>Set the database administrator name.</p> <p>The username must meet the following requirements:</p> <ul style="list-style-type: none">• Consists of lowercase letters, digits, or underscores.• Starts with a lowercase letter or an underscore.• Contains 6 to 64 characters.• Cannot be a keyword of the GaussDB(DWS) database. For details about the keywords of the GaussDB(DWS) database, see Keyword in the <i>Data Warehouse Service (DWS) Developer Guide</i>.	dbadmin
Administrator or Password	<p>Set the password of the database administrator account.</p> <p>The password complexity requirements are as follows:</p> <ul style="list-style-type: none">• Contains 12 to 32 characters.• Cannot be the username or the username spelled backwards.• Must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters (~!?,.,;_(){}[]/<>@#%&*&+ \=-)• Passes the weak password check. <p>NOTE Change the password regularly and keep it secure.</p>	-
Confirm Password	<p>Enter the database administrator password again.</p>	-
Database Port	<p>Set the port used when the client or application connects to the database in the cluster.</p> <p>The port number ranges from 8000 to 30000.</p>	8000
IPv6	<p>Specify whether to enable the IPv6 dual stack for the cluster. If this function is enabled, a client or application can connect to the database using an IPv6 address.</p> <p>NOTE To enable IPv6, the following conditions must be met:</p> <ul style="list-style-type: none">• The subnet configured in Step 6 is an IPv6 dual-stack subnet.• The cluster supports IPv6 addresses and a maximum of three NICs.• The cluster version must be 8.2.1.210 or later.	-

Parameter	Description	Example Value
Time Zone	You can set the time zone for the tenant cluster, including the system OS time zone and cluster data warehouse time zone.	-

Step 8 Select the enterprise project of the cluster. You can configure this parameter only when the Enterprise Project Management service is enabled. The default value is **default**.

An enterprise project facilitates project-level management and grouping of cloud resources and users.

You can select the default enterprise project **default** or other existing enterprise projects. To create an enterprise project, log in to the Enterprise Management console. For details, see [Enterprise Management User Guide](#).

Step 9 Configure advanced parameters. Select **Default** to keep the default values of the advanced parameters. You can also select **Custom** to modify the values.

- After this function is enabled, you can set the retention mode and the backup device used by the current cluster for automated snapshots. For more information, see [Table 3-9](#).

Table 3-9 Automated snapshot parameters

Parameter	Description
Backup Device	Select OBS or NFS from the drop-down list.
NFS Backup File System Address (NFS)	NFS shared IP address. To mount the SFS shared path, enter its IP address. If successful, a mount directory will be created in the /var/chroot/nfsbackup directory of the cluster instance.

- **CNs**

CNs receive access requests from the clients and return the execution results. In addition, a CN splits and distributes tasks to the DNs for parallel execution.

The value ranges from 3 to the number of cluster nodes. The maximum value is **20** and the default value is **3**. In a large-scale cluster, you are advised to deploy multiple CNs.

- **Tags**

A tag is a key-value pair used to identify a cluster. For details about the keys and values, see [Table 3-10](#). By default, no tag is added to the cluster.

If your organization has configured GaussDB(DWS) tag policies, you need to add tags to clusters based on the tag policies. If a tag does not comply with the tag policies, cluster creation may fail. Contact your organization administrator to learn more about tag policies.

For details about tags, see [Overview](#).

Table 3-10 Tag parameters

Parameter	Description	Example Value
Tag key	<p>You can:</p> <ul style="list-style-type: none">Select a predefined tag key or an existing resource tag key from the drop-down list of the text box. <p>NOTE To add a predefined tag, you need to create one on TMS and select it from the drop-down list of Tag key. You can click View predefined tags to enter the Predefined Tags page of TMS. Then, click Create Tag to create a predefined tag. For more information, see section Creating Predefined Tags in <i>Tag Management Service User Guide</i>.</p> <ul style="list-style-type: none">Enter a tag key in the text box. A tag key can contain a maximum of 36 characters. It cannot be an empty string or start or end with a space. The value cannot contain the following characters: =*<>\\, / <p>NOTE A key must be unique in a given cluster.</p>	key01
Value	<p>You can:</p> <ul style="list-style-type: none">Select a predefined tag value or resource tag value from the drop-down list of the text box.Enter a tag value in the text box. A tag value can contain a maximum of 43 characters, which can be an empty string. It cannot start or end with a space. The value cannot contain the following characters: =*<>\\, /	value01

Step 10 Click **Next: Confirm**. **NOTE**

If the number of requested nodes, vCPU (cores), or memory (GB) exceed the user's remaining quota, a warning dialog box is displayed, indicating that the quota is insufficient and displaying the detailed remaining quota and the current quota application. You can click **Increase quota** in the warning dialog box to submit a service ticket and apply for higher node quota. Once approved, we will update your resource quota accordingly and send you a notification. For details about quota operations, see [Quotas](#).

Step 11 Click **Buy Now**. If the billing mode is yearly/monthly billing, click **Buy Now**. The payment page is displayed.

After the submission is successful, the creation starts. Click **Back to Cluster List**. The cluster management page is displayed. The initial status of the cluster is **Creating**. Cluster creation takes some time. Wait for a while. Clusters in the **Available** state are ready for use.

----End

3.3 Purchasing a Discount Package

GaussDB(DWS) also supports discount packages. You can make a one-off payment according to the purchased service duration. The service duration ranges from one month to three years. It is economical and recommended for long-term users.

This section describes how to purchase a GaussDB(DWS) discount package.

NOTE

- A discount package is paid in advance by month or year. If you plan to use GaussDB(DWS) for a long term, nodes in a discount package are more cost-effective than nodes billed on a pay-per-use basis.
- A discount package is a billing concept. After you purchase a discount package, no cluster will be automatically created. Go to the GaussDB(DWS) console to create a cluster if no cluster exists.
- A discount package takes effect from the date of purchase and automatically expires when the subscription period ends. This period is fixed even if you do not run a GaussDB(DWS) cluster. To reduce costs, create a cluster immediately after purchasing a discount package, or purchase the package after creating a cluster.
- When a discount package expires or you unsubscribe from the package, the system will automatically charge you on a pay-per-use basis (by hour). The service will not be interrupted as long as your account balance is sufficient.
- Storage unit conversion: 1 PB = 1,024 TB; 1 TB = 1,024 GB; 1 GB = 1,024 MB; 1 MB = 1,024 KB

Purchasing a Discount Package

Step 1 Go to the [discount package purchase](#) page.

Step 2 On the displayed page, select a region.

Table 3-11 Region parameters

Parameter	Description	Example Value
Region	Select the AZ for the cluster nodes to run. Discount packages in different regions are isolated.	EU-Dublin

Step 3 (Optional) If you select **Cold data storage**, configure the specifications, quantity, and required duration. The estimated price will be displayed. Choose this if the data volume or OBS storage used in your cluster exceeds the package you purchased.

Table 3-12 Parameter description

Parameter	Description	Example Value
Storage Size	Package size. NOTE Multiple packages can be used together. If used resources exceed the package size, the excess part will be charged on a pay-per-use basis.	-
Purchase Quantity	Number of packages.	3
Validity Period	Validity period of a package. Before the package expires, the storage can be deducted from the package first. Excess usage will be charged on a pay-per-use basis.	-

Step 4 Click **Next** to switch to the **Confirm** page.

Step 5 Confirm the order information and click **Pay Now**.

Step 6 Select a payment method and complete the payment as prompted.

After the payment is successful, the order takes effect about 5 minutes later. You can use the purchased package only after the order takes effect. You can choose **Billing Center > My Orders** in the upper right corner of the console to go to the **My Orders** page and view the order status.

No cluster will be automatically created after you purchase a discount package. If you have not created a cluster, go to the GaussDB(DWS) console to create one. For details, see [Creating a Dedicated Cluster](#).

After the order takes effect, if you create a data warehouse cluster with the same flavor as the package cluster flavor you purchased, the cluster is automatically associated with the discount package. The nodes within the discount package range will not be charged during the validity period and the extra nodes will be charged on the pay-per-use basis.

----End

3.4 Deleting a Cluster

If you do not need to use a cluster, perform the operations in this section to delete it.

NOTE


- If your cluster is in arrears, this function may be unavailable. Please top up your account in time.
- A cluster that is read-only or being scaled cannot be deleted. You can delete it only after the scaling is complete or the read-only state is canceled.
- If a cluster has DR tasks, the cluster cannot be deleted. You need to delete the DR tasks and then delete the cluster.

Impact on the System

Deleted clusters cannot be recovered. Additionally, you cannot access user data and automated snapshots in a deleted cluster because the data and snapshots are automatically deleted. If you delete a cluster, its manual snapshots will not be deleted.

Deleting a Cluster

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Click  in the upper left corner of the management console to select a region.

Step 3 On the **Clusters > Dedicated Clusters** page, locate the cluster to be deleted.

Step 4 In the row of a cluster, choose **More > Delete**.

Step 5 In the displayed dialog box, confirm the deletion. You can determine whether to perform the following operations:

- Create a snapshot for the cluster.

If the cluster status is normal, you can click **Create Snapshot**. In the dialog box that is displayed, enter the snapshot name and click **OK** to create a snapshot for the cluster to be deleted. In the row of a cluster, choose **More > Delete**.

- Resource

– Release the EIP bound to the cluster.

If an EIP is bound to the cluster, you are advised to select **EIP** to release the EIP. If you do not release the EIP, you can bind it to another cluster or cloud resource and it will be billed based on the EIP pricing rule of VPC.

– Automated Snapshots

– Manual Snapshots

If you have created a manual snapshot, you can select **Manual Snapshot** to delete it.

Step 6 After confirming that the information is correct, enter **DELETE** or click **Auto Enter** and click **OK** to delete the cluster. The cluster status in the cluster list will change to **Deleting**, and the cluster deletion progress will be displayed.

If the cluster to be deleted uses an automatically created security group that is not used by other clusters, the security group is automatically deleted when the cluster is deleted.

----End

4 Cluster Connection

4.1 Methods of Connecting to a Cluster

If you have created a GaussDB(DWS) cluster, you can use the SQL client tool or a third-party driver such as JDBC or ODBC to connect to the cluster and access the database in the cluster.

The procedure for connecting to a cluster is as follows:

1. [Obtaining the Cluster Connection Address](#)
2. If SSL encryption is used, perform the operations in [Establishing Secure TCP/IP Connections in SSL Mode](#).
3. Connect to the cluster and access the database in the cluster. You can choose any of the following methods to connect to a cluster:

NOTICE

- You are advised to use the officially recommended method for connecting to the database.
- Compatibility with other clients cannot be guaranteed, so it may be necessary to verify it.
- If an error occurs due to incompatibility with another client and the client cannot be replaced, try replacing the libpq driver on the client. To replace the **libpg.so** file on the client, download and extract the gsql client package, locate the **gsql** directory, and obtain the file. Then, replace the existing **libpg.so** file in the designated directory on the client.

-
- Use the SQL client tool to connect to the cluster.
 - [Using the Linux gsql Client to Connect to a Cluster](#)
 - [Using the Windows gsql Client to Connect to a Cluster](#)
 - [Using the Data Studio GUI Client to Connect to a Cluster](#)

- Use a JDBC, psycopg2, or PyGreSQL driver to connect to the cluster.
 - [Using JDBC to Connect to a Cluster](#)
 - [Using ODBC to Connect to a Cluster](#)
 - [Using the Third-Party Function Library psycopg2 of Python to Connect to a Cluster](#)
 - [Using the Python Library PyGreSQL to Connect to a Cluster](#)
 - [Configuring JDBC to Connect to a Cluster \(IAM Authentication Mode\)](#)

4.2 Obtaining the Cluster Connection Address

Scenario

You can access GaussDB(DWS) clusters by different methods and the connection address of each connection method varies. This section describes how to view and obtain the private network address on the Huawei Cloud platform, public network address on the Internet, and JDBC connection strings.

To obtain the cluster connection address, use either of the following methods:

- [Obtaining the cluster connection address on the Client Connections Page](#)
- [Obtaining the Cluster Access Addresses on the Cluster Information Page](#)

Obtaining the cluster connection address on the Client Connections Page

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation tree on the left, choose **Management > Client Connections**.

Step 3 In the **Data Warehouse Connection Information** area, select an available cluster.

You can only select clusters in the **Available** state.

Step 4 View and obtain the cluster connection information.

- **Private Network IP Address**
- **Public Network IP Address**
- **ELB Address**
- **JDBC URL (Private Network)**
- **JDBC URL (Public Network)**
- **ODBC URL**

 NOTE

- If no EIP is automatically assigned during cluster creation, **Public Network Address** is empty. If you want to use a public network address (consisting of an EIP and the database port) to access the cluster from the Internet, click **Bind EIP** to bind one.
- If an EIP is bound during cluster creation but you do not want to use the public network address to access the cluster, click **Unbind EIP** to unbind the EIP. After the EIP is unbound, **Public Network Address** is empty.
- If a cluster was not bound to ELB when it was created, the **ELB Address** parameter will be left blank. You can bind the cluster to ELB to avoid single CN failures.
- If a cluster has been bound to ELB, use the ELB address to connect to the cluster for high availability purposes.
- If the IPv6 dual stack is enabled for a GaussDB(DWS) cluster, private IPv4 and IPv6 addresses can be used. You can connect to the cluster via IPv4 or IPv6.

----End

Obtaining the Cluster Access Addresses on the Cluster Information Page

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 3** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 4** In the **Connection** area, view and obtain the cluster's access address information, including the private network address and public network address.

Table 4-1 Connection

Parameter	Description
Private Network Domain Name	<p>Domain name for accessing the cluster database through the internal network. The domain name corresponds to all CN IP addresses. The private network domain address is automatically generated when a cluster is created.</p> <p>NOTE</p> <ul style="list-style-type: none">• If the cluster name does not comply with the domain name standards, the prefix of the default access domain name will be adjusted accordingly.• Load balancing is not supported. <p>You can click Modify to change the private network domain name. The access domain name contains 4 to 63 characters, which consists of letters, digits, and hyphens (-), and must start with a letter.</p> <p>For details, see Managing Access Domain Names.</p>

Parameter	Description
Private Network IP Address	<p>IP address for accessing the database in the cluster over the private network.</p> <p>NOTE</p> <ul style="list-style-type: none">• A private IP address is automatically generated when you create a cluster. The IP address is fixed.• The number of private IP addresses equals the number of CNs. You can log in to any node to connect to the cluster.• If you access a fixed IP address over the internal network, all the resource pools will run on a single CN.• If IPv6 is enabled for a cluster, both IPv4 and IPv6 private addresses will be displayed. You can use either of them as needed.
Public Network Domain Name	<p>Name of the domain for accessing the database in the cluster over the public network. For details, see Managing Access Domain Names.</p> <p>NOTE</p> <p>Load balancing is not supported.</p>
Public Network IP Address	<p>IP address for accessing the database in the cluster over the public network.</p> <p>NOTE</p> <ul style="list-style-type: none">• If no EIP is assigned during cluster creation and Public Network IP Address is empty, click Edit to bind an EIP to the cluster.• If an EIP is bound during cluster creation, click Edit to unbind the EIP.
Initial Administrator	<p>Database administrator specified during cluster creation. When you connect to the cluster for the first time, you need to use the initial database administrator and password to connect to the default database.</p>
Port	<p>Port number for accessing the cluster database through the public network or private network. The port number is specified when the cluster is created.</p>
Default Database	<p>Database name specified when the cluster is created. When you connect to the cluster for the first time, connect to the default database.</p>
ELB Address	<p>To achieve high availability and avoid single-CN failures, a new cluster needs to be bound to ELB. You are advised to use the ELB address to connect to the cluster.</p>

----End

4.3 Using the Data Studio GUI Client to Connect to a Cluster

Data Studio is a SQL client tool running on the Windows operating system. It provides various GUIs for you to manage databases and database objects, as well

as edit, run, and debug SQL scripts, and view execution plans. Download the Data Studio software package from the GaussDB(DWS) console. The package can be used without installation after being decompressed.

Data Studio versions include **Windows x86** (32-bit Windows system) and **Windows x64** (64-bit Windows system).

Preparations Before Connecting to a Cluster

- You have obtained the administrator username and password for logging in to the database in the data warehouse cluster.
- You have obtained the public network address, including the IP address and port number in the data warehouse cluster. For details, see [Obtaining the Cluster Connection Address](#).
- You have configured the security group of the GaussDB(DWS) cluster and added an inbound rule that allows users' IP addresses to access ports using the TCP.

For details, see [Adding a Security Group Rule](#) in the *Virtual Private Cloud User Guide*.

Connecting to the Cluster Database Using Data Studio

- Step 1** GaussDB(DWS) provides a Windows-based Data Studio client and the tool depends on the JDK. You need to install the JDK on the client host first.

NOTICE

Only JDK 1.8 is supported.

In the Windows operating system, you can download the required JDK version from the [official website of SDK](#), and install it by following the installation guidance.

- Step 2** Log in to the GaussDB(DWS) console.

- Step 3** Choose **Management > Client Connections**.

- Step 4** On the **Download Client and Driver** page, download **Data Studio GUI Client**.

- Select **Windows x86** or **Windows x64** based on the OS type and click **Download** to download a Data Studio version that matches the current cluster.

If clusters of different versions are available, you will download the Data Studio matching the earliest cluster version after clicking **Download**. If there is no cluster, you will download the Data Studio tool of the earliest version after clicking **Download**. GaussDB(DWS) clusters are compatible with earlier versions of Data Studio.

- Click **Historical Version** to download the corresponding Data Studio version. You are advised to download Data Studio based on the cluster version.

If you have clusters of different versions, the system displays a dialog box, prompting you to select the cluster version and download the corresponding client. In the cluster list on the **Clusters > Dedicated Clusters** page, click the

name of the specified cluster to go to the **Cluster Information** page and view the cluster version.


Table 4-2 Data Studio download links

Applicable OS	Download Link	Verification File
Windows x64	Data_Studio_8.2.x_64.zip	Data_Studio_8.2.x_64.zip.sha256
	Data_Studio_8.1.x_64.zip	Data_Studio_8.1.x_64.zip.sha256
	Data_Studio_8.0.x_64.zip	Data_Studio_8.0.x_64.zip.sha256
Windows x86	Data_Studio_8.2.x_32.zip	Data_Studio_8.2.x_32.zip.sha256
	Data_Studio_8.1.x_32.zip	Data_Studio_8.1.x_32.zip.sha256
	Data_Studio_8.0.x_32.zip	Data_Studio_8.0.x_32.zip.sha256

Step 5 Decompress the downloaded client software package (32-bit or 64-bit) to the installation directory.

Step 6 Open the installation directory and double-click **Data Studio.exe** to start the Data Studio client. See [Figure 4-1](#).

Figure 4-1 Starting the client




NOTE

If your computer blocks the running of the application, you can unlock the **Data Studio.exe** file to start the application.

Step 7 Choose **File > New Connection** from the main menu. See [Figure 4-2](#).

Figure 4-2 New connection



Step 8 In the displayed **New Database Connection** window, enter the connection parameters.

Table 4-3 Connection parameters

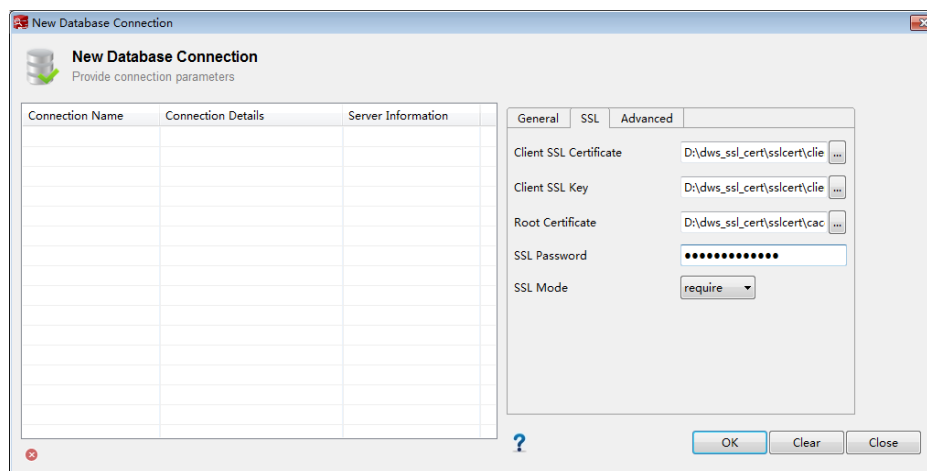
Field	Description	Example Value
Database Type	Select HUAWEI CLOUD DWS .	HUAWEI CLOUD DWS
Connection Name	Name of the connection	dws-demo
Host	IP address (IPv4) or domain name of the cluster to be connected	-
Port Number	Database port	8000
Database Name	Database name	gaussdb
Username	Username for connecting to the database	-
Password	Password for logging in to the database to be connected	-
Save Password	Select an option from the drop-down list: <ul style="list-style-type: none">● Current Session Only: The password is saved only in the current session.● Do Not Save: The password is not saved.	-
Enable SSL	If Enable SSL is selected, the client can use SSL to encrypt connections. The SSL connection mode is more secure than common modes, so you are advised to enable SSL connection.	-

When **Enable SSL** is selected, download the SSL certificate and decompress it by referring to [Downloading SSL Certificate](#). Click the **SSL** tab and configure the following parameters:

Table 4-4 Configuring SSL parameters

Field	Description
Client SSL Certificate	Select the sslcert\client.crt file in the decompressed SSL certificate directory.
Client SSL Key	Only the PK8 format is supported. Select the sslcert\client.key.pk8 file in the directory where the SSL certificate is decompressed.
Root Certificate	When SSL Mode is set to verify-ca , the root certificate must be configured. Select the sslcert\cacert.pem file in the decompressed SSL certificate directory.
SSL Cipher	Set the password for the client SSL key in PK8 format.
SSL Mode	GaussDB(DWS) supports the following SSL modes: <ul style="list-style-type: none"> • require • verify-ca GaussDB(DWS) does not support the verify-full mode.

Figure 4-3 Configuring SSL parameters



Step 9 Click **OK** to establish the database connection.

If SSL is enabled, click **Continue** in the displayed **Connection Security Alert** dialog box.

After the login is successful, the **RECENT LOGIN ACTIVITY** dialog box is displayed, indicating that Data Studio is connected to the database. You can run the SQL statement in the **SQL Terminal** window on the Data Studio page.

For details about how to use other functions of Data Studio, press **F1** to view the Data Studio user manual.

 NOTE

- Data cannot be rolled back after being added, deleted, modified, or queried in Data Studio.
- Data Studio can save connection information, excluding passwords.
- DDL/DDI and data cannot be exported in batches for the following objects:
 - **Export DDL:**
Connection, database, foreign table, sequence, column, index, constraint, partition, function/procedure group, regular tables group, views group, schemas group, and system catalog group.
 - **Export DDL and Data:**
Connection, database, namespace, foreign table, sequence, column, index, constraint, partition, function/procedure, view, regular tables group, schemas group, and system catalog group.

----End

4.4 Using the gsql CLI Client to Connect to a Cluster

4.4.1 Downloading the Client

GaussDB(DWS) provides client tool packages that match the cluster versions. You can download the desired client tool package on the GaussDB(DWS) management console.

The client tool package contains the following:

- **Linux database connection tool gsql and the script for testing sample data**

Linux gsql is a Linux command line client running in Linux. It is used to connect to the database in a data warehouse cluster.

The script for testing sample data is used to execute the introductory example.

- **Windows gsql**

Windows gsql is a command line client running on the Windows OS. It is used to connect to the database in a data warehouse cluster.

 NOTE

Only 8.1.3.101 and later cluster versions can be downloaded from the console.

- **GDS tool package**

Gauss Data Service (GDS) is a data service tool. You can use the GDS tool to import a data file in a common file system to the GaussDB(DWS) database. The GDS tool package must be installed on the server where the data source file is located. The server where the data source file is located is called a data server or GDS server.

Downloading the Client

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation tree on the left, choose **Management > Client Connections**.

Step 3 Select the GaussDB(DWS) client of the corresponding version from the drop-down list of **gsqL CLI Client**.

Choose a corresponding client version according to the cluster version and operating system to which the client is to be installed.

Table 4-5 gsql download links

OS Type	Applicable OS	Download Link	Verification File
Windows	Windows x86_64: <ul style="list-style-type: none"> Windows 7 or later Windows Server 2008 or later 	dws_8.1.x_gsql_for_windows.zip	dws_8.1.x_gsql_for_windows.zip.sha256
		dws_8.2.x_gsql_for_windows.zip	dws_8.2.x_gsql_for_windows.zip.sha256
Redhat x86_64	RHEL 6.4~7.6	dws_client_8.2.x_redhat_x64.zip	dws_client_8.2.x_redhat_x64.zip.sha256
		dws_client_8.1.x_redhat_x64.zip	dws_client_8.1.x_redhat_x64.zip.sha256
		dws_client_8.0.x_redhat_x64.zip	dws_client_8.0.x_redhat_x64.zip.sha256
SUSE x86_64	SLES 11.1~11.4, SLES 12.0~12.3	dws_client_8.2.x_suse_x64.zip	dws_client_8.2.x_suse_x64.zip.sha256
		dws_client_8.1.x_suse_x64.zip	dws_client_8.1.x_suse_x64.zip.sha256
		dws_client_8.0.x_suse_x64.zip	dws_client_8.0.x_suse_x64.zip.sha256
Euler Kunpeng_64	EulerOS 2.0 SP8	dws_client_8.1.x_euler_kunpeng_x64.zip	dws_client_8.1.x_euler_kunpeng_x64.zip.sha256
Redhat Kunpeng_64	CentOS-7.6-aarch64 and NeoKylin-7.6-aarch64 (adapted to Kunpeng 920 CPU)	dws_client_8.1.x_redhat_kunpeng_x64.zip	dws_client_8.1.x_redhat_kunpeng_x64.zip.sha256

 NOTE

The CPU architecture of the client must be the same as that of the cluster. If the cluster uses x86 servers, select an x86 client.

Step 4 Click **Download** to download the gsql tool matching the 8.1.x cluster version. Click **Historical Version** to download the gsql tool corresponding to the cluster version.

- You are advised to download the gsql tool that matches the cluster version. That is, use gsql 8.1.x for clusters of 8.1.0 or later, and use gsql 8.2.x for clusters of 8.2.0 or later.
- The following table describes the files and folders in the Linux gsql tool package.

Table 4-6 Files and folders in the Linux gsql tool package

File or Folder	Description
bin	This folder holds the Linux executable files for gsql, which include tools gsql, GDS, gs_dump, gs_dumpall, and gs_restore. For details, see Server Tool .
gds	This folder contains the files of the GDS data service tool. The GDS tool is used for parallel data loading and can import the data files stored in a common file system to a GaussDB(DWS) database.
lib	This folder contains the lib library required for executing the gsql client.
sample	This folder contains the following directories and files: <ul style="list-style-type: none">– setup.sh: script file for configuring the AK/SK before using gsql to import sample data– tpcds_load_data_from_obs.sql: script file for importing the TPC-DS sample data using the gsql client– query_sql directory: script file for querying the TPC-DS sample data
gsql_env.sh	Script file for configuring environment variables before running the gsql client.

- The following table describes the files and folders in the Windows gsql tool package.

Table 4-7 Files and folders in the Windows gsql tool package

File or Folder	Description
x64	This folder contains the 64-bit Windows gsql execution binary file and the dynamic library.

File or Folder	Description
x86	This folder contains the 32-bit Windows gsql execution binary file and the dynamic library.

 NOTE

In the cluster list on the **Clusters > Dedicated Clusters** page, click the name of the specified cluster to go to the **Cluster Information** page and view the cluster version.

----End

4.4.2 Using the Linux gsql Client to Connect to a Cluster

This section describes how to connect to a database through an SQL client after you create a data warehouse cluster and before you use the cluster's database. GaussDB(DWS) provides the Linux gsql client that matches the cluster version for you to access the cluster through the cluster's public or private network address.

The gsql command line client provided by GaussDB(DWS) runs on Linux. Before using it to remotely connect to a GaussDB(DWS) cluster, you need to prepare a Linux server for installing and running the gsql client. If you use a public network address to access the cluster, you can install the Linux gsql client on your own Linux server. Ensure that the Linux server has a public network address. If no EIPs are configured for your GaussDB(DWS) cluster, you are advised to create a Linux ECS for convenience purposes. For more information, see [\(Optional\) Preparing an ECS as the gsql Client Server](#).

(Optional) Preparing an ECS as the gsql Client Server

For details about how to purchase an ECS, see [Purchasing an ECS](#) in the *Elastic Cloud Server Getting Started*.

The created ECS must meet the following requirements:

- ECS and GaussDB(DWS) clusters must belong to the same region and AZ.
- If you use the gsql client provided by GaussDB(DWS) to connect to the GaussDB(DWS) cluster, the ECS image must meet the following requirements:

The image's OS must be one of the following Linux OSs supported by the gsql client:

- The **Redhat x86_64** client can be used on the following OSs:
 - RHEL 6.4~7.6
 - CentOS 6.4~7.4
 - EulerOS 2.3
- The **SUSE x86_64** client can be used on the following OSs:
 - SLES 11.1~11.4

- SLES 12.0~12.3
- The **Euler Kunpeng_64** client can be used on the following OS:
 - EulerOS 2.8
- The **Stream Euler x86_64** client can be used on the following OS:
EulerOS 2.2
- The **Stream Euler Kunpeng_64** client can be used on the following OS:
 - EulerOS 2.8
- If the client accesses the cluster using the private network address, ensure that the created ECS is in the same VPC as the GaussDB(DWS) cluster.
For details about VPC operations, see [VPC and Subnet](#) in the *Virtual Private Cloud User Guide*.
- If the client accesses the cluster using the public network address, ensure that both the created ECS and GaussDB(DWS) cluster have an EIP.
When purchasing an ECS, set **EIP** to **Buy now** or **Specify**.
- The security group rules of the ECS must enable communication between the ECS and the port that the GaussDB(DWS) cluster uses to provide services.
For details about security group operations, see [Security Group](#) in the *Virtual Private Cloud User Guide*.
Ensure that the security group of the ECS contains rules meeting the following requirements. If the rules do not exist, add them to the security group:
 - **Transfer Direction: Outbound**
 - **Protocol:** The protocol must contain TCP. For example, **TCP** or **All**.
 - **Port:** The value must contain the database port that provides services in the GaussDB(DWS) cluster. For example, set this parameter to **1-65535** or a specific GaussDB(DWS) database port.
 - **Destination:** The IP address set here must contain the IP address of the GaussDB(DWS) cluster to be connected. **0.0.0.0/0** indicates any IP address.
- The security group rules of the data warehouse cluster must ensure that GaussDB(DWS) can receive network access requests from clients.
Ensure that the cluster's security group contains rules meeting the following requirements. If the rules do not exist, add them to the security group:
 - **Transfer Direction: Inbound**
 - **Protocol:** The protocol must contain TCP. For example, **TCP** or **All**.
 - **Port:** Set this parameter to the servicing database port of the GaussDB(DWS) cluster. Example: 8000.
 - **Source IP Address:** The IP address set here must contain the IP address of the GaussDB(DWS) client host. Example: 192.168.0.10/32.

Downloading the Linux gsql Client and Connecting to a Cluster

- Step 1** Download the Linux gsql client by referring to [Downloading the Client](#), and use an SSH file transfer tool (such as WinSCP) to upload the client to a target Linux server.

You are advised to download the gsql tool that matches the cluster version. That is, use gsql 8.1.x for clusters of 8.1.0 or later, and use gsql 8.2.x for clusters of 8.2.0 or later. To download gsql 8.2.x, replace **dws_client_8.1.x_redhat_x64.zip** with **dws_client_8.2.x_redhat_x64.zip**. The **dws_client_8.1.x_redhat_x64.zip** is used as an example.

The user who uploads the client must have the full control permission on the target directory on the host to which the client is uploaded.

Step 2 Use the SSH tool to remotely manage the host where the client is installed.

For details about how to log in to an ECS, see "[Login Using an SSH Password](#)" in the *Elastic Cloud Server User Guide*.

Step 3 (Optional) To connect to the cluster in SSL mode, configure SSL authentication parameters on the host where the client is installed. For details, see [Establishing Secure TCP/IP Connections in SSL Mode](#).

 **NOTE**

The SSL connection mode is more secure than the non-SSL mode. You are advised to connect the client to the cluster in SSL mode.

Step 4 Run the following commands to decompress the client:

```
cd <Path for saving the client>  
unzip dws_client_8.1.x_redhat_x64.zip
```

In the preceding commands:

- *<Path_for_storing_the_client>*: Replace it with the actual path.
- *dws_client_8.1.x_redhat_x64.zip*: This is the client tool package name of **RedHat x86**. Replace it with the actual name.

Step 5 Run the following command to configure the GaussDB(DWS) client:

```
source gsql_env.sh
```

If the following information is displayed, the gsql client is successfully configured:

```
All things done.
```

Step 6 Connect to the database in the GaussDB(DWS) cluster using the gsql client. Replace the values of each parameter with actual values.

```
gsql -d <Database_name> -h <Cluster_address> -U <Database_user> -p <Database_port> -W  
<Cluster_password> -r
```

The parameters are described as follows:

- *Database_name*: Enter the name of the database to be connected. If you use the client to connect to the cluster for the first time, enter the default database **gaussdb**.
- *Cluster_address*: For details about how to obtain this address, see [Obtaining the Cluster Connection Address](#). If a public network address is used for connection, set this parameter to **Public Network Address** or **Public Network Domain Name**. If a private network address is used for connection, set this parameter to **Private Network Address** or **Private Network Domain Name**. If ELB is used for connection, set this parameter to **ELB Address**.
- *Database_user*: Enter the username of the cluster's database. If you use the client to connect to the cluster for the first time, set this parameter to the

default administrator configured during cluster creation, for example, **dbadmin**.

- *Database_port*: Enter the database port set during cluster creation.

For example, run the following command to connect to the default database **gaussdb** in the GaussDB(DWS) cluster:

```
gsql -d gaussdb -h 10.168.0.74 -U dbadmin -p 8000 -W password -r
```

If the following information is displayed, the connection succeeded:

```
gaussdb=>
```

```
----End
```

gsql Command Reference

For more information about the gsql commands, see the [Data Warehouse Service \(DWS\) Tool Guide](#).

(Optional) Importing TPC-DS Sample Data Using gsql

GaussDB(DWS) users can import data from external sources to data warehouse clusters. This section describes how to import sample data from OBS to a data warehouse cluster and perform querying and analysis operations on the sample data. The sample data is generated based on the standard TPC-DS benchmark test.

TPC-DS is the benchmark for testing the performance of decision support. With TPC-DS test data and cases, you can simulate complex scenarios, such as big data set statistics, report generation, online query, and data mining, to better understand functions and performance of database applications.

- Step 1** Use the SSH remote connection tool to log in to the server where the gsql client is installed and go to the gsql directory. The **/opt** directory is used as an example for storing the gsql client.

```
cd /opt
```

- Step 2** Switch to the specified directory and set the AK and SK for importing sample data and the OBS access address.

```
cd sample  
/bin/bash setup.sh -ak <Access_Key_Id> -sk <Secret_Access_Key> -obs_location obs.eu-dublin.myhuaweicloud.com
```

If the following information is displayed, the settings are successful:

```
setup successfully!
```

NOTE

<Access_Key_Id> and <Secret_Access_Key>: indicate the AK and SK, respectively. For how to obtain the AK and SK, see [Creating Access Keys \(AK and SK\)](#). Then, replace the parameters in the statements with the obtained values.

- Step 3** Go back to previous directory and run the gsql environment variables.

```
cd ..  
source gsql_env.sh  
cd bin
```


Step 4 Import the sample data to the data warehouse.

Command format:

```
gsql -d <Database name> -h <Public network address of the cluster> -U <Administrator> -p <Data warehouse port number> -f <Path for storing the sample data script> -r
```

Sample command:

```
gsql -d gaussdb -h 10.168.0.74 -U dbadmin -p 8000 -f /opt/sample/tpcds_load_data_from_obs.sql -r
```

 **NOTE**

In the preceding command, sample data script **tpcds_load_data_from_obs.sql** is stored in the sample directory (for example, **/opt/sample/**) of the GaussDB(DWS) client.

After you enter the administrator password and successfully connect to the database in the cluster, the system will automatically create a foreign table to associate the sample data outside the cluster. Then, the system creates a target table for saving the sample data and imports the data to the target table using the foreign table.

The time required for importing a large dataset depends on the current GaussDB(DWS) cluster specifications. Generally, the import takes about 10 to 20 minutes. If information similar to the following is displayed, the import is successful.

```
Time:1845600.524 ms
```

Step 5 In the Linux command window, run the following commands to switch to a specific directory and query the sample data:

```
cd /opt/sample/query_sql/  
/bin/bash tpcds100x.sh
```

Step 6 Enter the cluster's public network IP address, access port, database name, user who accesses the database, and password of the user as prompted.

- The default database name is **gaussdb**.
- Use the administrator username and password configured during cluster creation as the username and password for accessing the database.

After the query is complete, a directory for storing the query result, such as **query_output_20170914_072341**, will be generated in the current query directory, for example, **sample/query_sql/**.

----End

4.4.3 Using the Windows gsql Client to Connect to a Cluster

This section describes how to connect to a database through an SQL client after you create a data warehouse cluster and before you use the cluster's database. GaussDB(DWS) provides the Windows gsql client that matches the cluster version for you to access the cluster through the cluster's public or private network address.

Procedure

Step 1 Install and run the gsql client on the local Windows server (in Windows CLI). Windows Server 2008/Windows 7 and later are supported.

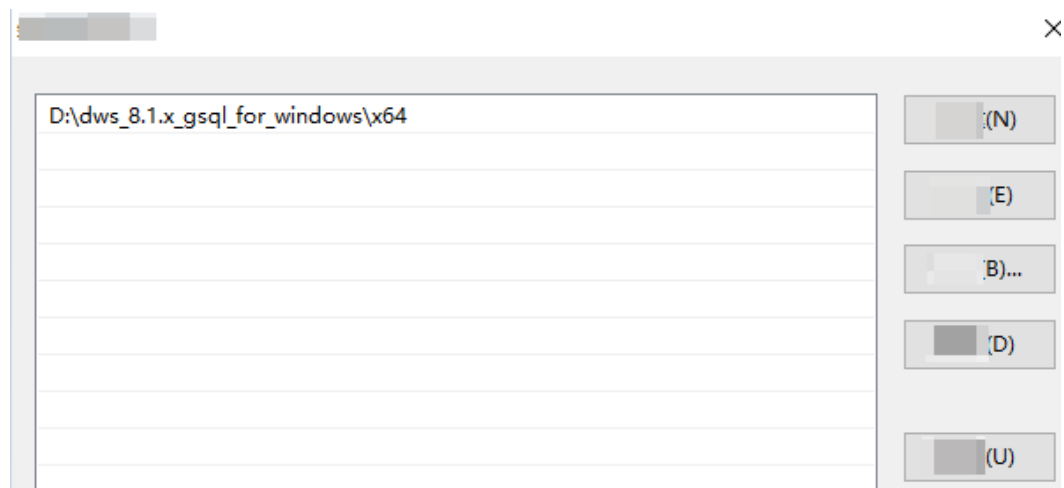
- Step 2** Download the Windows gsql client by referring to [Downloading the Client](#) and decompress the package to a local folder.
- Step 3** On the local server, click **Start**, search for **cmd**, and run the program as the administrator. Alternatively, press **Win+R** to open the Windows CLI.
- Step 4** Set environment variables. For a 32-bit OS, select the **x86** folder. For a 64-bit OS, select the **x64** folder.

Method 1: Configure environment variables in the Windows CLI. Open the command prompt and run the **set path=<window_gsql>%path%** command, where **<window_gsql>** indicates the folder path where the Windows gsql client was decompressed to in the previous step. For example:

```
set path=C:\Users\xx\Desktop\dws_8.1.x_gsql_for_windows\x64;%path%
```

Method 2: In the **Control Panel** window, search for **System** and click **View advanced system settings**. Click the **Advanced** tab, and click **Environment Variables**. Select the **Path** parameter and click **Edit**. Add the gsql path in the parameter value. For example:

Figure 4-4 Configuring Windows environment variables



- Step 5** (Optional) To connect to the cluster in SSL mode, configure SSL authentication parameters on the server where the client is installed. For details, see [Establishing Secure TCP/IP Connections in SSL Mode](#).

NOTE

The SSL connection mode is more secure than the non-SSL mode. You are advised to connect the client to the cluster in SSL mode.

- Step 6** In the Windows CLI, run the following command to connect to the database in the GaussDB(DWS) cluster using the gsql client:

```
gsql -d <Database_name> -h <Cluster_address> -U <Database_user> -p <Database_port> -W <Cluster_password> -r
```

The parameters are as follows:

- **Database name:** Enter the name of the database to be connected. If you use the client to connect to the cluster for the first time, enter the default database **gaussdb**.

- **Cluster address:** For details about how to obtain this address, see [Obtaining the Cluster Connection Address](#). If a public network address is used for connection, set this parameter to the public network domain name. If a private network address is used for connection, set this parameter to the private network domain name. If ELB is used for connection, set this parameter to **ELB Address**.
- **Database user:** Enter the username of the cluster's database. If you use the client to connect to the cluster for the first time, set this parameter to the default administrator configured during cluster creation, for example, **dbadmin**.
- **Database port:** Enter the database port set during cluster creation.

For example, run the following command to connect to the default database **gaussdb** in the GaussDB(DWS) cluster:

```
gsql -d gaussdb -h 10.168.0.74 -U dbadmin -p 8000 -W password -r
```

If the following information is displayed, the connection succeeded:

```
gaussdb=>
```

----End

Precautions

1. The default character encoding of the Windows command prompt is GBK, and the default value of **client_encoding** of Windows gsql is **GBK**. Some characters encoded using UTF-8 cannot be displayed in Windows gsql.
Suggestion: Ensure the file specified using **-f** uses UTF-8 encoding, and set the default encoding format to **UTF-8 (set client_encoding='utf-8')**;

2. Paths in Windows gsql must be separated by slashes (/), or an error will be reported. In a meta-command, the backslash (\) indicates the start of a meta-command. If the backslash is enclosed in single quotation marks ('\'), it is used for escape.

```
gaussdb=> \i D:\test.sql
D:: Permission denied
postgres=> \i D:/test.sql
id
----
1
(1 row)
```

3. To use the **\!** metacommand to run a system command in Windows gsql, be sure to use the path separator required by the system command. Generally, the path separator is a backslash (\).

```
gaussdb=> \! type D:/test.sql
Incorrect syntax.
gaussdb=> \! type D:\test.sql
select 1 as id;
```

4. Windows gsql does not support the **\parallel** meta-command.

```
gaussdb=> \parallel
ERROR: "\parallel" is not supported in Windows.
```

5. In Linux shell, single quotation marks (') and double quotation marks (") can be used to enclose strings. In Windows, only double quotation marks can be used.

```
gsql -h 192.168.233.189 -p 8109 -d postgres -U odbcuser -W password -c "select 1 as id"
id
----
```

- ```
1
(1 row)
```
- If single quotation marks are used, an error will be reported and the input will be ignored.
- ```
gsql -h 192.168.233.189 -p 8109 -d postgres -U odbcuser -W password -c 'select 1 as id'
gsql: warning: extra command-line argument "1" ignored
gsql: warning: extra command-line argument "as" ignored
gsql: warning: extra command-line argument "id" ignored
ERROR: unterminated quoted string at or near "'select"
LINE 1: 'select
```
- If Windows gsql is idle for a long time after a connection is established, the connection session times out, and an SSL error is reported. In this case, you need to log in again. The following error is reported:
SSL SYSCALL error: Software caused connection abort (0x00002745/10053), remote datanode <NULL>, error: Result too large
 - In Windows, press **Ctrl+C** to exit gsql. If **Ctrl+C** are pressed during input, the input will be ignored and you will be forced to exit gsql.
Enter **as** and press **Ctrl+C**. After **\q** is displayed, exit gsql.
gaussdb=> select 1
gaussdb=> as \q
 - Windows gsql cannot connect to a database using the LATIN1 character encoding. The error information is as follows:
gsql: FATAL: conversion between GBK and LATIN1 is not supported
 - The location of the **gsqlrc.conf** file:
The default **gsqlrc** path is **%APPDATA%/postgresql/gsqlrc.conf**. You can also set the path using the **PSQLRC** variable.
set PSQLRC=C:\Users\xx\Desktop\dws_8.1.x_gsql_for_windows\x64\gsqlrc.conf

gsql Command Reference

For more information about the gsql commands, see the [Data Warehouse Service \(DWS\) Tool Guide](#).

4.4.4 Establishing Secure TCP/IP Connections in SSL Mode

GaussDB(DWS) supports the standard SSL. As a highly secure protocol, SSL authenticates bidirectional identification between the server and client using digital signatures and digital certificates to ensure secure data transmission. To support SSL connection, GaussDB(DWS) has obtained the formal certificates and keys for the server and client from the CA certification center. It is assumed that the key and certificate for the server are **server.key** and **server.crt** respectively; the key and certificate for the client are **client.key** and **client.crt** respectively, and the name of the CA root certificate is **cacert.pem**.

The SSL connection mode is more secure. By default, the SSL feature in a cluster allows SSL and non-SSL connections from the client. For security purposes, you are advised to connect to the cluster via SSL from the client. Ensure the certificate, private key, and root certificate of the GaussDB(DWS) server have been configured by default. To forcibly use an SSL connection, configure the **require_ssl** parameter in the **Require SSL Connection** area of the cluster's **Security Settings** page on the GaussDB(DWS) management console. Require SSL Connection on the Security Settings page of the cluster. For more information, see [Configuring SSL Connection](#) and [Combinations of SSL Connection Parameters on the Client and Server](#).

The client or JDBC/ODBC driver needs to use SSL connection. Configure related SSL connection parameters in the client or application code. The GaussDB(DWS) console provides the SSL certificate required by the client. The SSL certificate contains the default certificate, private key, root certificate, and private key password encryption file required by the client. Download the SSL certificate to the host where the client is installed, and specify the path of the certificate on the client. For more information, see [Configuring Digital Certificate Parameters Related to SSL Authentication on the gsql Client](#) and [SSL Authentication Modes and Client Parameters](#).

NOTE

Using the default certificate may pose security risks. To improve system security, you are advised to periodically change the certificate to prevent password cracking. If you need to replace the certificate, contact the database customer service.

Configuring SSL Connection

Prerequisites

- Changes made to security configuration parameters require a cluster restart to take effect. Otherwise, the cluster will be temporarily unavailable.
- To modify the cluster's security configuration, ensure that the following conditions are met:
 - The cluster status is **Available** or **Unbalanced**.
 - The **Task Information** cannot be set to **Creating snapshot**, **Scaling out**, **Configuring**, or **Restarting**.

Procedure


Step 1 Log in to the GaussDB(DWS) console.


Step 2 In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.

Step 3 In the cluster list, click the name of a cluster. On the page that is displayed, click **Security Settings**.

By default, **Configuration Status** is **Synchronized**, which indicates that the latest database result is displayed.

Step 4 In the **SSL Connection** area, enable **Require SSL Connection** (recommended).

 indicates the function is enabled. The **require_ssl** is set to **1**, indicating that the server forcibly requires the SSL connection.

 indicates the function is disabled (default value). The **require_ssl** parameter is set to **0**, indicating that the server does not require SSL connections. For details about how to configure the **require_ssl** parameter, see [require_ssl \(Server\)](#).

 NOTE

- If the gsql client or ODBC driver provided by GaussDB(DWS) is used, GaussDB(DWS) supports the TLSv1.2 SSL protocol.
- If the JDBC driver provided by GaussDB(DWS) is used, GaussDB(DWS) supports SSL protocols, such as SSLv3, TLSv1, TLSv1.1, and TLSv1.2. The SSL protocol used between the client and the database depends on the Java Development Kit (JDK) version used by the client. Generally, JDK supports multiple SSL protocols.

Step 5 Click **Apply**.

The system automatically saves the SSL connection settings. On the **Security Settings** page, **Configuration Status** is **Applying**. After **Configuration Status** changes to **Synchronized**, the settings have been saved and taken effect.

----End

Configuring Digital Certificate Parameters Related to SSL Authentication on the gsql Client

After a data warehouse cluster is deployed, the SSL authentication mode is enabled by default. The server certificate, private key, and root certificate have been configured by default. You need to configure the client parameters.

Step 1 Log in to the GaussDB(DWS) console. In the navigation pane, choose **Client Connections**.

Step 2 In the **Driver** area, click **download an SSL certificate**.

Step 3 Use a file transfer tool (such as WinSCP) to upload the SSL certificate to the host where the client is installed.

For example, save the downloaded certificate **dws_ssl_cert.zip** to the **/home/dbadmin/dws_ssl/** directory.

Step 4 Use an SSH remote connection tool (such as PuTTY) to log in to the host where the gsql client is installed and run the following commands to go to the directory where the SSL certificate is stored and decompress the SSL certificate:

```
cd /home/dbadmin/dws_ssl/  
unzip dws_ssl_cert.zip
```

Step 5 Run the export command and configure digital certificate parameters related to SSL authentication on the host where the gsql client is installed.

There are two SSL authentication modes: bidirectional authentication and unidirectional authentication. The client environment variables to be configured vary according to the authentication mode. For details, see [SSL Authentication Modes and Client Parameters](#).

The following parameters must be configured for bidirectional authentication:

```
export PGSSLCERT="/home/dbadmin/dws_ssl/sslcert/client.crt"  
export PGSSLKEY="/home/dbadmin/dws_ssl/sslcert/client.key"  
export PGSSLMODE="verify-ca"  
export PGSSLROOTCERT="/home/dbadmin/dws_ssl/sslcert/cacert.pem"
```

The following parameters must be configured for unidirectional authentication:

```
export PGSSLMODE="verify-ca"  
export PGSSLROOTCERT="/home/dbadmin/dws_ssl/sslcert/cacert.pem"
```

NOTICE

- You are advised to use bidirectional authentication for security purposes.
- The environment variables configured for a client must contain the absolute file paths.

Step 6 Change the client private key permissions.

The permissions on the client's root certificate, private key, certificate, and encrypted private key file must be **600**. If the permissions do not meet the requirement, the client cannot connect to the cluster in SSL mode.

```
chmod 600 client.key
chmod 600 client.crt
chmod 600 client.key.cipher
chmod 600 client.key.rand
chmod 600 cacert.pem
```

----End

SSL Authentication Modes and Client Parameters

There are two SSL authentication modes: bidirectional authentication and unidirectional authentication. Table [Table 4-8](#) shows the differences between these two modes. You are advised to use bidirectional authentication for security purposes.

Table 4-8 Authentication modes

Authentication Mode	Description	Environment Variables Configured on a Client	Maintenance
Bidirectional authentication (recommended)	The client verifies the server's certificate and the server verifies the client's certificate. The connection can be set up only after the verifications are successful.	Set the following environment variables: <ul style="list-style-type: none"> • PGSSLCERT • PGSSLKEY • PGSSLROTCERT • PGSSLMODE 	This authentication mode is applicable to scenarios that require high data security. When using this mode, you are advised to set the PGSSLMODE client variable to verify-ca for network data security purposes.

Authentication Mode	Description	Environment Variables Configured on a Client	Maintenance
Unidirectional authentication	The client verifies the server's certificate, whereas the server does not verify the client's certificate. The server loads the certificate information and sends it to the client. The client verifies the server's certificate according to the root certificate.	Set the following environment variables: <ul style="list-style-type: none"> PGSSLOTCERT PGSSLMODE 	To prevent TCP-based security attacks, you are advised to use the SSL certificate authentication. In addition to configuring the client root certificate, you are advised to set the PGSSLMODE variable to verify-ca on the client.

Configure environment variables related to SSL authentication on the client. For details, see [Table 4-9](#).

 **NOTE**

The path of environment variables is set to `/home/dbadmin/dws_ssl/` as an example. Replace it with the actual path.

Table 4-9 Client parameters

Environment Variable	Description	Value Description
PGSSLCERT	Specifies the certificate files for a client, including the public key. Certificates prove the legal identity of the client and the public key is sent to the remote end for data encryption.	The absolute path of the files must be specified, for example: <pre>export PGSSLCERT='/home/dbadmin/dws_ssl/sslcert/client.crt'</pre> (No default value)
PGSSLKEY	Specifies the client private key file used to decrypt the digital signatures and the data encrypted using the public key.	The absolute path of the files must be specified, for example: <pre>export PGSSLKEY='/home/dbadmin/dws_ssl/sslcert/client.key'</pre> (No default value)

Environment Variable	Description	Value Description
PGSSLMODE	Specifies whether to negotiate with the server about SSL connection and specifies the priority of the SSL connection.	<p>Values and meanings:</p> <ul style="list-style-type: none"> • disable: only tries to establish a non-SSL connection. • allow: tries to establish a non-SSL connection first, and then an SSL connection if the first attempt fails. • prefer: tries to establish an SSL connection first, and then a non-SSL connection if the first attempt fails. • require: only tries to establish an SSL connection. If there is a CA file, perform the verification according to the scenario in which the parameter is set to verify-ca. • verify-ca: tries to establish an SSL connection and check whether the server certificate is issued by a trusted CA. • verify-full: GaussDB(DWS) does not support this mode. <p>Default value: prefer</p> <p>NOTE When an external client accesses a cluster, the error message "ssl SYSCALL error" is displayed on some nodes. In this case, run export PGSSLMODE="allow" or export PGSSLMODE="prefer".</p>
PGSSLROOTCERT	Specifies the root certificate file for issuing client certificates. The root certificate is used to verify the server certificate.	<p>The absolute path of the files must be specified, for example: <code>export PGSSLROOTCERT='/home/dbadmin/dws_ssl/sslcert/certca.pem'</code></p> <p>Default value: null</p>
PGSSLCRL	Specifies the certificate revocation list file, which is used to check whether a server certificate is in the list. If the certificate is in the list, it is invalid.	<p>The absolute path of the files must be specified, for example: <code>export PGSSLCRL='/home/dbadmin/dws_ssl/sslcert/sslcert-file.crt'</code></p> <p>Default value: null</p>

Combinations of SSL Connection Parameters on the Client and Server

Whether the client uses the SSL encryption connection mode and whether to verify the server certificate depend on client parameter **sslmode** and server (cluster) parameters **ssl** and **require_ssl**. The parameters are as follows:

- **ssl (Server)**

The **ssl** parameter indicates whether to enable the SSL function. **on** indicates that the function is enabled, and **off** indicates that the function is disabled.

- The default value is **on** for clusters whose version is 1.3.1 or later, and you cannot set this parameter on the GaussDB(DWS) console.
- For clusters whose version is earlier than 1.3.1, the default value is **on**. You can set this parameter in the **SSL Connection** area on the cluster's **Security Settings** page of the GaussDB(DWS) console.

- **require_ssl (Server)**

The **require_ssl** parameter specifies whether the server forcibly requires SSL connection. This parameter is valid only when **ssl** is set to **on**. **on** indicates that the server forcibly requires SSL connection. **off** indicates that the server does not require SSL connection.

- The default value is **off** for clusters whose version is 1.3.1 or later. You can set the **require_ssl** parameter in the **Require SSL Connection** area of the cluster's **Security Settings** page on the GaussDB(DWS) console.
- For clusters whose version is earlier than 1.3.1, the default value is **off**, and you cannot set this parameter on the GaussDB(DWS) console.

- **sslmode (Client)**

You can set this parameter in the SQL client tool.

- In the gsql command line client, this parameter is the **PGSSLMODE** parameter.
- On the Data Studio client, this parameter is the **SSL Mode** parameter.

The combinations of client parameter **sslmode** and server parameters **ssl** and **require_ssl** are as follows.

Table 4-10 Combinations of SSL connection parameters on the client and server

ssl (Server)	sslmode (Client)	require_ssl (Server)	Result
on	disable	on	The server requires SSL, but the client disables SSL for the connection. As a result, the connection cannot be set up.
	disable	off	The connection is not encrypted.
	allow	on	The connection is encrypted.
	allow	off	The connection is not encrypted.
	prefer	on	The connection is encrypted.
	prefer	off	The connection is encrypted.

ssl (Server)	sslmode (Client)	require_ssl (Server)	Result
	require	on	The connection is encrypted.
	require	off	The connection is encrypted.
	verify-ca	on	The connection is encrypted and the server certificate is verified.
	verify-ca	off	The connection is encrypted and the server certificate is verified.
off	disable	on	The connection is not encrypted.
	disable	off	The connection is not encrypted.
	allow	on	The connection is not encrypted.
	allow	off	The connection is not encrypted.
	prefer	on	The connection is not encrypted.
	prefer	off	The connection is not encrypted.
	require	on	The client requires SSL, but SSL is disabled on the server. Therefore, the connection cannot be set up.
	require	off	The client requires SSL, but SSL is disabled on the server. Therefore, the connection cannot be set up.
	verify-ca	on	The client requires SSL, but SSL is disabled on the server. Therefore, the connection cannot be set up.
	verify-ca	off	The client requires SSL, but SSL is disabled on the server. Therefore, the connection cannot be set up.

4.5 Using the JDBC and ODBC Drivers to Connect to a Cluster

4.5.1 Development Specifications

If the connection pool mechanism is used during application development, comply with the following specifications: If you do not do so, the status of connections in the connection pool will remain, which affects subsequent operations using the connection pool.

- If the GUC parameter is set in a connection, you must execute **SET SESSION AUTHORIZATION DEFAULT;RESET ALL;** to clear the connection status before returning the connection to the connection pool.
- If a temporary table is used, it must be deleted before the connection is returned to the connection pool.

4.5.2 JDBC Version Description

Version 8.3.1.200

New features

- JDBC load balancing can automatically detect the live CN list.
 - The **cnListRefreshSwitch** parameter is added to enable automatic detection of the live CN list. The default value is **off**.
 - The **cnListRefreshDelay** parameter is added to specify the delay for enabling automatic detection of the live CN list. The default value is 1,800,000 milliseconds. This parameter only works when **cnListRefreshSwitch** is set to **on**.
 - The **cnListRefreshPeriod** parameter is added to specify the period for automatically detecting the live CN list. The default value is 1,800,000 milliseconds. This parameter only works when **cnListRefreshSwitch** is set to **on**.
- JDBC supports certificate revocation.

The **sslCrl** parameter is added to specify the path of the revoked certificate. By default, this parameter is left blank.
- JDBC supports database reconnection.
 - The **autoReconnect** parameter is added to specify whether to enable automatic database reconnection. The default value is **false**.
 - The **reConnectCount** parameter is added to specify the number of automatic reconnection times. The default value is **10**. This parameter takes effect when **autoReconnect** is set to **true**. If the number of connection attempts exceeds this parameter's value, reconnection will fail.

Version 8.3.0.202

- New features:

The **tcpKeepAlive** configuration is added. When **tcpKeepAlive** is set to **true**, the following parameters take effect:

Default value:

 - a. **TCP_KEEPIDLE=30**: The detection starts after the connection is idle for 30 seconds.
 - b. **TCP_KEEPCOUNT=9**: A total of nine detections are performed.
 - c. **TCP_KEEPINTERVAL=30**: The detection interval is 30s.

NOTE

The JDK varies according to the operating system. Some platforms, such as Windows, Red Hat, and SUSE, may not support this parameter.

- Fixed vulnerability:
CVE-2024-1597

Version 8.3.0.201

- Fixed bug:
Multiple functions cannot be automatically split when they are executed at a time.

Version 8.3.0

- Fixed bug:
loadBalanceHosts=false does not take effect.

Version 8.2.1.300

- Fixed bug:
The NVARCHAR array type is incompatible.
- Fixed vulnerability:
CVE-2022-41946

Version 8.2.1.1

The **defaultQueryMetaData** parameter is added to specify whether to query SQL metadata by default. The default value is **false**.

JDBC supports the raw type, which requires the querying of metadata. If you want to use JDBC to perform operations on the raw type, set **defaultQueryMetaData** to true.

If this parameter is enabled, **prepareStatement** is incompatible with the syntax **create table as**. You can replace it by **Statement**.

Version 8.2.1

- Fixed bug:
 - a. An error is reported when **reWriteBatchedInserts** is used to insert data in batches.
 - b. "Invalid input syntax for type oid: 03032VLM" is reported when data is imported from Spark to GaussDB(DWS).

Version 8.2.0

- New feature: Compatibility with the Oracle Raw data type. The usage is as follows:

- Insertion or Modification

```
byte[] bytes = oracleResultSet.getBytes(2)
prepareStatement.setBytes(bytes)
// Or
prepareStatement.setObject(bytes)
```
- Query

```
resultSet.getBytes()
resultSet.getObject()
```

- Fixed bug:
The field length obtained by the `getColumnDisplaySize()` method is incorrect.
- Fixed vulnerabilities:
CVE-2022-26520
CVE-2022-31197

Version 8.1.3.100

- New features
The `nvarchar2` object can be obtained through `resultSet.getNString`.
- Fixed vulnerabilities:
The dependency package `fastjson` is upgraded to 1.2.83.

Version 8.1.3

Upgrade to the open source version 42.2.23.

- New features
 - The `nvarchar2` type is supported.
 - The `nvarchar2` object can be obtained through `resultSet.getObject`.
- Fixed vulnerabilities
CVE-2022-21724

NOTE

For JDBC 8.1.3 and later versions, JDK 1.8 is required.

Version 8.1.1.300

- New features
 - The `nvarchar2` type is supported.
 - The `nvarchar2` object can be obtained through `resultSet.getObject`.
- Fixed vulnerabilities

Version 8.1.1.100

- New features
By default, the driver reports the OS user. To disable this function, you can set **`connectionExtraInfo=false`**.

```
jdbc:postgresql://host:port/database?connectionExtraInfo=false
```
- Fixed vulnerabilities
Jackson was upgraded.

4.5.3 Downloading the JDBC or ODBC Driver

The JDBC or ODBC driver is used to connect to data warehouse clusters. You can download the JDBC or ODBC driver provided by GaussDB(DWS) from the management console or use the open-source JDBC or ODBC driver.

Open-Source JDBC or ODBC Driver

GaussDB(DWS) also supports open-source JDBC and ODBC drivers: PostgreSQL JDBC 9.3-1103 or later; PostgreSQL ODBC 09.01.0200 or later

Downloading the JDBC or ODBC Driver

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation tree on the left, choose **Client Connections**.

Step 3 In the **Driver** area, choose a driver that you want to download.

- **JDBC Driver**

Method 1:

Select **DWS JDBC Driver** and click **Download** to download the JDBC driver matching the current cluster version. The driver package name is **dws_8.1.x_jdbc_driver.zip**. After the package is decompressed, there will be two JAR packages **gsjdbc4.jar** and **gsjdbc200.jar**.

- **gsjdbc4.jar**: The **gsjdbc4.jar** driver package is compatible with PostgreSQL. Its class names and class structures are the same as those of the PostgreSQL driver. Applications that run in PostgreSQL can be directly migrated to the current system.
- **gsjdbc200.jar**: If a JVM process needs to access PostgreSQL and GaussDB(DWS) at the same time, this driver package must be used. In this package, the main class name is **com.huawei.gauss200.jdbc.Driver** (that is, **org.postgresql** is replaced with **com.huawei.gauss200.jdbc**). The URL prefix of the database connection is **jdbc:gaussdb**. Other parameters are the same as those of **gsjdbc4.jar**.

If clusters of different versions are available, you will download the JDBC driver matching the earliest cluster version after clicking **Download**. If there is no cluster, you will download the JDBC driver of the earliest version after clicking **Download**. GaussDB(DWS) clusters are compatible with earlier versions of JDBC drivers.

Click **Historical Version** to download the corresponding JDBC driver version. You are advised to download the JDBC driver based on the cluster version.

The JDBC driver can be used on all platforms and depends on JDK 1.6 or later.

If you have clusters of different versions, the system displays a dialog box, prompting you to select the cluster version and download the driver corresponding to the cluster version. In the cluster list on the **Clusters > Dedicated Clusters** page, click the name of the specified cluster to go to the **Cluster Information** page and view the cluster version.

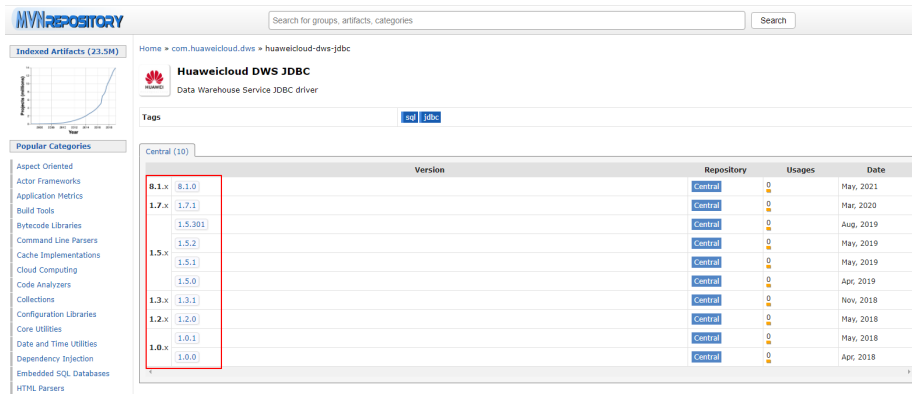
Table 4-11 JDBC driver download address

Driver	Download Link	Verification File
DWS JDBC Driver	dws_8.2.x_jdbc_driver.zip	dws_8.2.x_jdbc_driver.zip.sha256
	dws_8.1.x_jdbc_driver.zip	dws_8.1.x_jdbc_driver.zip.sha256

Method 2:

Download the SDK software package by configuring the Maven repository. Click **Add Maven Dependency**. The following page is displayed.

Figure 4-5 Maven page



In the list shown in **Figure 4-5**, the first column indicates the cluster version, and the second column indicates the version number of the GaussDB(DWS) JDBC driver package. Select the driver package based on the cluster version and go to the following page:

Figure 4-6 Maven dependency



Copy the Maven repository information and add it to the **pom.xml** file. For example, add the following code configuration to the **pom.xml** file:

- gsjdbc4.jar


```
<dependency>
  <groupId>com.huaweicloud.dws</groupId>
  <artifactId>huaweicloud-dws-jdbc</artifactId>
  <version>8.1.0</version>
</dependency>
```
- gsjdbc200.jar


```
<dependency>
  <groupId>com.huaweicloud.dws</groupId>
  <artifactId>huaweicloud-dws-jdbc</artifactId>
```



```
<version>8.1.1.1-200</version>
</dependency>
```

- **ODBC Driver**

Select a corresponding version and click **Download** to download the ODBC driver matching the current cluster version. If clusters of different versions are available, you will download the ODBC driver matching the earliest cluster version after clicking **Download**. If there is no cluster, you will download the ODBC driver of the earliest version after clicking **Download**. GaussDB(DWS) clusters are compatible with earlier versions of ODBC drivers.

Click **Historical Version** to download the corresponding ODBC driver version. You are advised to download the ODBC driver based on the cluster version.

 **NOTE**

The ODBC driver is incompatible with Windows Server 2016.

Table 4-12 ODBC driver download address

Applicable OS	Download Link	Verification File
Windows	dws_8.2.x_odbc_driver_for_windows.zip	dws_8.2.x_odbc_driver_for_windows.zip.sha256
	dws_8.1.x_odbc_driver_for_windows.zip	dws_8.1.x_odbc_driver_for_windows.zip.sha256
Redhat x86_64	dws_8.2.x_odbc_driver_for_x86_redhat.zip	dws_8.2.x_odbc_driver_for_x86_redhat.zip.sha256
	dws_8.1.x_odbc_driver_for_x86_redhat.zip	dws_8.1.x_odbc_driver_for_x86_redhat.zip.sha256
SUSE x86_64	dws_8.2.x_odbc_driver_for_x86_suse.zip	dws_8.1.x_odbc_driver_for_x86_suse.zip.sha256
	dws_8.1.x_odbc_driver_for_x86_suse.zip	dws_8.1.x_odbc_driver_for_x86_suse.zip.sha256

----End

4.5.4 Using JDBC to Connect to a Cluster

In GaussDB(DWS), you can use a JDBC driver to connect to a database on Linux or Windows. The driver can connect to the database through an ECS on the Huawei Cloud platform or over the Internet.

When using the JDBC driver to connect to the data warehouse cluster, determine whether to enable SSL authentication. SSL authentication is used to encrypt communication data between the client and the server. It safeguards sensitive

data transmitted over the Internet. You can download a self-signed certificate file on the GaussDB(DWS) console. To make the certificate take effect, you must configure the client program using the OpenSSL tool and the Java keytool.

NOTE

The SSL mode delivers higher security than the common mode. You are advised to enable SSL connection when using JDBC to connect to a GaussDB(DWS) cluster.

For details about how to use the JDBC API, see the official documentation.

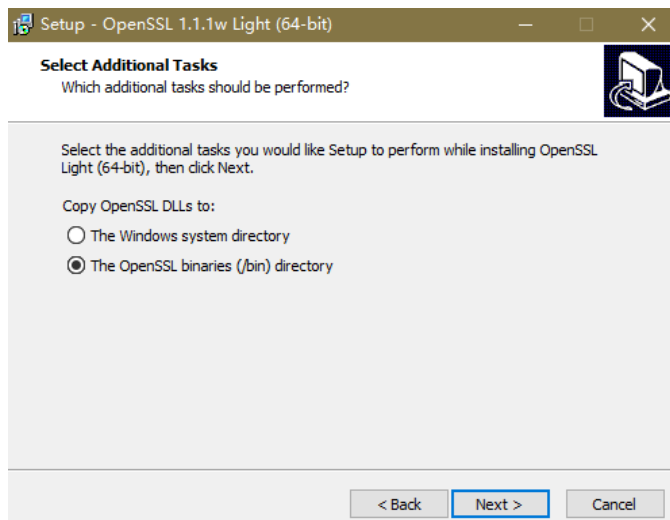
Prerequisites

- You have installed JDK 1.6 or later and configured environment variables.
- You have downloaded the JDBC driver. For details, see [Downloading the JDBC or ODBC Driver](#).
GaussDB(DWS) also supports open-source JDBC driver: PostgreSQL JDBC 9.3-1103 or later.
- You have downloaded the SSL certificate file. For details, see [Downloading an SSL Certificate](#).

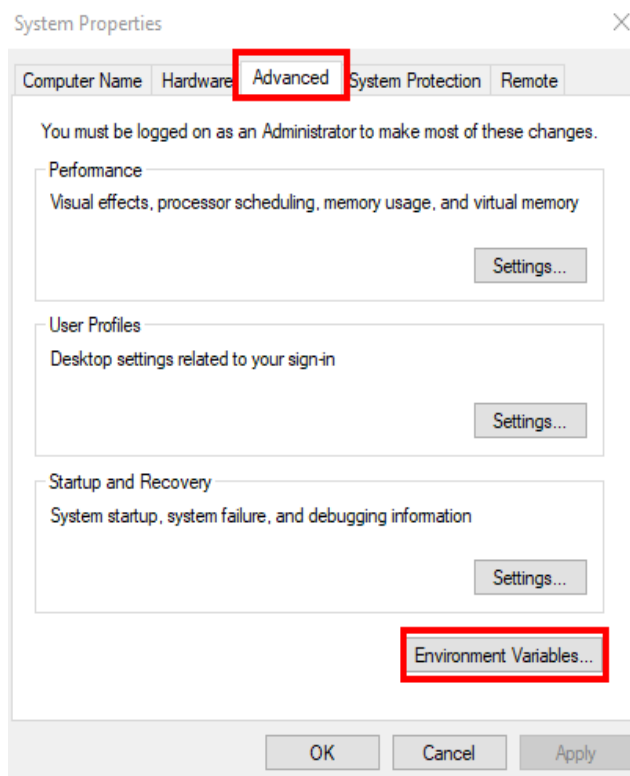
Using a JDBC Driver to Connect to a Database

The procedure for connecting to the database using a JDBC driver in a Linux environment is similar to that in a Windows environment. The following describes the connection procedure in a Windows environment.

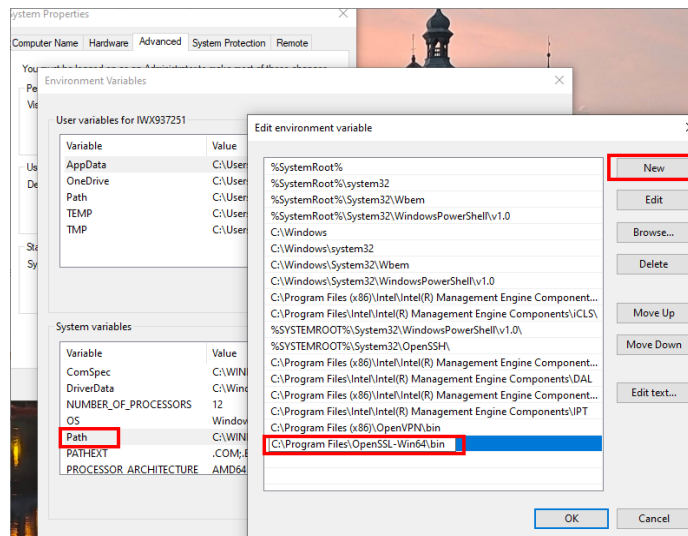
- Step 1** Determine whether you want to use the SSL mode to connect to the GaussDB(DWS) cluster.
- If yes, enable SSL connection by referring to [Configuring SSL Connection](#). SSL connection is enabled by default. Then go to [Step 2](#).
 - If no, disable SSL connection by referring to [Configuring SSL Connection](#) and go to [Step 4](#).
- Step 2** (Optional) On Linux, use WinSCP to upload the downloaded SSL certificate file to the Linux environment.
- Step 3** Configure the certificate to enable SSL connection.
1. Download the OpenSSL toolkit for Windows at <https://slproweb.com/products/Win32OpenSSL.html>. OpenSSL 3.0.0 is currently not supported. Download Win64 OpenSSL v1.1.1w Light instead.
 2. Double-click the installation package **Win64OpenSSL_Light-1_1_1w.exe** and install it to the default path on drive C. Copy the DLLs to the OpenSSL directory, as shown in the following figure. Retain the default settings in the remaining steps until the installation is successful.



3. Install an environment variable. Click **Start** in the lower left corner of the local PC, right-click **This PC**, choose **More > Properties > View advanced system settings**. Switch to the **Advanced** tab and click **Environment Variables**.



4. In the **System variables** area, double-click **Path** and click **New** in the window displayed. Add the OpenSSL **bin** path to the last line, for example, **C:\Program Files\OpenSSL-Win64\bin**, and click **OK**. Click **OK** again and the variable is configured successfully.



- Decompress the package to obtain the certificate file. Decompression path `C:\` is used as an example.

You are advised to store the certificate file in a path of the English version and can specify the actual path when configuring the certificate. If the path is incorrect, a message stating that the file does not exist will be prompted.

- Open **Command Prompt** and switch to the `C:\dws_ssl_cert\sslcert` path. Run the following commands to import the root license to the truststore:

```
openssl x509 -in cacert.pem -out cacert.crt.der -outform der  
keytool -keystore mytruststore -alias cacert -import -file cacert.crt.der
```

 - `cacert.pem` indicates the root certificate obtained after decompression.
 - `cacert.crt.der` indicates the generated intermediate file. You can store the file to another path and change the file name to your desired one.
 - `mytruststore` indicates the generated truststore name and `cacert` indicates the alias name. Both parameters can be modified.

Enter the truststore password as prompted and answer **y**.

- Convert the format of the client private key.

```
openssl pkcs12 -export -out client.pkcs12 -in client.crt -inkey client.key
```

Enter the client private key password **Gauss@MppDB**. Then enter and confirm the self-defined private key password.


- Import the private key to the keystore.

```
keytool -importkeystore -deststorepass Gauss@MppDB -destkeystore client.jks -srckeystore client.pkcs12 -srcstorepass Password -srcstoretype PKCS12 -alias 1
```

 NOTE

- In the preceding command, *Password* is an example. Replace it with the actual password.
- If information similar to the following is displayed and no error is reported, the import is successful. The target key file **client.jks** will be generated in **C:\dws_ssl_cert\sslcert**.

```
C:\dws_ssl_cert\sslcert>keytool -importkeystore -deststorepass Gauss@1000 -destkeystore client.jks -srckeystore client.pkcs12 -srcstorepass key123 -srcstoretype PKCS12 -alias 1  
Importing keystore client.pkcs12 to client.jks...
```

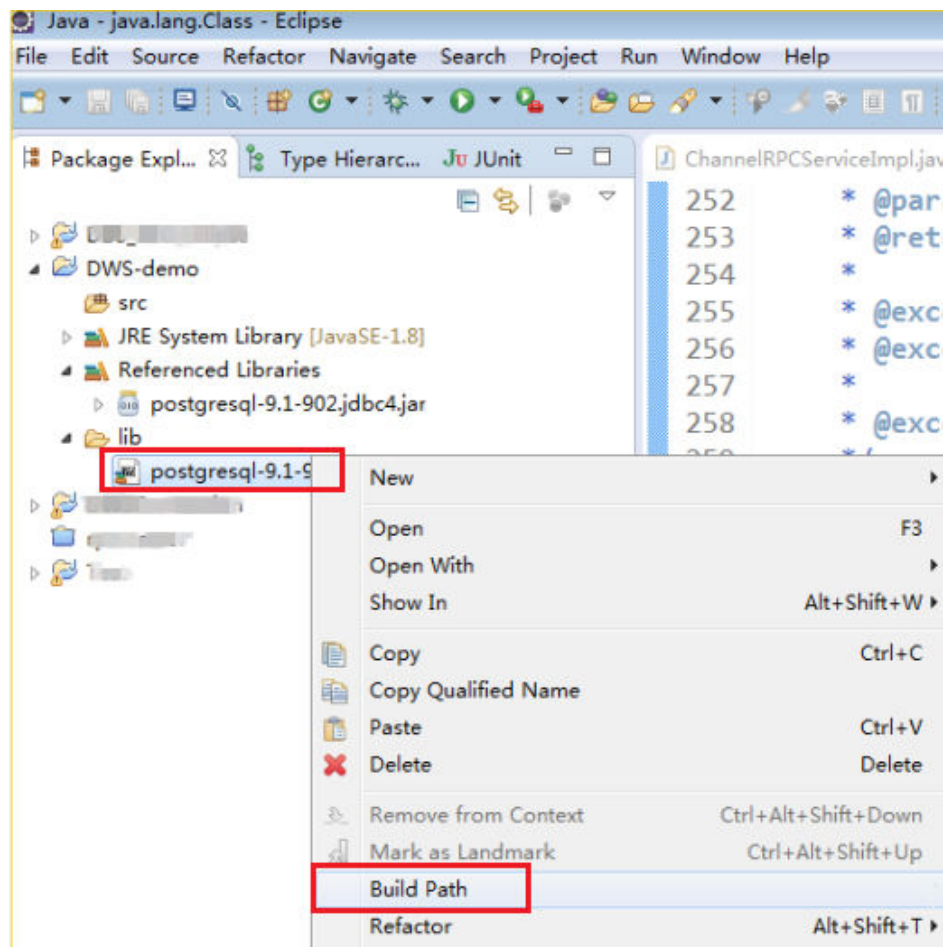


Step 4 Download the driver package **dws_8.1.x_jdbc_driver.zip** and decompress it. There will be two JDBC drive JAR packages, **gsjdbc4.jar** and **gsjdbc200.jar**. Use either of them as required.

Step 5 Add the JAR file to the application project so that applications can reference the JAR file.

Take the Eclipse project as an example. Store the JAR file to the project directory, for example, the **lib** directory in the project directory. In the Eclipse project, right-click the JAR file in the **lib** directory and choose **Build Path** to reference the JAR file.

Figure 4-7 Referencing a JAR file



Alternatively, you can use another method. In the Maven project, you can directly add the GaussDB(DWS) JDBC driver as a dependency item to the POM file. The following shows an example:

Step 6 Load the driver.

The following methods are available:

- Using a code: `Class.forName("org.postgresql.Driver");`
- Using a parameter during the JVM startup: `java -Djdbc.drivers=org.postgresql.Driver jdbctest`

NOTE

The JDBC driver package downloaded on GaussDB(DWS) contains both **gsjdbc4.jar** and **gsjdbc200.jar**.

- **gsjdbc4.jar**: The **gsjdbc4.jar** driver package is compatible with PostgreSQL. Its class names and class structures are the same as those of the PostgreSQL driver. Applications that run in PostgreSQL can be directly migrated to the current system.
- **gsjdbc200.jar**: If a JVM process needs to access PostgreSQL and GaussDB(DWS) at the same time, this driver package must be used. In this package, the main class name is `com.huawei.gauss200.jdbc.Driver` (replace `org.postgresql` with `com.huawei.gauss200.jdbc`). The URL prefix of the database connection is `jdbc:gaussdb`. Other parameters are the same as those of **gsjdbc4.jar**.
- The GaussDB(DWS) driver package downloaded from the Maven repository is the same as the **gsjdbc4** driver package.

Step 7 Call the **DriverManager.getConnection()** method of JDBC to connect to GaussDB(DWS) databases.

The JDBC API does not provide the connection retry capability. You need to implement the retry processing in the service code.

DriverManager.getConnection() methods:

- DriverManager.getConnection(String url);
- DriverManager.getConnection(String url, Properties info);
- DriverManager.getConnection(String url, String user, String password);

Table 4-13 Database connection parameters

Parameter	Description
url	<p>Specifies the database connection descriptor, which can be viewed on the management console. For details, see Obtaining the Cluster Connection Address.</p> <p>The URL format is as follows:</p> <ul style="list-style-type: none"> • jdbc:postgresql:database • jdbc:postgresql://host/database • jdbc:postgresql://host:port/database • jdbc:postgresql://host:port[,host:port][...]/database <p>NOTE</p> <ul style="list-style-type: none"> • If gsjdbc200.jar is used, change jdbc:postgresql to jdbc:gaussdb. <ul style="list-style-type: none"> – database indicates the name of the database to be connected. – host indicates the name or IP address of the database server. If an ELB is bound to the cluster, set host to the IP address of the ELB. – port indicates the port number of the database server. By default, the database running on port 8000 of the local host is connected. – Multiple IP addresses and ports can be configured. JDBC balances load by random access and failover, and will automatically ignore unreachable IP addresses. Separate multiple pairs of IP addresses and ports by commas (,). Example: jdbc:postgresql://10.10.0.13:8000,10.10.0.14:8000/database • If JDBC is used to connect to a cluster, only JDBC connection parameters can be configured in a cluster address. Variables cannot be added.

Parameter	Description
info	<p>Specifies database connection properties. Common properties include the following:</p> <ul style="list-style-type: none">● user: a string type. It indicates the database user who creates the connection task.● password: a string type. It indicates the password of the database user.● ssl: a boolean type. It indicates whether to use the SSL connection.● loggerLevel: string type. It indicates the volume of log data sent to the LogStream or LogWriter specified in the DriverManager. Currently, OFF, DEBUG, and TRACE are supported. DEBUG indicates that only logs of DEBUG or a higher level are printed, generating little log information. TRACE indicates that logs of the DEBUG and TRACE levels are displayed, generating detailed log information. The default value is OFF, indicating that no logs will be displayed.● prepareThreshold: integer type. It indicates the number of PreparedStatement executions required before requests are converted to prepared statements in servers. The default value is 5.● batchMode: boolean type. It indicates whether to connect the database in batch mode.● fetchsize: integer type. It indicates the default fetch size for statements in the created connection.● ApplicationName: string type. It indicates an application name. The default value is PostgreSQL JDBC Driver.● allowReadOnly: boolean type. It indicates whether to enable the read-only mode for connection. The default value is false. If the value is not changed to true, the execution of connection.setReadOnly does not take effect.● blobMode: string type. It is used to set the setBinaryStream method to assign values to different data types. The value on indicates that values are assigned to the BLOB data type and off indicates that values are assigned to the BYTEA data type. The default value is on.● currentSchema: string type. It specifies the schema used for connecting to the database.● defaultQueryMetaData: Boolean. It specifies whether to query SQL metadata by default. The default value is false. After this function is enabled, raw data operations are supported. However, it is incompatible with the create table as and select into operations in PrepareStatement.● connectionExtraInfo: boolean type. This parameter indicates whether the JDBC driver reports the driver deployment path and process owner to the database.

Parameter	Description
	<p>NOTE The value can be true or false. The default value is true. If connectionExtraInfo is set to true, the JDBC driver reports the driver deployment path and process owner to the database and displays the information in the connection_info parameter. In this case, you can query the information from PG_STAT_ACTIVITY or PGXC_STAT_ACTIVITY.</p> <ul style="list-style-type: none">• TCP_KEEPIDLE=30: The detection starts after the connection is idle for 30s. This parameter is valid only when tcpKeepAlive is set to true.• TCP_KEEPCOUNT=9: A total of nine detections are performed. This parameter is valid only when tcpKeepAlive is set to true.• TCP_KEEPINTERVAL=30: The detection interval is 30s. This parameter is valid only when tcpKeepAlive is set to true.• cnListRefreshSwitch (string): determines if JDBC automatically detects the CN list. Set to on for automatic detection, and off to disable it. The default value is off.• cnListRefreshDelay (integer): specifies the start time for scanning the CN liveness list. The default value is 1800000 (in milliseconds). This parameter is valid only when cnListRefreshSwitch is set to on.• cnListRefreshPeriod (integer): specifies the interval for scanning the CN list. The default value is 1800000 (in milliseconds). This parameter is valid only when cnListRefreshSwitch is set to on.• autoReconnect (boolean): enables or disables automatic reconnection for database connections. The value true enables it. The default value is false.• reConnectCount (integer): specifies the number of automatic reconnection attempts. The default value is 10. This parameter is valid only when autoReconnect is set to true. If reconnections exceed this number, the reconnection fails.• sslCrl: a string type that sets the path for the revoked certificate used by JDBC. The default value is null.
user	Specifies the database user.
password	Specifies the password of the database user.

The following describes the sample code used to encrypt the connection using the SSL certificate:

```
// The following code obtains the database SSL connection operation and encapsulates the operation as an API.  
public static Connection GetConnection(String username, String passwd) {  
    // Define the driver class.  
    String driver = "org.postgresql.Driver";  
    //Set keyStore.  
    System.setProperty("javax.net.ssl.trustStore", "mytruststore");  
    System.setProperty("javax.net.ssl.keyStore", "client.jks");  
    System.setProperty("javax.net.ssl.trustStorePassword", "password");  
}
```

```
System.setProperty("javax.net.ssl.keyStorePassword", "password");

Properties props = new Properties();
props.setProperty("user", username);
props.setProperty("password", passwd);
props.setProperty("ssl", "true");

String url = "jdbc:postgresql://" + "10.10.0.13" + ':' + "8000" + '/' + "gaussdb";
Connection conn = null;

try {
    // Load the driver.
    Class.forName(driver);
} catch (Exception e) {
    e.printStackTrace();
    return null;
}

try {
    // Create a connection.
    conn = DriverManager.getConnection(url, props);
    System.out.println("Connection succeed!");
} catch (SQLException throwables) {
    throwables.printStackTrace();
    return null;
}

return conn;
}
```

Step 8 Run SQL statements.

1. Run the following command to create a statement object:
Statement stmt = con.createStatement();
2. Run the following command to execute the statement object:
int rc = stmt.executeUpdate("CREATE TABLE tab1(id INTEGER, name VARCHAR(32));");
3. Run the following command to release the statement object:
stmt.close();

Step 9 Call close() to close the connection.

----End

Sample Code

This code sample illustrates how to develop applications based on the JDBC API provided by GaussDB(DWS).

NOTE

Before completing the following example, you need to create a stored procedure. For details, see [Tutorial: Development Using JDBC or ODBC](#).

```
create or replace procedure testproc
(
    psv_in1 in integer,
    psv_in2 in integer,
    psv_inout in out integer
)
as
begin
    psv_inout := psv_in1 + psv_in2 + psv_inout;
end;
/

//DBtest.java
//gsjdbc4.jar is used as an example. If gsjdbc200.jar is used, replace the driver class name org.postgresql
with com.huawei.gauss200.jdbc and replace the URL prefix jdbc:postgresql with jdbc:gaussdb.
//Demonstrate the main steps for JDBC development, including creating databases, creating tables, and
```

```
inserting data.

import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.PreparedStatement;
import java.sql.SQLException;
import java.sql.Statement;
import java.sql.CallableStatement;
import java.sql.Types;

public class DBTest {
//Create a database connection. Replace the following IP address and database with the actual database
connection address and database name.
public static Connection GetConnection(String username, String passwd) {
    String driver = "org.postgresql.Driver";
    String sourceURL = "jdbc:postgresql://10.10.0.13:8000/database";
    Connection conn = null;
    try {
        // Load the database driver.
        Class.forName(driver).newInstance();
    } catch (Exception e) {
        e.printStackTrace();
        return null;
    }

    try {
        //Create a database connection.
        conn = DriverManager.getConnection(sourceURL, username, passwd);
        System.out.println("Connection succeed!");
    } catch (Exception e) {
        e.printStackTrace();
        return null;
    }

    return conn;
};

//Run the common SQL statements to create table customer_t1.
public static void CreateTable(Connection conn) {
    Statement stmt = null;
    try {
        stmt = conn.createStatement();

        //Run the common SQL statements.
        int rc = stmt
            .executeUpdate("CREATE TABLE customer_t1(c_customer_sk INTEGER, c_customer_name
VARCHAR(32));");

        stmt.close();
    } catch (SQLException e) {
        if (stmt != null) {
            try {
                stmt.close();
            } catch (SQLException e1) {
                e1.printStackTrace();
            }
        }
        e.printStackTrace();
    }
}

//Run the prepared statements and insert data in batches.
public static void BatchInsertData(Connection conn) {
    PreparedStatement pst = null;

    try {
        //Generate the prepared statements.
        pst = conn.prepareStatement("INSERT INTO customer_t1 VALUES (?,?)");
        for (int i = 0; i < 3; i++) {
```

```
//Add parameters.
pst.setInt(1, i);
pst.setString(2, "data " + i);
pst.addBatch();
}
//Execute batch processing.
pst.executeBatch();
pst.close();
} catch (SQLException e) {
if (pst != null) {
try {
pst.close();
} catch (SQLException e1) {
e1.printStackTrace();
}
}
e.printStackTrace();
}
}

//Run the precompiled statement to update the data.
public static void ExecPreparedSQL(Connection conn) {
PreparedStatement pstmt = null;
try {
pstmt = conn
.prepareStatement("UPDATE customer_t1 SET c_customer_name = ? WHERE c_customer_sk = 1");
pstmt.setString(1, "new Data");
int rowcount = pstmt.executeUpdate();
pstmt.close();
} catch (SQLException e) {
if (pstmt != null) {
try {
pstmt.close();
} catch (SQLException e1) {
e1.printStackTrace();
}
}
e.printStackTrace();
}
}

//Execute the storage procedure.
public static void ExecCallableSQL(Connection conn) {
CallableStatement cstmt = null;
try {

cstmt=conn.prepareCall("{? = CALL TESTPROC(?,?,?)}");
cstmt.setInt(2, 50);
cstmt.setInt(1, 20);
cstmt.setInt(3, 90);
cstmt.registerOutParameter(4, Types.INTEGER); //Register a parameter of the out type. Its value is an
integer.
cstmt.execute();
int out = cstmt.getInt(4); //Obtain the out parameter.
System.out.println("The CallableStatment TESTPROC returns:"+out);
cstmt.close();
} catch (SQLException e) {
if (cstmt != null) {
try {
cstmt.close();
} catch (SQLException e1) {
e1.printStackTrace();
}
}
e.printStackTrace();
}
}
}
```

```
/**
 * Main program, which gradually invokes each static method.
 * @param args
 */
public static void main(String[] args) {
    //Create a database connection. Replace User and Password with the actual database user name and
    password.
    Connection conn = GetConnection("User", "Password");

    //Create a table.
    CreateTable(conn);

    //Insert data in batches.
    BatchInsertData(conn);

    //Run the precompiled statement to update the data.
    ExecPreparedSQL(conn);

    //Execute the storage procedure.
    ExecCallableSQL(conn);

    //Close the database connection.
    try {
        conn.close();
    } catch (SQLException e) {
        e.printStackTrace();
    }
}
}
```

4.5.5 Configuring JDBC to Connect to a Cluster (Load Balancing Mode)

Context

If you use JDBC to connect to only one CN in the cluster, this CN may be overloaded and other CN resources wasted. It also incurs single-node failure risks.

To avoid these problems, you can use JDBC to connect to multiple CNs. Two modes are available:

- Connection using ELB: An ELB distributes access traffic to multiple ECSs for traffic control based on forwarding policies. It improves the fault tolerance capability of application programs.
- Connection in multi-host mode: Use JDBC to configure multiple nodes, which is similar to ELB.

Method 1: Using ELB to Connect to a Cluster

Step 1 Obtain the Elastic Load Balance address. On the console, go to the details page of a cluster and obtain the ELB IP address.

Step 2 Configure the driver. For details, see [Downloading the JDBC or ODBC Driver](#).

```
<dependency>
  <groupId>com.huaweicloud.dws</groupId>
  <artifactId>huaweicloud-dws-jdbc</artifactId>
  <version>8.1.1.1</version>
</dependency>
```

Step 3 Obtain the database connection.

```
private static final String USER_NAME = "dbadmin";
private static final String PASSWORD = "password";
// jdbc:postgresql://ELB_IP:PORT/dbName"
private static final String URL = "jdbc:postgresql://100.95.153.169:8000/gaussdb";
private static Properties properties = new Properties();
static {
    properties.setProperty("user", USER_NAME);
    properties.setProperty("password", PASSWORD);
}
/**
 * Obtain the database connection.
 */
public static Connection getConnection() {
    Connection connection = null;
    try {
        connection = DriverManager.getConnection(URL, properties);
    } catch (SQLException e) {
        e.printStackTrace();
    }
    return connection;
}
```

----End

Method 2: Connecting to the Cluster in Multi-host Mode

Step 1 Obtain the EIP. Go to the details page of a cluster on the console and obtain the EIP.

Step 2 Configure the driver. For details, see [Downloading the JDBC or ODBC Driver](#).

```
<dependency>
  <groupId>com.huaweicloud.dws</groupId>
  <artifactId>huaweicloud-dws-jdbc</artifactId>
  <version>8.1.1.1</version>
</dependency>
```

Step 3 Obtain the database connection.

```
private static final String USER_NAME = "dbadmin";
private static final String PASSWORD = "password";
// jdbc:postgresql://host1:port1,host2:port2/dbName"
private static final String URL = "jdbc:postgresql://
100.95.146.194:8000,100.95.148.220:8000,100.93.0.221:8000/gaussdb?loadBalanceHosts=true";
private static Properties properties = new Properties();
static {
    properties.setProperty("user", USER_NAME);
    properties.setProperty("password", PASSWORD);
}
/**
 * Obtain the database connection.
 */
public static Connection getConnection() {
    Connection connection = null;
    try {
        connection = DriverManager.getConnection(URL, properties);
    } catch (SQLException e) {
        e.printStackTrace();
    }
    return connection;
}
```

----End

4.5.6 Configuring JDBC to Connect to a Cluster (IAM Authentication Mode)

Overview

GaussDB(DWS) allows you to access databases using IAM authentication. When you use the JDBC application program to connect to a cluster, set the IAM username, credential, and other information as you configure the JDBC URL. After doing this, when you try to access a database, the system will automatically generate a temporary credential and a connection will be set up.

NOTE

- Currently, only clusters 1.3.1 and later versions and their corresponding JDBC drivers can access the databases in IAM authentication mode. Download the JDBC driver. For details, see [Downloading the JDBC or ODBC Driver](#).
- Stream data warehouses do not support the connection to a cluster in IAM authentication mode.

IAM supports two types of user credential: password and Access Key ID/Secret Access Key (AK/SK). JDBC connection requires the latter.

The IAM account you use to access a database must be granted with the **DWS Database Access** permission. Only users with both the **DWS Administrator** and **DWS Database Access** permissions can connect to GaussDB(DWS) databases using the temporary database user credentials generated based on IAM users.

The **DWS Database Access** permission can only be granted to user groups. Ensure that your IAM account is in a user group with this permission.

On IAM, only users in the **admin** group have the permissions to manage users. This requires that your IAM account be in the **admin** user group. Otherwise, contact the IAM account administrator to grant your IAM account this permission.

The process of accessing a database is as follows:

1. [Granting an IAM Account the GaussDB\(DWS\) Database Access Permission](#)
2. [Creating an IAM User Credential](#)
3. [Configuring the JDBC Connection to Connect to a Cluster Using IAM Authentication](#)

Granting an IAM Account the GaussDB(DWS) Database Access Permission

Step 1 Log in to the Huawei Cloud management console. In the service list, choose **Management & Governance > Identity and Access Management** to enter the IAM management console.

Step 2 Modify the user group to which your IAM user belongs. Set a policy for, grant the **DWS Database Access** permission to, and add your IAM user to it.

Only users in the **admin** user group of IAM can perform this step. In IAM, only users in the **admin** user group can manage users, including creating user groups and users and setting user group rights.

You can also create an IAM user group, and set a policy for, grant the **DWS Administrator** and **DWS Database Access** permissions to, and add your IAM user to it.

----End

Creating an IAM User Credential

You can log in to the management console to create an AK/SK pair or use an existing one.

- Step 1** Log in to the management console.
- Step 2** Move the cursor to the username in the upper right corner and choose **My Credentials**.
- Step 3** Choose **Access Keys** to view the existing access keys. You can also click **Create Access Key** to create a new one.

The AK/SK pair is so important that you can download the private key file containing the AK/SK information only when you create the pair. On the management console, you can only view the AKs. If you have not downloaded the file, obtain it from your administrator or create an AK/SK pair again.

NOTE

Each user can create a maximum of two AK/SK pairs, which are valid permanently. To ensure account security, change your AK/SK pairs periodically and keep them safe.

----End

Configuring the JDBC Connection to Connect to a Cluster Using IAM Authentication

Configuring JDBC Connection Parameters

Table 4-14 Database connection parameters

Parameter	Description
url	<p>gsjdbc4.jar/gsjdbc200.jar database connection descriptor. The JDBC API does not provide the connection retry capability. You need to implement the retry processing in the service code. The URL example is as follows:</p> <pre>jdbc:dws:iam://dws-IAM-demo:eu-dublin-1/gaussdb? AccessKeyID=XXXXXXXXXXXXXXXXXXXX&SecretAccessKey=XXXXXXXXXXXXXXXXXXXX XXXXXXXXXXXX&DbUser=user_test&AutoCreate=true</pre> <p>JDBC URL parameters:</p> <ul style="list-style-type: none"> • jdbc:dws:iam is a prefix in the URL format. • dws-IAM-demo indicates the name of the cluster containing the database. • eu-dublin-1 indicates the region where the cluster resides. JDBC accesses the GaussDB(DWS) cluster in the corresponding region and delivers the IAM certificate to the cluster for IAM user authentication. The GaussDB(DWS) service address has been recorded in the JDBC configuration file. • gaussdb indicates the name of the database to which you want to connect. • AccessKeyID and SecretAccessKey are the access key ID and secret access key corresponding to the IAM user specified by DbUser. • Set DbUser to the IAM username. Note that the current version does not support hyphens (-) in the IAM username. <ul style="list-style-type: none"> – If the user specified by DbUser exists in the database, the temporary user credential has the same permissions as the existing user. – If the user specified by DbUser does not exist in the database and the value of AutoCreate is true, a new user named by the value of DbUser is automatically created. The created user is a common database user by default. • Parameter AutoCreate is optional. The default value is false. This parameter indicates whether to automatically create a database user named by the value of DbUser in the database. <ul style="list-style-type: none"> – The value true indicates that a user is automatically created. If the user already exists, the user will not be created again. – The value false indicates that a user is not created. If the username specified by DbUser does not exist in the database, an error is returned.

Parameter	Description
info	<p>Database connection properties. Common properties include the following:</p> <ul style="list-style-type: none"> • ssl: a boolean type. It indicates whether the SSL connection is used. • loglevel: an integer type. It sets the log amount recorded in DriverManager for LogStream or LogWriter. Currently, org.postgresql.Driver.DEBUG and org.postgresql.Driver.INFO logs are supported. If the value is 1, only org.postgresql.Driver.INFO (little information) is recorded. If the value is greater than or equal to 2, org.postgresql.Driver.DEBUG and org.postgresql.Driver.INFO logs are printed, and detailed log information is generated. Its default value is 0, which indicates that no logs are printed. • charSet: a string type. It indicates character sets used when data is sent from the database or the database receives data. • prepareThreshold: an integer type. It is used to determine the execution times of PreparedStatement before the information is converted into prepared statements on the server. The default value is 5.

Example

```
//The following uses gsjdbc4.jar as an example.
// The following code encapsulates the database connection obtaining operations into an API. You can
connect to the database by specifying the region where the cluster is located, cluster name, access key ID,
secret access key, and the corresponding IAM username.
public static Connection GetConnection(String clustername, String regionname, String AK, String SK,
String username) {
    String driver = "org.postgresql.Driver";
    // Driver class.
    String driver = "org.postgresql.Driver";
    // Database connection descriptor.
    String sourceURL = "jdbc:dws:iam://" + clustername + ":" + regionname + "/postgresgaussdb?" +
"AccessKeyID="
    + AK + "&SecretAccessKey=" + SK + "&DbUser=" + username + "&autoCreate=true";

    Connection conn = null;

    try {
        // Load the driver.
        Class.forName(driver);
    } catch (ClassNotFoundException e) {
        return null;
    }
    try {
        // Create a connection.
        conn = DriverManager.getConnection(sourceURL);
        System.out.println("Connection succeed!");
    } catch (SQLException e) {
        return null;
    }
    return conn;
}
```

4.5.7 Using ODBC to Connect to a Cluster

GaussDB(DWS) allows you to use an ODBC driver to connect to the database through an ECS on the Huawei Cloud platform or over the Internet.

For details about how to use the ODBC API, see the official document.

Prerequisites

- You have downloaded ODBC driver packages **dws_x.x.x_odbc_driver_for_xxx.zip** (for Linux) and **dws_odbc_driver_for_windows.zip** (for Windows). For details, see [Downloading the JDBC or ODBC Driver](#).

GaussDB(DWS) also supports open-source ODBC driver: PostgreSQL ODBC 09.01.0200 or later.

- You have downloaded the open-source unixODBC code file 2.3.0 from <https://sourceforge.net/projects/unixodbc/files/unixODBC/2.3.0/unixODBC-2.3.0.tar.gz/download>.
- You have downloaded the SSL certificate file. For details, see [Downloading an SSL Certificate](#).

Using an ODBC Driver to Connect to a Database (Linux)

Step 1 Upload the ODBC package and code file to the Linux environment and decompress them to the specified directory.

Step 2 Log in to the Linux environment as user **root**.

Step 3 Prepare **unixODBC**.

- Decompress the **unixODBC** code file.

```
tar -xvf unixODBC-2.3.0.tar.gz
```

- Compile the code file and install the driver.

```
cd unixODBC-2.3.0
./configure --enable-gui=no
make
make install
```

NOTE

- After the unixODBC is compiled and installed, the ***.so.2** library file will be in the installation directory. To create the ***.so.1** library file, change **LIB_VERSION** in the configure file to **1:0:0**.

```
LIB_VERSION="1:0:0"
```
- This driver dynamically loads the **libodbcinst.so.*** library files. If one of the library files is successfully loaded, the library file is loaded. The loading priority is **libodbcinst.so > libodbcinst.so.1 > libodbcinst.so.1.0.0 > libodbcinst.so.2 > libodbcinst.so.2.0.0**.

For example, a directory can be dynamically linked to **libodbcinst.so.1**, **libodbcinst.so.1.0.0**, and **libodbcinst.so.2**. The driver file loads **libodbcinst.so** first. If **libodbcinst.so** cannot be found in the current environment, the driver file searches for **libodbcinst.so.1**, which has a lower priority. After **libodbcinst.so.1** is loaded, the loading is complete.

Step 4 Replace the driver file. (This document uses the **dws_8.1.x_odbc_driver_for_x86_redhat.zip** package of Red Hat as an example.)

1. Decompress the **dws_8.1.x_odbc_driver_for_x86_redhat.zip** package.
`unzip dws_8.1.x_odbc_driver_for_x86_redhat.zip`
2. Copy all files in the **lib** directory to **/usr/local/lib**. If there are files with the same name, overwrite them.
3. Copy **psqlodbcw.la** and **psqlodbcw.so** in the **odbc/lib** directory to **/usr/local/lib**.

Step 5 Run the following command to modify the configuration of the driver file:

```
vi /usr/local/etc/odbcinst.ini
```

Copy the following content to the file:

```
[DWS]  
Driver64=/usr/local/lib/psqlodbcw.so
```

The parameters are as follows:

- **[DWS]**: indicates the driver name. You can customize the name.
- **Driver64** or **Driver**: indicates the path where the dynamic library of the driver resides. For a 64-bit operating system, search for **Driver64** first. If **Driver64** is not configured, search for **Driver**.

Step 6 Run the following command to modify the data source file:

```
vi /usr/local/etc/odbc.ini
```

Copy the following content to the configuration file, save the modification, and exit.

```
[DWSODBC]  
Driver=DWS  
Servername=10.10.0.13  
Database=gaussdb  
Username=dbadmin  
Password=password  
Port=8000  
Sslmode=allow
```

Parameter	Description	Example Value
[DSN]	Data source name.	[DWSODBC]
Driver	Driver name, corresponding to DriverName in odbcinst.ini .	Driver=DWS
Servername	IP address of the server. When the cluster is bound to an ELB, set this parameter to the IP address of the ELB.	Servername=10.10.0.13
Database	Name of the database to be connected to.	Database=gaussdb
Username	Database username.	Username=dbadmin
Password	Database user password.	Password= <i>password</i>
Port	Port number of the server.	Port=8000

Parameter	Description	Example Value
Sslmode	<p>SSL certification mode. This parameter is enabled for the cluster by default.</p> <p>Values and meanings:</p> <ul style="list-style-type: none">• disable: only tries to establish a non-SSL connection.• allow: tries establishing a non-SSL connection first, and then an SSL connection if the attempt fails.• prefer: tries establishing an SSL connection first, and then a non-SSL connection if the attempt fails.• require: only tries establishing an SSL connection. If there is a CA file, perform the verification according to the scenario in which the parameter is set to verify-ca.• verify-ca: tries establishing an SSL connection and checks whether the server certificate is issued by a trusted CA.• verify-full: not supported by GaussDB(DWS) <p>NOTE The SSL mode delivers higher security than the common mode. By default, the SSL function is enabled in a cluster to allow SSL or non-SSL connections from the client. You are advised to use the SSL mode when using ODBC to connect to a GaussDB (DWS) cluster.</p>	Sslmode=allow

 **NOTE**

- You can view the values of **Servername** and **Port** on the GaussDB(DWS) management console. Log in to the GaussDB(DWS) management console and click **Client Connections**. In the **Data Warehouse Connection String** area, select the target cluster and obtain **Private Network Address** or **Public Network Address**. For details, see [Obtaining the Cluster Connection Address](#).

Step 7 Configure environment variables.

```
vi ~/.bashrc
```

Add the following information to the configuration file:

```
export LD_LIBRARY_PATH=/usr/local/lib/:$LD_LIBRARY_PATH
export ODBCYSINI=/usr/local/etc
export ODBCINI=/usr/local/etc/odbc.ini
```

Step 8 Import environment variables.

```
source ~/.bashrc
```

Step 9 Run the following commands to connect to the database:

```
/usr/local/bin/isql -v DWSODBC
```

If the following information is displayed, the connection is successful:

```
+-----+
| Connected!          |
|                    |
| sql-statement      |
| help [tablename]   |
| quit               |
|                    |
+-----+
SQL>
```

----End

Using an ODBC Driver to Connect to a Database (Windows)

Step 1 Decompress ODBC driver package **dws_odbc_driver_for_windows.zip** (for Windows) and install **psqlodbc.msi**.

Step 2 Decompress the SSL certificate package to obtain the certificate file.

You have the option to deploy the certificate either automatically or manually, depending on your requirements.

- Automatic deployment:

Double-click the **sslcert_env.bat** file to trigger automatic deployment of the certificate to a default location.

 **NOTE**

- The **sslcert_env.bat** file ensures the purity of the certificate environment. When the **%APPDATA%\postgresql** directory exists, a message will be prompted asking you whether you want to remove related directories. If you want to remove related directories, back up files in the directory.
- Manual deployment:
 - Create a new folder named **postgresql** in the **%APPDATA%** directory.
 - Copy files **client.crt**, **client.key**, **client.key.cipher**, and **client.key.rand** to the **%APPDATA%\postgresql** directory and change **client** in the file name to **postgres**. For example, change the name of **client.key** to **postgres.key**.
 - Copy **cacert.pem** to **%APPDATA%\postgresql** and change the name of **cacert.pem** to **root.crt**.

Step 3 Open Driver Manager.

GaussDB(DWS) provides 32-bit and 64-bit ODBC drivers. Choose the version suitable for your system when configuring the data source. (Assume the Windows system drive is drive C. If another disk drive is used, modify the path accordingly.)

- If you want to develop 32-bit programs in the 64-bit OS and have installed the 32-bit driver, open the 32-bit Driver Manager at **C:\Windows\SysWOW64\odbcad32.exe**.

Do not choose **Control Panel > System and Security > Administrative Tools > Data Sources (ODBC)** directly.

 **NOTE**

WOW64 is the acronym for Windows 32-bit on Windows 64-bit. **C:\Windows\SysWOW64** stores the 32-bit environment on a 64-bit system.

- If you want to develop 64-bit programs in the 64-bit OS and have installed the 64-bit driver, open the 64-bit Driver Manager at **C:\Windows\System32\odbcad32.exe**.

Do not choose **Control Panel > System and Security > Administrative Tools > Data Sources (ODBC)** directly.

 **NOTE**

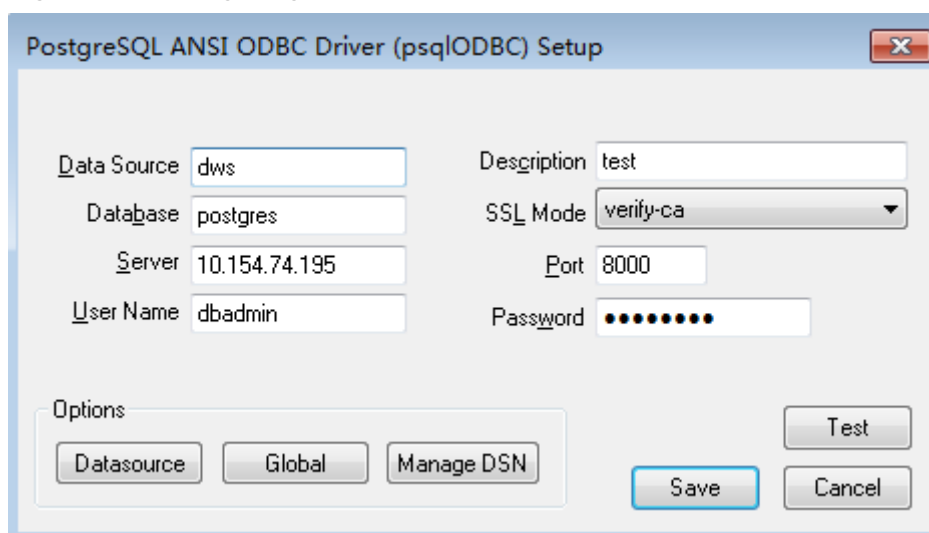
C:\Windows\System32 stores the environment consistent with the current OS. For technical details, see Windows technical documents.

- In a 32-bit OS, open **C:\Windows\System32\odbcad32.exe**.
Alternatively, click **Computer**, and choose **Control Panel**. Click **Administrative Tools** and click **Data Sources (ODBC)**.

Step 4 Configure a data source to be connected to.

1. On the **User DSN** tab, click **Add** and choose **PostgreSQL Unicode** for setup.

Figure 4-8 Configuring a data source to be connected to



You can view the values of **Server** and **Port** on the GaussDB(DWS) management console. Log in to the GaussDB(DWS) management console and click **Client Connections**. In the **Data Warehouse Connection String** area, select the target cluster and obtain **Private Network Address** or **Public**

Network Address. For details, see [Obtaining the Cluster Connection Address](#).

- Click **Test** to verify that the connection is correct. If **Connection successful** is displayed, the connection is correct.

Step 5 Compile an ODBC sample program to connect to the data source.

The ODBC API does not provide the database connection retry capability. You need to implement the connection retry processing in the service code.

The sample code is as follows:

```
// This example shows how to obtain GaussDB(DWS) data through the ODBC driver.
// DBtest.c (compile with: libodbc.so)
#include <stdlib.h>
#include <stdio.h>
#include <sqlext.h>
#ifdef WIN32
#include <windows.h>
#endif
SQLHENV    V_OD_Env;    // Handle ODBC environment
SQLHSTMT   V_OD_hstmt; // Handle statement
SQLHDBC    V_OD_hdbc;  // Handle connection
char       typename[100];
SQLINTEGER value = 100;
SQLINTEGER V_OD_erg,V_OD_buffer,V_OD_err,V_OD_id;
int main(int argc,char *argv[])
{
    // 1. Apply for an environment handle.
    V_OD_erg = SQLAllocHandle(SQL_HANDLE_ENV,SQL_NULL_HANDLE,&V_OD_Env);
    if ((V_OD_erg != SQL_SUCCESS) && (V_OD_erg != SQL_SUCCESS_WITH_INFO))
    {
        printf("Error AllocHandle\n");
        exit(0);
    }
    // 2. Set environment attributes (version information).
    SQLSetEnvAttr(V_OD_Env, SQL_ATTR_ODBC_VERSION, (void*)SQL_OV_ODBC3, 0);
    // 3. Apply for a connection handle.
    V_OD_erg = SQLAllocHandle(SQL_HANDLE_DBC, V_OD_Env, &V_OD_hdbc);
    if ((V_OD_erg != SQL_SUCCESS) && (V_OD_erg != SQL_SUCCESS_WITH_INFO))
    {
        SQLFreeHandle(SQL_HANDLE_ENV, V_OD_Env);
        exit(0);
    }
    // 4. Set connection attributes.
    SQLSetConnectAttr(V_OD_hdbc, SQL_ATTR_AUTOCOMMIT, SQL_AUTOCOMMIT_ON, 0);
    // 5. Connect to a data source. You do not need to enter the username and password if you have
    // configured them in the odbc.ini file. If you have not configured them, specify the name and password of
    // the user who wants to connect to the database in the SQLConnect function.
    V_OD_erg = SQLConnect(V_OD_hdbc, (SQLCHAR*) "gaussdb", SQL_NTS,
        (SQLCHAR*) "", SQL_NTS, (SQLCHAR*) "", SQL_NTS);
    if ((V_OD_erg != SQL_SUCCESS) && (V_OD_erg != SQL_SUCCESS_WITH_INFO))
    {
        printf("Error SQLConnect %d\n",V_OD_erg);
        SQLFreeHandle(SQL_HANDLE_ENV, V_OD_Env);
        exit(0);
    }
    printf("Connected !\n");
    // 6. Set statement attributes.
    SQLSetStmtAttr(V_OD_hstmt,SQL_ATTR_QUERY_TIMEOUT,(SQLPOINTER *)3,0);
    // 7. Apply for a statement handle.
    SQLAllocHandle(SQL_HANDLE_STMT, V_OD_hdbc, &V_OD_hstmt);
    // 8. Executes an SQL statement directly.
    SQLExecDirect(V_OD_hstmt,"drop table IF EXISTS testtable",SQL_NTS);
    SQLExecDirect(V_OD_hstmt,"create table testtable(id int)",SQL_NTS);
    SQLExecDirect(V_OD_hstmt,"insert into testtable values(25)",SQL_NTS);
    // 9. Prepare for execution.
    SQLPrepare(V_OD_hstmt,"insert into testtable values(?)",SQL_NTS);
```



```
// 10. Bind parameters.
SQLBindParameter(V_OD_hstmt,1,SQL_PARAM_INPUT,SQL_C_SLONG,SQL_INTEGER,0,0,
    &value,0,NULL);
// 11. Execute the ready statement.
SQLExecute(V_OD_hstmt);
SQLExecDirect(V_OD_hstmt,"select id from testtable",SQL_NTS);
// 12. Obtain the attributes of a certain column in the result set.
SQLColAttribute(V_OD_hstmt,1,SQL_DESC_TYPE,typename,100,NULL,NULL);
printf("SQLColAttribute %s\n",typename);
// 13. Bind the result set.
SQLBindCol(V_OD_hstmt,1,SQL_C_SLONG, (SQLPOINTER)&V_OD_buffer,150,
    (SQLLEN *)&V_OD_err);
// 14. Collect data using SQLFetch.
V_OD_erg=SQLFetch(V_OD_hstmt);
// 15. Obtain and return data using SQLGetData.
while(V_OD_erg != SQL_NO_DATA)
{
    SQLGetData(V_OD_hstmt,1,SQL_C_SLONG,(SQLPOINTER)&V_OD_id,0,NULL);
    printf("SQLGetData ----ID = %d\n",V_OD_id);
    V_OD_erg=SQLFetch(V_OD_hstmt);
};
printf("Done !\n");
// 16. Disconnect from the data source and release handles.
SQLFreeHandle(SQL_HANDLE_STMT,V_OD_hstmt);
SQLDisconnect(V_OD_hdbc);
SQLFreeHandle(SQL_HANDLE_DBC,V_OD_hdbc);
SQLFreeHandle(SQL_HANDLE_ENV, V_OD_Env);
return(0);
}
```

----End

4.6 Using the Third-Party Function Library psycopg2 of Python to Connect to a Cluster

After creating a data warehouse cluster and using the third-party function library psycopg2 to connect to the cluster, you can use Python to access GaussDB(DWS) and perform various operations on data tables.

Preparations Before Connecting to a Cluster

- An EIP has been bound to the data warehouse cluster.
- You have obtained the administrator username and password for logging in to the database in the data warehouse cluster.

MD5 algorithms may be vulnerable to collision attacks and cannot be used for password verification. Currently, GaussDB(DWS) uses the default security design. By default, MD5 password verification is disabled, and this may cause failures of connections from open source clients. You are advised to set **password_encryption_type** to **1**. For details, see [Modifying Database Parameters](#).

 NOTE

- For security purposes, GaussDB(DWS) no longer uses MD5 to store password digests by default. As a result, the open-source drivers and clients may fail to connect to the database. To use the MD5 algorithm used in an open-source protocol, you must modify your password policy and create a new user, or change the password of an existing user.
- The database stores the hash digest of passwords instead of password text. During password verification, the system compares the hash digest with the password digest sent from the client (salt operations are involved). If you change your cryptographic algorithm policy, the database cannot generate a new hash digest for your existing password. For connectivity purposes, you must manually change your password or create a new user. The new password will be encrypted using the hash algorithm and stored for authentication in the next connection.
- You have obtained the public network address, including the IP address and port number in the data warehouse cluster. For details, see [Obtaining the Cluster Connection Address](#).
- You have installed the third-party function library psycopg2. Download address: <https://pypi.org/project/psycopg2/>. For details about installation and deployment, see <https://www.psycopg.org/install/>.

 NOTE

- In CentOS and Red Hat OS, run the following **yum** command:
yum install python-psycopg2
- psycopg2 depends on the libpq dynamic library of PostgreSQL (32-bit or 64-bit version, whichever matches the psycopg2 bit version). In Linux, you can run the **yum** command and do not need to install the library. Before using psycopg2 in Windows, you need to install libpq in either of the following ways:
 - Install PostgreSQL and configure the libpq, ssl, and crypto dynamic libraries in the environment variable **PATH**.
 - Install psqldb and use the libpq, ssl, and crypto dynamic libraries carried by the PostgreSQL ODBC driver.

Constraints

psycopg2 is a PostgreSQL-based client interface, and its functions are not fully supported by GaussDB(DWS). For details, see [Table 4-15](#).

 NOTE

The following APIs are supported based on Python 3.8.5 and psycopg 2.9.1.

Table 4-15 psycopg2 APIs supported by GaussDB(DWS)

Class Name	Usage	Function/Member Variable	Yes	Remarks
connections	basic	<i>cursor(name=None, cursor_factory=None, scrollable=None, withhold=False)</i>	Y	-
		commit()	Y	-
		rollback()	Y	-

Class Name	Usage	Function/Member Variable	Yes	Remarks
		close()	Y	-
	Two-phase commit support methods	xid(<i>format_id, gtrid, bqual</i>)	Y	-
		tpc_begin(<i>xid</i>)	Y	-
		tpc_prepare()	N	The kernel does not support explicit PREPARE TRANSACTION .
		tpc_commit(<i>[xid]</i>)	Y	-
		tpc_rollback(<i>[xid]</i>)	Y	-
		tpc_recover()	Y	-
		closed	Y	-
		cancel()	Y	-
		reset()	N	DISCARD ALL is not supported.
		dsn	Y	-
	Transaction control methods and attributes.	set_session(<i>isolation_level=None, readonly=None, deferrable=None, autocommit=None</i>)	Y	The database does not support the setting of default_transaction_read_only in a session.
		autocommit	Y	-
		isolation_level	Y	-
		readonly	N	The database does not support the setting of default_transaction_read_only in a session.
		deferrable	Y	-

Class Name	Usage	Function/Member Variable	Yes	Remarks
		set_isolation_level(<i>level</i>)	Y	-
		encoding	Y	-
		set_client_encoding(enc)	Y	-
		notices	N	The database does not support listen/notify .
		notifies	Y	-
		cursor_factory	Y	-
		info	Y	-
		status	Y	-
		lobjct	N	The database does not support operations related to large objects.
	Methods related to asynchronous support	poll()	Y	-
		fileno()	Y	-
		isexecuting()	Y	-
	Interoperation with other C API modules	pgconn_ptr	Y	-
		get_native_connection()	Y	-
	informative methods of the native connection	get_transaction_status()	Y	-
protocol_version		Y	-	
server_version		Y	-	
get_backend_pid()		Y	The obtained PID is not the background PID, but the ID of the logical connection.	
get_parameter_status(parameter)		Y	-	

Class Name	Usage	Function/Member Variable	Yes	Remarks
		get_dsn_parameters()	Y	-
cursor	basic	description	Y	-
		close()	Y	-
		closed	Y	-
		connection	Y	-
		name	Y	-
		scrollable	N	The database does not support SCROLL CURSOR .
		withhold	N	The withhold cursor needs to be closed before the commit operation.
	Command s execution methods	execute(<i>query</i> , <i>vars=None</i>)	Y	-
		executemany(<i>query</i> , <i>vars_list</i>)	Y	-
		callproc(<i>procname</i> [, <i>parameters</i>])	Y	-
		mogrify(<i>operation</i> [, <i>parameters</i>])	Y	-
		setinputsizes(<i>sizes</i>)	Y	-
		fetchone()	Y	-
		fetchmany([<i>size=cursor.arraysize</i>])	Y	-
		fetchall()	Y	-
		scroll(<i>value</i> [, <i>mode='relative'</i>])	N	The database does not support SCROLL CURSOR .
		arraysize	Y	-
		itersize	Y	-
		rowcount	Y	-

Class Name	Usage	Function/Member Variable	Yes	Remarks
		rownumber	Y	-
		lastrowid	Y	-
		query	Y	-
		statusmessage	Y	-
		cast(<i>oid</i> , <i>s</i>)	Y	-
		tzinfo_factory	Y	-
		nextset()	Y	-
		setoutputsize(<i>size</i> [, <i>column</i>])	Y	-
	COPY-related methods	copy_from(<i>file</i> , <i>table</i> , <i>sep</i> =' t', <i>null</i> =' N', <i>size</i> =8192, <i>columns</i> =None)	Y	-
		copy_to(<i>file</i> , <i>table</i> , <i>sep</i> =' t', <i>null</i> =' N', <i>columns</i> =None)	Y	-
		copy_expert(<i>sql</i> , <i>file</i> , <i>size</i> =8192)	Y	-
	Interoper- ation with other C API modules	pgresult_ptr	Y	-

Using the Third-Party Function Library psycopg2 to Connect to a Cluster (Linux)

Step 1 Log in to the Linux environment as user **root**.

Step 2 Run the following command to create the **python_dws.py** file:

```
vi python_dws.py
```

Copy and paste the following content to the **python_dws.py** file:

```
#!/usr/bin/python
# -*- coding: UTF-8 -*-

from __future__ import print_function

import psycopg2

def create_table(connection):
    print("Begin to create table")
    try:
        cursor = connection.cursor()
        cursor.execute("drop table if exists test;"
            "create table test(id int, name text);")
```

```
        connection.commit()
    except psycopg2.ProgrammingError as e:
        print(e)
    else:
        print("Table created successfully")
        cursor.close()

def insert_data(connection):
    print("Begin to insert data")
    try:
        cursor = connection.cursor()
        cursor.execute("insert into test values(1,'number1');")
        cursor.execute("insert into test values(2,'number2');")
        cursor.execute("insert into test values(3,'number3');")
        connection.commit()
    except psycopg2.ProgrammingError as e:
        print(e)
    else:
        print("Insert data successfully")
        cursor.close()

def update_data(connection):
    print("Begin to update data")
    try:
        cursor = connection.cursor()
        cursor.execute("update test set name = 'numberupdated' where id=1;")
        connection.commit()
        print("Total number of rows updated :", cursor.rowcount)
        cursor.execute("select * from test order by 1;")
        rows = cursor.fetchall()
        for row in rows:
            print("id = ", row[0])
            print("name = ", row[1], "\n")
    except psycopg2.ProgrammingError as e:
        print(e)
    else:
        print("After Update, Operation done successfully")

def delete_data(connection):
    print("Begin to delete data")
    try:
        cursor = connection.cursor()
        cursor.execute("delete from test where id=3;")
        connection.commit()
        print("Total number of rows deleted :", cursor.rowcount)
        cursor.execute("select * from test order by 1;")
        rows = cursor.fetchall()
        for row in rows:
            print("id = ", row[0])
            print("name = ", row[1], "\n")
    except psycopg2.ProgrammingError as e:
        print(e)
    else:
        print("After Delete, Operation done successfully")

def select_data(connection):
    print("Begin to select data")
    try:
        cursor = connection.cursor()
        cursor.execute("select * from test order by 1;")
        rows = cursor.fetchall()
        for row in rows:
            print("id = ", row[0])
            print("name = ", row[1], "\n")
    except psycopg2.ProgrammingError as e:
```

```
print(e)
print("select failed")
else:
    print("Operation done successfully")
    cursor.close()

if __name__ == '__main__':
    try:
        conn = psycopg2.connect(host='10.154.70.231',
                                port='8000',
                                database='gaussdb', # Database to be connected
                                user='dbadmin',
                                password='password') # Database user password
    except psycopg2.DatabaseError as ex:
        print(ex)
        print("Connect database failed")
    else:
        print("Opened database successfully")
        create_table(conn)
        insert_data(conn)
        select_data(conn)
        update_data(conn)
        delete_data(conn)
        conn.close()
```

- Step 3** Change the public network address, cluster port number, database name, database username, and database password in the `python_dws.py` file based on the actual cluster information.

The `psycopg2` API does not provide the connection retry capability. You need to implement the retry processing in the service code.

```
conn = psycopg2.connect(host='10.154.70.231',
                        port='8000',
                        database='gaussdb', # Database to be connected
                        user='dbadmin',
                        password='password') # Database user password
```

- Step 4** Run the following command to connect to the cluster using the third-party function library `psycopg`:

```
python python_dws.py
```

----End

Using the Third-Party Function Library `psycopg2` to Connect to a Cluster (Windows)

- Step 1** In the Windows operating system, click the **Start** button, enter `cmd` in the search box, and click `cmd.exe` in the result list to open the command-line interface (CLI).

- Step 2** In the CLI, run the following command to create the `python_dws.py` file:

```
type nul> python_dws.py
```

Copy and paste the following content to the `python_dws.py` file:

```
#!/usr/bin/python
# -*- coding:UTF-8 -*-

from __future__ import print_function

import psycopg2

def create_table(connection):
```



```
print("Begin to create table")
try:
    cursor = connection.cursor()
    cursor.execute("drop table if exists test;"
                  "create table test(id int, name text);")
    connection.commit()
except psycopg2.ProgrammingError as e:
    print(e)
else:
    print("Table created successfully")
    cursor.close()

def insert_data(connection):
    print("Begin to insert data")
    try:
        cursor = connection.cursor()
        cursor.execute("insert into test values(1,'number1');")
        cursor.execute("insert into test values(2,'number2');")
        cursor.execute("insert into test values(3,'number3');")
        connection.commit()
    except psycopg2.ProgrammingError as e:
        print(e)
    else:
        print("Insert data successfully")
        cursor.close()

def update_data(connection):
    print("Begin to update data")
    try:
        cursor = connection.cursor()
        cursor.execute("update test set name = 'numberupdated' where id=1;")
        connection.commit()
        print("Total number of rows updated :", cursor.rowcount)
        cursor.execute("select * from test order by 1;")
        rows = cursor.fetchall()
        for row in rows:
            print("id = ", row[0])
            print("name = ", row[1], "\n")
    except psycopg2.ProgrammingError as e:
        print(e)
    else:
        print("After Update, Operation done successfully")

def delete_data(connection):
    print("Begin to delete data")
    try:
        cursor = connection.cursor()
        cursor.execute("delete from test where id=3;")
        connection.commit()
        print("Total number of rows deleted :", cursor.rowcount)
        cursor.execute("select * from test order by 1;")
        rows = cursor.fetchall()
        for row in rows:
            print("id = ", row[0])
            print("name = ", row[1], "\n")
    except psycopg2.ProgrammingError as e:
        print(e)
    else:
        print("After Delete, Operation done successfully")

def select_data(connection):
    print("Begin to select data")
    try:
        cursor = connection.cursor()
        cursor.execute("select * from test order by 1;")
```

```
rows = cursor.fetchall()
for row in rows:
    print("id = ", row[0])
    print("name = ", row[1], "\n")
except psycopg2.ProgrammingError as e:
    print(e)
    print("select failed")
else:
    print("Operation done successfully")
    cursor.close()

if __name__ == '__main__':
    try:
        conn = psycopg2.connect(host='10.154.70.231',
                                port='8000',
                                database='postgresgaussdb', # Database to be connected
                                user='dbadmin',
                                password='password') # Database user password
    except psycopg2.DatabaseError as ex:
        print(ex)
        print("Connect database failed")
    else:
        print("Opened database successfully")
        create_table(conn)
        insert_data(conn)
        select_data(conn)
        update_data(conn)
        delete_data(conn)
        conn.close()
```

Step 3 Change the public network address, cluster port number, database name, database username, and database password in the `python_dws.py` file based on the actual cluster information.

```
conn = psycopg2.connect(host='10.154.70.231',
                        port='8000',
                        database='gaussdb', # Database to be connected
                        user='dbadmin',
                        password='password') # Database user password
```

Step 4 On the CLI, run the following command to use `psycopg2` to connect to the cluster:

```
python python_dws.py
```

----End

Why CN Retry Is Not Supported When `psycopg2` Is Connected to a Cluster?

With the CN retry feature, GaussDB(DWS) retries a statement that failed to be executed and identifies the failure type. However, in a session connected using `psycopg2`, a failed SQL statement will report an error and stop to be executed. In a primary/standby switchover, if a failed SQL statement is not retried, the following error will be reported. If the switchover is complete during an automatic retry, the correct result will be returned.

```
psycopg2.errors.ConnectionFailure: pooler: failed to create 1 connections, Error Message: remote node dn_6003_6004, detail: could not connect to server: Operation now in progress
```

Error causes:

1. `psycopg2` sends the **BEGIN** statement to start a transaction before sending an SQL statement.
2. CN retry does not support statements in transaction blocks.

Solution:

- In synchronous connection mode, end the transaction started by the driver.
cursor = conn.cursor()
End the transaction started by the driver.
cursor.execute("end; select * from test order by 1;")
rows = cursor.fetchall()
- Start a transaction in an asynchronous connection. For details, visit the PyScopg official website at: <https://www.pyscopg.org/docs/advanced.html?highlight=async>

```
#!/usr/bin/env python3
# -*- encoding=utf-8 -*-

import psycopg2
import select

# Wait function provided by psycopg2 in asynchronous connection mode
#For details, see https://www.pyscopg.org/docs/advanced.html?highlight=async.
def wait(conn):
    while True:
        state = conn.poll()
        if state == psycopg2.extensions.POLL_OK:
            break
        elif state == psycopg2.extensions.POLL_WRITE:
            select.select([], [conn.fileno()], [])
        elif state == psycopg2.extensions.POLL_READ:
            select.select([conn.fileno()], [], [])
        else:
            raise psycopg2.OperationalError("poll() returned %s" % state)

def psycopg2_cnretry_sync():
    # Create a connection.
    conn = psycopg2.connect(host='10.154.70.231',
                           port='8000',
                           database='gaussdb', # Database to be connected
                           user='dbadmin',
                           password='password', # Database user password
                           async=1) # Use the asynchronous connection mode.

    wait(conn)

    # Execute a query.
    cursor = conn.cursor()
    cursor.execute("select * from test order by 1;")
    wait(conn)
    rows = cursor.fetchall()
    for row in rows:
        print(row[0], row[1])

    # Close the connection.
    conn.close()

if __name__ == '__main__':
    psycopg2_cnretry_async()
```

4.7 Using the Python Library PyGreSQL to Connect to a Cluster

After creating a data warehouse cluster and using the third-party function library PyGreSQL to connect to the cluster, you can use Python to access GaussDB(DWS) and perform various operations on data tables.

Preparations Before Connecting to a Cluster

- An EIP has been bound to the data warehouse cluster.

- You have obtained the administrator username and password for logging in to the database in the data warehouse cluster.

MD5 algorithms may be vulnerable to collision attacks and cannot be used for password verification. Currently, GaussDB(DWS) uses the default security design. By default, MD5 password verification is disabled, and this may cause failures of connections from open source clients. You are advised to set **password_encryption_type** to **1**. For details, see [Modifying Database Parameters](#).

NOTE

- For security purposes, GaussDB(DWS) no longer uses MD5 to store password digests by default. As a result, the open-source drivers and clients may fail to connect to the database. To use the MD5 algorithm used in an open-source protocol, you must modify your password policy and create a new user, or change the password of an existing user.
- The database stores the hash digest of passwords instead of password text. During password verification, the system compares the hash digest with the password digest sent from the client (salt operations are involved). If you change your cryptographic algorithm policy, the database cannot generate a new hash digest for your existing password. For connectivity purposes, you must manually change your password or create a new user. The new password will be encrypted using the hash algorithm and stored for authentication in the next connection.
- You have obtained the public network address, including the IP address and port number in the data warehouse cluster. For details, see [Obtaining the Cluster Connection Address](#).
- You have installed the third-party function library PyGreSQL.
Download address: <http://www.pygresql.org/download/index.html>
- For details about the installation and deployment operations, see <http://www.pygresql.org/contents/install.html>

NOTE

- In CentOS and Red Hat OS, run the following **yum** command:

```
yum install PyGreSQL
```
- PyGreSQL depends on the libpq dynamic library of PostgreSQL (32-bit or 64-bit version, whichever matches the PyGreSQL bit version). In Linux, you can run the **yum** command and do not need to install the library. Before using PyGreSQL in Windows, you need to install libpq in either of the following ways:
 - Install PostgreSQL and configure the libpq, ssl, and crypto dynamic libraries in the environment variable **PATH**.
 - Install **psqlodbc** and use the **libpq**, **ssl**, and **crypto** dynamic libraries carried by the PostgreSQL ODBC driver.

Constraints

PyGreSQL is a PostgreSQL-based client interface, and its functions are not fully supported by GaussDB(DWS). For details, see [Table 4-16](#).

NOTE

The following APIs are supported based on Python 3.8.5 and PyGreSQL 5.2.4.

Table 4-16 PyGreSQL APIs supported by DWS

PyGreSQL	Yes	Remarks	
Module functions and constants	connect – Open a PostgreSQL connection	Y	-
	get_pqlib_version – get the version of libpq	Y	-
	get/set_defhost – default server host [DV]	Y	-
	get/set_defport – default server port [DV]	Y	-
	get/set_defopt – default connection options [DV]	Y	-
	get/set_defbase – default database name [DV]	Y	-
	get/set_defuser – default database user [DV]	Y	-
	get/set_defpasswd – default database password [DV]	Y	-
	escape_string – escape a string for use within SQL	Y	-
	escape_bytea – escape binary data for use within SQL	Y	-
	unescape_bytea – unescape data that has been retrieved as text	Y	-
	get/set_namedresult – conversion to named tuples	Y	-
	get/set_decimal – decimal type to be used for numeric values	Y	-
	get/set_decimal_point – decimal mark used for monetary values	Y	-
	get/set_bool – whether boolean values are returned as bool objects	Y	-
get/set_array – whether arrays are returned as list objects	Y	-	

PyGreSQL		Yes	Remarks
	get/set_bytea_escaped – whether bytea data is returned escaped	Y	-
	get/set_jsondecode – decoding JSON format	Y	-
	get/set_cast_hook – fallback typecast function	Y	-
	get/set_datestyle – assume a fixed date style	Y	-
	get/set_typecast – custom typecasting	Y	-
	cast_array/record – fast parsers for arrays and records	Y	-
	Type helpers	Y	-
	Module constants	Y	-
Connection – The connection object	query – execute a SQL command string	Y	-
	send_query – executes a SQL command string asynchronously	Y	-
	query_prepared – execute a prepared statement	Y	-
	prepare – create a prepared statement	Y	-
	describe_prepared – describe a prepared statement	Y	-
	reset – reset the connection	Y	-
	poll – completes an asynchronous connection	Y	-
	cancel – abandon processing of current SQL command	Y	-
	close – close the database connection	Y	-
	transaction – get the current transaction state	Y	-

PyGreSQL		Yes	Remarks
	parameter – get a current server parameter setting	Y	-
	date_format – get the currently used date format	Y	-
	fileno – get the socket used to connect to the database	Y	-
	set_non_blocking - set the non-blocking status of the connection	Y	-
	is_non_blocking - report the blocking status of the connection	Y	-
	getnotify – get the last notify from the server	N	The database does not support listen/notify .
	inserttable – insert a list into a table	Y	Use double quotation marks (""") to quote \n in the copy command.
	get/set_notice_receiver – custom notice receiver	Y	-
	putline – write a line to the server socket [DA]	Y	-
	getline – get a line from server socket [DA]	Y	-
	endcopy – synchronize client and server [DA]	Y	-
	locreate – create a large object in the database [LO]	N	Operations related to large objects
	getlo – build a large object from given oid [LO]	N	Operations related to large objects
	loimport – import a file to a large object [LO]	N	Operations related to large objects
	Object attributes	Y	-

PyGreSQL		Yes	Remarks
The DB wrapper class	Initialization	Y	-
	pkey – return the primary key of a table	Y	-
	get_databases – get list of databases in the system	Y	-
	get_relations – get list of relations in connected database	Y	-
	get_tables – get list of tables in connected database	Y	-
	get_attnames – get the attribute names of a table	Y	-
	has_table_privilege – check table privilege	Y	-
	get/set_parameter – get or set run-time parameters	Y	-
	begin/commit/rollback/savepoint/release – transaction handling	Y	-
	get – get a row from a database table or view	Y	-
	insert – insert a row into a database table	Y	-
	update – update a row in a database table	Y	-
	upsert – insert a row with conflict resolution	Y	-
	query – execute a SQL command string	Y	-
	query_formatted – execute a formatted SQL command string	Y	-
	query_prepared – execute a prepared statement	Y	-
prepare – create a prepared statement	Y	-	

PyGreSQL		Yes	Remarks
	describe_prepared – describe a prepared statement	Y	-
	delete_prepared – delete a prepared statement	Y	-
	clear – clear row values in memory	Y	-
	delete – delete a row from a database table	Y	A tuple must have unique key or primary key.
	truncate – quickly empty database tables	Y	-
	get_as_list/dict – read a table as a list or dictionary	Y	-
	escape_literal/identifier/string/bytea – escape for SQL	Y	-
	unescape_bytea – unescape data retrieved from the database	Y	-
	encode/decode_json – encode and decode JSON data	Y	-
	use_regtypes – determine use of regular type names	Y	-
	notification_handler – create a notification handler	N	The database does not support listen/notify .
	Attributes of the DB wrapper class	Y	-
Query methods	getresult – get query values as list of tuples	Y	-
	dictresult/dictiter – get query values as dictionaries	Y	-
	namedresult/namediter – get query values as named tuples	Y	-

PyGreSQL		Yes	Remarks
	scalarresult/scalariter – get query values as scalars	Y	-
	one/onedict/onenamed/onescalar – get one result of a query	Y	-
	single/singledict/singlenamed/singlescalar – get single result of a query	Y	-
	listfields – list fields names of previous query result	Y	-
	fieldname, fieldnum – field name/number conversion	Y	-
	fieldinfo – detailed info about query result fields	Y	-
	ntuples – return number of tuples in query object	Y	-
	memsize – return number of bytes allocated by query result	Y	-
LargeObject – Large Objects	open – open a large object	N	Operations related to large objects
	close – close a large object	N	Operations related to large objects
	read, write, tell, seek, unlink – file-like large object handling	N	Operations related to large objects
	size – get the large object size	N	Operations related to large objects
	export – save a large object to a file	N	Operations related to large objects
	Object attributes	N	Operations related to large objects
The Notification Handler	Instantiating the notification handler	N	The database does not support listen/notify .

PyGreSQL		Yes	Remarks
	Invoking the notification handler	N	The database does not support listen/notify .
	Sending notifications	N	The database does not support listen/notify .
	Auxiliary methods	N	The database does not support listen/notify .
pgdb			
Module functions and constants	connect – Open a PostgreSQL connection	Y	-
	get/set/reset_typecast – Control the global typecast functions	Y	-
	Module constants	Y	-
	Errors raised by this module	Y	-
Connection – The connection object	close – close the connection	Y	-
	commit – commit the connection	Y	-
	rollback – roll back the connection	Y	-
	cursor – return a new cursor object	Y	-
	Attributes that are not part of the standard	Y	-
Cursor – The cursor object	description – details regarding the result columns	Y	-
	rowcount – number of rows of the result	Y	-
	close – close the cursor	Y	-
	execute – execute a database operation	Y	-

PyGreSQL		Yes	Remarks
	executemany – execute many similar database operations	Y	-
	callproc – Call a stored procedure	Y	-
	fetchone – fetch next row of the query result	Y	-
	fetchmany – fetch next set of rows of the query result	Y	-
	fetchall – fetch all rows of the query result	Y	-
	arraysize - the number of rows to fetch at a time	Y	-
	Methods and attributes that are not part of the standard	Y	-
Type – Type objects and constructors	Type constructors	Y	-
	Type objects	Y	-

Using the Third-Party Function Library PyGreSQL to Connect to a Cluster (Linux)

Step 1 Log in to the Linux environment as user **root**.

Step 2 Run the following command to create the **python_dws.py** file:

```
vi python_dws.py
```

Copy and paste the following content to the **python_dws.py** file:

```
#!/usr/bin/env python3
# *_ encoding:utf-8 *_

from __future__ import print_function

import pg

def create_table(connection):
    print("Begin to create table")
    try:
        connection.query("drop table if exists test;"
            "create table test(id int, name text);")
    except pg.InternalError as e:
        print(e)
    else:
        print("Table created successfully")

def insert_data(connection):
```

```
print("Begin to insert data")
try:
    connection.query("insert into test values(1,'number1');")
    connection.query("insert into test values(2,'number2');")
    connection.query("insert into test values(3,'number3');")
except pg.InternalError as e:
    print(e)
else:
    print("Insert data successfully")

def update_data(connection):
    print("Begin to update data")
    try:
        result = connection.query("update test set name = 'numberupdated' where id=1;")
        print("Total number of rows updated :", result)
        result = connection.query("select * from test order by 1;")
        rows = result.getresult()
        for row in rows:
            print("id = ", row[0])
            print("name = ", row[1], "\n")
    except pg.InternalError as e:
        print(e)
    else:
        print("After Update, Operation done successfully")

def delete_data(connection):
    print("Begin to delete data")
    try:
        result = connection.query("delete from test where id=3;")
        print("Total number of rows deleted :", result)
        result = connection.query("select * from test order by 1;")
        rows = result.getresult()
        for row in rows:
            print("id = ", row[0])
            print("name = ", row[1], "\n")
    except pg.InternalError as e:
        print(e)
    else:
        print("After Delete, Operation done successfully")

def select_data(connection):
    print("Begin to select data")
    try:
        result = connection.query("select * from test order by 1;")
        rows = result.getresult()
        for row in rows:
            print("id = ", row[0])
            print("name = ", row[1])
    except pg.InternalError as e:
        print(e)
        print("select failed")
    else:
        print("Operation done successfully")

if __name__ == '__main__':
    try:
        conn = pg.DB(host='10.154.70.231',
                    port=8000,
                    dbname='gaussdb', # Database to be connected
                    user='dbadmin',
                    passwd='password') # Database user password
    except pg.InternalError as ex:
        print(ex)
        print("Connect database failed")
    else:
```

```
print("Opened database successfully")
create_table(conn)
insert_data(conn)
select_data(conn)
update_data(conn)
delete_data(conn)
conn.close()
```

Alternatively, use the dbapi interface.

```
#!/usr/bin/python
# -*- coding: UTF-8 -*-

from __future__ import print_function

import pg
import pgdb

def create_table(connection):
    print("Begin to create table")
    try:
        cursor = connection.cursor()
        cursor.execute("drop table if exists test;"
                       "create table test(id int, name text);")
        connection.commit()
    except pg.InternalError as e:
        print(e)
    else:
        print("Table created successfully")
        cursor.close()

def insert_data(connection):
    print("Begin to insert data")
    try:
        cursor = connection.cursor()
        cursor.execute("insert into test values(1,'number1');")
        cursor.execute("insert into test values(2,'number2');")
        cursor.execute("insert into test values(3,'number3');")
        connection.commit()
    except pg.InternalError as e:
        print(e)
    else:
        print("Insert data successfully")
        cursor.close()

def update_data(connection):
    print("Begin to update data")
    try:
        cursor = connection.cursor()
        cursor.execute("update test set name = 'numberupdated' where id=1;")
        connection.commit()
        print("Total number of rows updated :", cursor.rowcount)
        cursor.execute("select * from test;")
        rows = cursor.fetchall()
        for row in rows:
            print("id = ", row[0])
            print("name = ", row[1], "\n")
    except pg.InternalError as e:
        print(e)
    else:
        print("After Update, Operation done successfully")

def delete_data(connection):
    print("Begin to delete data")
    try:
        cursor = connection.cursor()
```

```
cursor.execute("delete from test where id=3;")
connection.commit()
print("Total number of rows deleted :", cursor.rowcount)
cursor.execute("select * from test;")
rows = cursor.fetchall()
for row in rows:
    print("id = ", row[0])
    print("name = ", row[1], "\n")
except pg.InternalError as e:
    print(e)
else:
    print("After Delete,Operation done successfully")

def select_data(connection):
    print("Begin to select data")
    try:
        cursor = connection.cursor()
        cursor.execute("select * from test;")
        rows = cursor.fetchall()
        for row in rows:
            print("id = ", row[0])
            print("name = ", row[1], "\n")
    except pg.InternalError as e:
        print(e)
        print("select failed")
    else:
        print("Operation done successfully")
        cursor.close()

if __name__ == '__main__':
    try:
        conn = pgdb.connect(host='10.154.70.231',
                            port='8000',
                            database='gaussdb', # Database to be connected
                            user='dbadmin',
                            password='password') # Database user password
    except pg.InternalError as ex:
        print(ex)
        print("Connect database failed")
    else:
        print("Opened database successfully")
        create_table(conn)
        insert_data(conn)
        select_data(conn)
        update_data(conn)
        delete_data(conn)
        conn.close()
```

Step 3 Change the public network address, cluster port number, database name, database username, and database password in the `python_dws.py` file based on the actual cluster information.

 **NOTE**

The PyGreSQL API does not provide the connection retry capability. You need to implement the retry processing in the service code.

```
conn = pgdb.connect(host='10.154.70.231',
                    port='8000',
                    database='gaussdb', # Database to be connected
                    user='dbadmin',
                    password='password') # Database user password
```

Step 4 Run the following command to connect to the cluster using the third-party function library PyGreSQL:

```
python python_dws.py
```

```
----End
```

Using the Third-Party Function Library PyGreSQL to Connect to a Cluster (Windows)

Step 1 In the Windows operating system, click the **Start** button, enter **cmd** in the search box, and click **cmd.exe** in the result list to open the command-line interface (CLI).

Step 2 In the CLI, run the following command to create the **python_dws.py** file:

```
type nul> python_dws.py
```

Copy and paste the following content to the **python_dws.py** file:

```
#!/usr/bin/env python3
# -*- encoding:utf-8 -*-

from __future__ import print_function

import pg

def create_table(connection):
    print("Begin to create table")
    try:
        connection.query("drop table if exists test;"
                          "create table test(id int, name text);")
    except pg.InternalError as e:
        print(e)
    else:
        print("Table created successfully")

def insert_data(connection):
    print("Begin to insert data")
    try:
        connection.query("insert into test values(1,'number1');")
        connection.query("insert into test values(2,'number2');")
        connection.query("insert into test values(3,'number3');")
    except pg.InternalError as e:
        print(e)
    else:
        print("Insert data successfully")

def update_data(connection):
    print("Begin to update data")
    try:
        result = connection.query("update test set name = 'numberupdated' where id=1;")
        print("Total number of rows updated :", result)
        result = connection.query("select * from test order by 1;")
        rows = result.getresult()
        for row in rows:
            print("id = ", row[0])
            print("name = ", row[1], "\n")
    except pg.InternalError as e:
        print(e)
    else:
        print("After Update, Operation done successfully")

def delete_data(connection):
    print("Begin to delete data")
    try:
        result = connection.query("delete from test where id=3;")
        print("Total number of rows deleted :", result)
```



```
result = connection.query("select * from test order by 1;")
rows = result.getresult()
for row in rows:
    print("id = ", row[0])
    print("name = ", row[1], "\n")
except pg.InternalError as e:
    print(e)
else:
    print("After Delete,Operation done successfully")

def select_data(connection):
    print("Begin to select data")
    try:
        result = connection.query("select * from test order by 1;")
        rows = result.getresult()
        for row in rows:
            print("id = ", row[0])
            print("name = ", row[1])
    except pg.InternalError as e:
        print(e)
        print("select failed")
    else:
        print("Operation done successfully")

if __name__ == '__main__':
    try:
        conn = pg.DB(host='10.154.70.231',
                    port=8000,
                    dbname='gaussdb', # Database to be connected
                    user='dbadmin',
                    passwd='password') # Database user password
    except pg.InternalError as ex:
        print(ex)
        print("Connect database failed")
    else:
        print("Opened database successfully")
        create_table(conn)
        insert_data(conn)
        select_data(conn)
        update_data(conn)
        delete_data(conn)
        conn.close()
```

Alternatively, use the dbapi interface.

```
#!/usr/bin/python
# -*- coding: UTF-8 -*-

from __future__ import print_function

import pg
import pgdb

def create_table(connection):
    print("Begin to create table")
    try:
        cursor = connection.cursor()
        cursor.execute("drop table if exists test;"
                      "create table test(id int, name text);")
        connection.commit()
    except pg.InternalError as e:
        print(e)
    else:
        print("Table created successfully")
        cursor.close()
```

```
def insert_data(connection):
    print("Begin to insert data")
    try:
        cursor = connection.cursor()
        cursor.execute("insert into test values(1,'number1');")
        cursor.execute("insert into test values(2,'number2');")
        cursor.execute("insert into test values(3,'number3');")
        connection.commit()
    except pg.InternalError as e:
        print(e)
    else:
        print("Insert data successfully")
        cursor.close()

def update_data(connection):
    print("Begin to update data")
    try:
        cursor = connection.cursor()
        cursor.execute("update test set name = 'numberupdated' where id=1;")
        connection.commit()
        print("Total number of rows updated :", cursor.rowcount)
        cursor.execute("select * from test;")
        rows = cursor.fetchall()
        for row in rows:
            print("id = ", row[0])
            print("name = ", row[1], "\n")
    except pg.InternalError as e:
        print(e)
    else:
        print("After Update, Operation done successfully")

def delete_data(connection):
    print("Begin to delete data")
    try:
        cursor = connection.cursor()
        cursor.execute("delete from test where id=3;")
        connection.commit()
        print("Total number of rows deleted :", cursor.rowcount)
        cursor.execute("select * from test;")
        rows = cursor.fetchall()
        for row in rows:
            print("id = ", row[0])
            print("name = ", row[1], "\n")
    except pg.InternalError as e:
        print(e)
    else:
        print("After Delete, Operation done successfully")

def select_data(connection):
    print("Begin to select data")
    try:
        cursor = connection.cursor()
        cursor.execute("select * from test;")
        rows = cursor.fetchall()
        for row in rows:
            print("id = ", row[0])
            print("name = ", row[1], "\n")
    except pg.InternalError as e:
        print(e)
        print("select failed")
    else:
        print("Operation done successfully")
        cursor.close()

if __name__ == '__main__':
```

```
try:
    conn = pgdb.connect(host='10.154.70.231',
                        port='8000',
                        database='gaussdb', # Database to be connected
                        user='dbadmin',
                        password='password') # Database user password
except pg.InternalError as ex:
    print(ex)
    print("Connect database failed")
else:
    print("Opened database successfully")
    create_table(conn)
    insert_data(conn)
    select_data(conn)
    update_data(conn)
    delete_data(conn)
    conn.close()
```

- Step 3** Change the public network address, cluster port number, database name, database username, and database password in the `python_dws.py` file based on the actual cluster information.

The PyGreSQL API does not provide the connection retry capability. You need to implement the retry processing in the service code.

```
conn = pgdb.connect(host='10.154.70.231',
                    port='8000',
                    database='gaussdb', # Database to be connected
                    user='dbadmin',
                    password='password') # Database user password
```

- Step 4** Run the following command to connect to the cluster using the third-party function library PyGreSQL:

```
python python_dws.py
```

----End

4.8 Managing Database Connections

Scenario

By default, a database supports a certain number of connections. Administrators can manage database connections to learn about the connection performance of the current database or increase the connection limit so that more users or applications can connect to the database at the same time.

Maximum Number of Connections

The number of connections supported by a cluster depends on its node flavor.

Table 4-17 Number of supported connections

Parameter	Description	Number of CN Connections	Number of DN Connections
max_connections	Specifies the maximum number of concurrent connections to the database.	800	Max (Number of vCPU cores/Number of DNs on a single node x 120 + 24, 5000)
max_pool_size	Specifies the maximum number of connections between the connection pool of a CN and another CN or DN.		
max_prepared_transactions	Specifies the maximum number of transactions that can stay in the prepared state simultaneously.		

 **NOTE**

For details about CNs and DNs, see [Logical Cluster Architecture](#).

Viewing the Maximum Number of Connections

Method 1: Post cluster creation, select the cluster name on the GaussDB(DWS) console to access the **Parameter Modification** page and view the value of **max_connections**.

Method 2: Use an SQL client tool to establish a database connection within the cluster and execute an SQL query to obtain the value of **max_connections**.
`SHOW max_connections;`

Information similar to the following is displayed, showing that the maximum number of database connections is **200** by default.

```
max_connections
-----
200
(1 row)
```

Viewing the Number of Used Connections

Step 1 Use the SQL client tool to connect to the database in a cluster.

Step 2 View the number of connections in scenarios described in [Table 4-18](#).

NOTICE

Except for database and user names that are enclosed with double quotation marks (") during creation, uppercase letters are not allowed in the database and user names in the commands in the following table.

Table 4-18 Viewing the number of connections

Description	Command
View the maximum number of sessions connected to a specific user.	<p>Run the following command to view the maximum number of sessions connected to user dbadmin.</p> <pre>SELECT ROLNAME,ROLCONNLIMIT FROM PG_ROLES WHERE ROLNAME='dbadmin';</pre> <p>Information similar to the following is displayed. -1 indicates that the number of sessions connected to user dbadmin is not limited.</p> <pre>rolname rolconlimit -----+----- dwsadmin -1 (1 row)</pre>
View the number of session connections that have been used by a user.	<p>Run the following command to view the number of session connections that have been used by dbadmin.</p> <pre>SELECT COUNT(*) FROM V\$SESSION WHERE USERNAME='dbadmin';</pre> <p>Information similar to the following is displayed. 1 indicates the number of session connections used by user dbadmin.</p> <pre>count ----- 1 (1 row)</pre>
View the maximum number of sessions connected to a specific database.	<p>View the upper limit of connections used by the gaussdb database.</p> <pre>SELECT DATNAME,DATCONNLIMIT FROM PG_DATABASE WHERE DATNAME='gaussdb';</pre> <p>Information similar to the following is displayed. -1 indicates that the number of sessions connected to the gaussdb database is not limited.</p> <pre>datname datconlimit -----+----- gaussdb -1 (1 row)</pre>
View the number of session connections that have been used by a database.	<p>View the number of session connections that have been used by the gaussdb database.</p> <pre>SELECT COUNT(*) FROM PG_STAT_ACTIVITY WHERE DATNAME='gaussdb';</pre> <p>Information similar to the following is displayed. 1 indicates the number of session connections used by the gaussdb database.</p> <pre>count ----- 1 (1 row)</pre>

Description	Command
View the number of session connections that have been used by all users.	Run the following command to view the number of session connections that have been used by all users: <pre>SELECT COUNT(*) FROM PG_STAT_ACTIVITY; count ----- 10 (1 row)</pre>

----End

5 Monitoring and Alarms

5.1 Dashboard

There are two types of Dashboard pages, [Dashboard Without Cluster Overview](#) and [Dashboard With Cluster Overview](#). The dashboard page displayed on your console is determined by whether you have purchased any cluster or not.


Dashboard Without Cluster Overview

The dashboard page consists of the following modules:

- Process

This module describes how to create a cluster, connect to a cluster, and import sample data from OBS to GaussDB(DWS). You can click **Learn more** to check more information.

Figure 5-1 Process



- Features

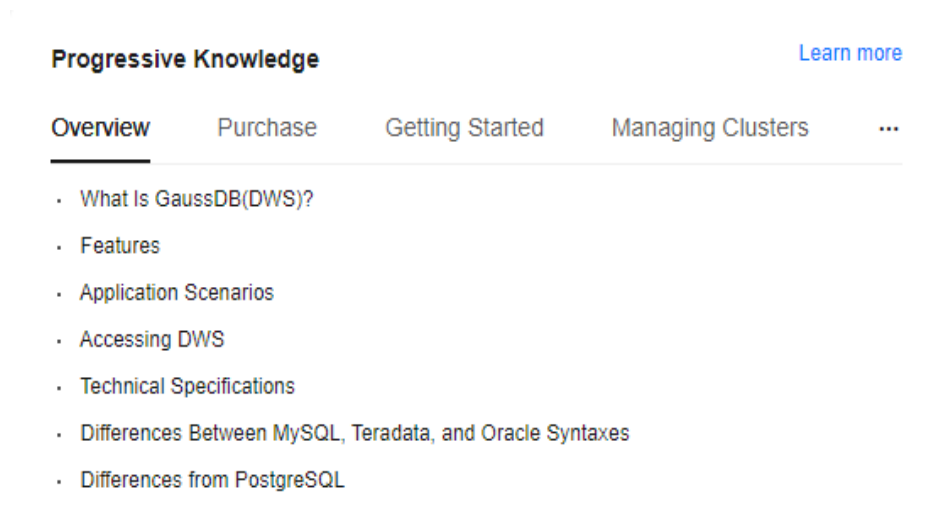
This module describes multiple powerful functions of GaussDB(DWS), including SQL compatibility, cluster snapshot, cluster disaster recovery, database monitoring, resource management, and online O&M. You can quickly learn how they work and use them as required.

Figure 5-2 Features



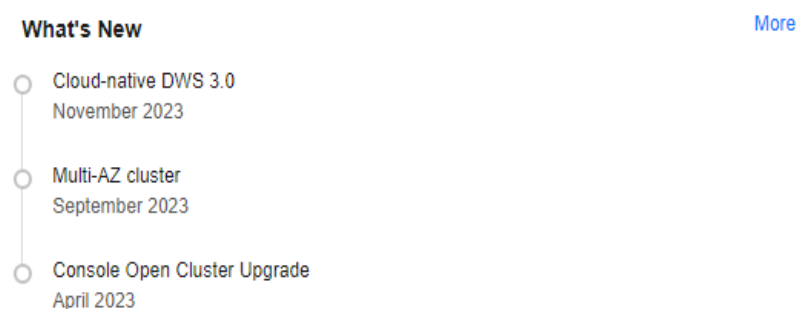
- Progressive Knowledge
This module helps you start as a beginner and become an expert in the real-time, secure, reliable enterprise-grade GaussDB(DWS) data warehouse.

Figure 5-3 Progressive knowledge



- What's New
By default, the latest three features of GaussDB(DWS) are displayed. You can click **More** to check more information.

Figure 5-4 What's new



Dashboard With Cluster Overview

The dashboard page contains Progressive Knowledge, What's New, Features, and the following modules:

- Resources

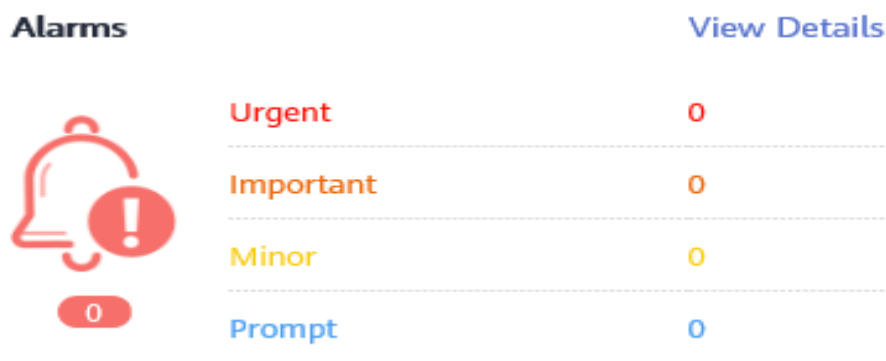
In the Resources module, you can view the number of available resources, including **Available/Total Clusters**, **Available/Total Nodes**, and **Total Capacity**.

Figure 5-5 Resources



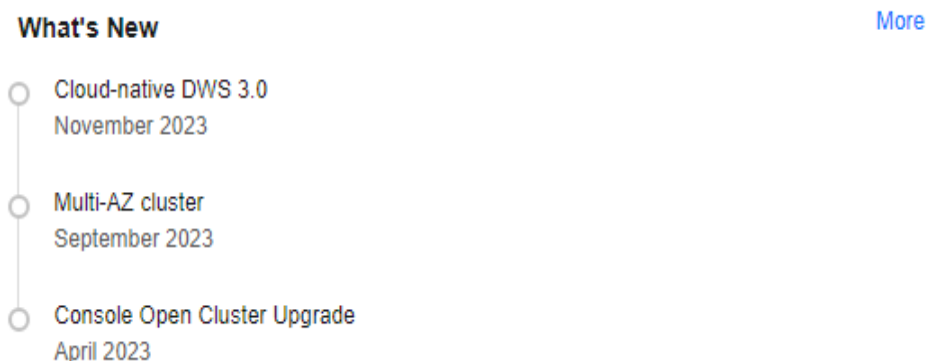
- Alarms
Alarms are classified by severity: **Urgent**, **Important**, **Minor**, **Prompt**. For details, see [Alarms](#).

Figure 5-6 Alarms



- Recent Events
Events are change records of user cluster status. Events can be triggered by user operations or cluster status changes. For details, see [Event Notifications](#).

Figure 5-7 Recent events



- Main cluster metrics:
 - Cluster CPU Usage
 - Cluster Memory Usage
 - Cluster Disk Usage

Figure 5-8 Main cluster metrics



- For details about Progressive Knowledge, see [Progressive Knowledge](#).
- For details about product changes, see [What' New](#).
- For details about GaussDB(DWS) features, see [Features](#).

5.2 Databases Monitoring (DMS)

5.2.1 Database Monitoring Overview

Overview

DMS is provided by GaussDB(DWS) to ensure the fast and stable running of databases. It collects, monitors, and analyzes the disk, network, and OS metric data used by the service database, as well as key performance metric data of cluster running. It also diagnoses database hosts, instances, and service SQL statements based on the collected metrics to expose key faults and performance problems in a database in a timely manner, and guides customers to optimize and resolve the problems.

NOTE

- Database monitoring is supported by 8.1.1.200 and later versions.
- The hybrid data warehouse (standalone) does not support database monitoring.
- The database monitoring function and Cloud Eye monitor different data sources. In database monitoring, the size of a database is the total disk space used by the database, including the space occupied due to bloating.

Entering the Database Monitoring Page

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
- Step 3** In the **Operation** column of the target cluster, choose **Monitoring Panel**. The database monitoring page is displayed.

----End

5.2.2 Monitoring Metrics

You can check the status and available resources of a cluster and learn about its real-time resource consumption through the GaussDB(DWS) monitoring items.

Table 5-1 describes GaussDB(DWS) monitoring metrics.

Table 5-1 GaussDB(DWS) monitoring metrics

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
Abnormal Job Monitoring Statistics	Query ID	String	30s	Collect data regarding abnormal jobs whose status is aborted in the PGXC_WLM_SESSION_HISTORY view.	N/A
	Statement executed for exception handling	String			
	Block time before the statement is executed	≥ 0			
	Elapsed time when the statement is executed	≥ 0			
	Total time used by the CPU on the DN when the statement is executed for exception handling	≥ 0			
	CPU usage skew on the DN when the statement is executed for exception handling	≥ 0			
	cgroups used for exception handling during statement execution	String			
	Status of a statement after exception handling	String			
	Exception handling action executed by the statement	String			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Reason why the statement is processed abnormally	String			
Node Status Statistics	Host name	String	60s	Gather information on the current state and specifics of each ECS host instance (VM) in a cluster.	N/A
	Host status	String			
Instances by Status	Host name	String	60s	Gather information on the current state and specifics of each CN/DN instance in the cluster.	N/A
	Instance type	String			
	Instance role	String			
	Instance status	String			
	Cause of the instance status.	String			
Cluster Status	Cluster status	String	30s	Monitor the cluster status.	N/A
	Whether an primary/standby switchover has occurred	String			
	Whether redistribution has occurred	String			
	Whether the current cluster is read-only	String			
CPU usage	The default value is ALL .	String	30s	Gather CPU usage data from ECS instances (VMs) to monitor node CPU usage. High CPU usage can lead to SQL queuing and slow queries.	85 %
	User-mode CPU time %	≥ 0.0			
	CPU time the process with a negative nice value (%)	≥ 0.0			
	Kernel mode time (%)	≥ 0.0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	I/O wait time (%)	≥ 0.0			
	Hard interruption time (%)	≥ 0.0			
	Software interrupt time (%)	≥ 0.0			
	Time spent by the virtual CPU in involuntary waiting when the virtual machine manager serves another virtual processor (%)	≥ 0.0			
	Time spent running the virtual processor (%)	≥ 0.0			
	Idle time except for disk waiting operations (%)	≥ 0.0			
	Hyper-threading capability	Yes/No			
	Hyper-threading Enabled (Yes/No)	Yes/No			
	Number of processes in the runnable status	≥ 0			
	Number of processes waiting for I/O completion	≥ 0			
Active Sessions	Database name	String	30s	Collect information on the current cluster's active sessions.	N/A
	Instance name	String			
	Number of all user sessions	≥ 0			
	Distinct username	≥ 0			
	Distinct application name	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Number of queries in the active or fastpathfunctioncall state.	≥ 0			
Disk Capacity Statistics	Instance name	String	86400s	Collect statistics on the disk space used by each database in the current cluster.	N/A
	Database name	String			
	Database size	≥ 0			
Transaction Status	Database name	String	60s	Gather data on the running status of databases in the current cluster, such as the number of updated, deleted, and inserted rows, transactions, and deadlocks.	N/A
	Instance name	String			
	Number of rows returned through a global database scan	≥ 0			
	Number of rows returned by querying indexes in the database	≥ 0			
	Number of rows inserted by queries in this database	≥ 0			
	Number of rows updated by queries in this database	≥ 0			
	Number of rows deleted by queries in this database	≥ 0			
	Number of transactions in this database that have been committed	≥ 0			
	Number of transactions in this database that have been rolled back	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Number of deadlocks detected in this database	≥ 0			
	Number of disk blocks read in this database	≥ 0			
	Number of disk blocks found in the buffer cache in the current database, that is, the number of blocks hit in the cache. (This only includes hits in the GaussDB(DWS) buffer cache, not in the file system cache.)	≥ 0			
	Time spent reading data file blocks by backends in this database, in milliseconds	≥ 0.0			
	Time spent reading data file blocks by backends in this database, in milliseconds	≥ 0.0			
	Number of queries canceled due to database recovery conflicts (conflicts that occur only on the standby server).	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Number of temporary files created by queries in this database. This parameter calculates all temporary files (such as sorting or hashing) and ignores the log_temp_files setting.	≥ 0			
	Total amount of data written to temporary files by queries in this database. This parameter calculates all temporary files and ignores the log_temp_files setting.	≥ 0			
	Database capacity, in bytes.	≥ 0			
	Number of rows returned by global database scanning in a unit time	≥ 0			
	Number of rows returned by querying indexes in the database in a unit time	≥ 0			
	Number of rows inserted through database query in a unit time	≥ 0			
	Number of rows updated through database query in unit time	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Number of rows deleted by querying the database in a unit time	≥ 0			
	Number of transactions that have been submitted in the database per unit time	≥ 0			
	Number of transactions that have been rolled back in the database per unit time	≥ 0			
	Number of deadlocks retrieved in the database per unit time	≥ 0			
File Handle	Name of the disk file system	String	30s	Gather data on the disk inode information of the cluster to monitor inode usage. High inode usage can pose risks.	90%
	Total inode capacity (unit: KB)	≥ 0			
	Used capacity (unit: KB)	≥ 0			
Node Disk Usage	Name of the disk file system	≥ 0	30s	Track the disk usage of every ECS instance and switch the cluster to read-only mode when it hits 90%.	90%
	Total space (unit: KB)	≥ 0			
	Used capacity (unit: KB)	≥ 0			
	Available capacity (unit: KB)	≥ 0			
	Disk usage	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
gsar NIC Usage Statistics	Node name	String	30s	Monitor the running status of the gsar NIC.	N/A
	NIC name	String			
	NIC IP address	String			
	Data received by the NIC (unit: KB)	≥ 0			
	Number of packets received by the NIC	≥ 0			
	Average length of received packets (unit: byte)	≥ 0			
	Number of received data packets that are dropped by the NIC	≥ 0			
	Port transmit discard rate	≥ 0.0			
	Data sent by the network adapter (unit: KB)	≥ 0			
	Number of packets sent by the NIC	≥ 0			
	Average length of received packets (unit: byte)	≥ 0			
gsar TCP Statistics	Number of retransmitted TCP packets due to timeout	≥ 0	30s	Monitor the TCP retransmission rate.	Number of retransmitted TCP packets > 0
	Number of sent TCP packets	≥ 0			
	Number of retransmitted TCP packets	≥ 0			
	TCP retransmission rate	≥ 0.0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
Node Disk I/O Statistics	Disk name (devicename)	String	30s	Keep track of the I/O status of every disk on a node, which is indicated by the data transfer rates and the number of read and write operations. Excessive data transfer rates could affect cluster services.	350 MB/s
	Number of transmissions per second (transferpersecond). The size of each transmission is unknown.	≥ 0.0			
	Amount of data read from the device per second (unit: KB)	≥ 0.0			
	Amount of data written to the device per second (unit: KB)	≥ 0.0			
	Total amount of read data (unit: KB)	≥ 0.0			
	Total amount of written data (unit: KB)	≥ 0.0			
	Number of times that read requests to the device are combined per second	≥ 0.0			
	Number of times that write requests to the device are combined per second	≥ 0.0			
	Number of completed reads per second	≥ 0.0			
	Number of completed writes per second	≥ 0.0			
	Amount of data read per second (unit: KB)	≥ 0.0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Amount of data written per second (unit: KB)	≥ 0.0			
	Average data volume of each I/O operation (unit: sector)	≥ 0.0			
	Average request queue length	≥ 0.0			
	Average waiting time for each I/O request (unit: ms)	≥ 0.0			
	Average processing time for each I/O request (unit: ms)	≥ 0.0			
	Percentage of the time when the I/O queue is not empty (I/O operation time divided by the total time)	≥ 0.0			
Instance Memory Monitoring Statistics	Instance name	String	60s	Monitor instance and dynamic memory, gather memory usage statistics for each CN and DN. If the instance memory usage goes beyond the threshold, there may not be enough instance memory in the cluster.	85 %
	Memory size occupied by the instance	≥ 0.0			
	Memory size used by a process	≥ 0.0			
	Maximum dynamic memory	≥ 0.0			
	Used dynamic memory	≥ 0.0			
	Dynamic peak memory	≥ 0.0			
	Maximum dynamic shared memory context	≥ 0.0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Dynamic peak value of the shared memory context	≥ 0.0			
	Maximum shared memory	≥ 0.0			
	Used shared memory	≥ 0.0			
	Maximum memory allowed by column store	≥ 0.0			
	Memory used in column store	≥ 0.0			
	Maximum memory that can be used by the communication library	≥ 0.0			
	Used memory size of the communication library	≥ 0.0			
	Peak memory usage of the communication library	≥ 0.0			
	Maximum memory that can be used by top SQLs to record historical job monitoring information	≥ 0.0			
	Peak memory usage of the top SQL that records historical job monitoring information	≥ 0.0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Memory used by top SQLs to record historical job monitoring information	≥ 0.0			
	Other used memory	≥ 0.0			
	Memory size occupied by pooler connections	≥ 0.0			
	Memory size occupied by pooler idle connections	≥ 0.0			
	Memory size used by column-store compression and decompression	≥ 0.0			
	Memory reserved for the UDFWorker process	≥ 0.0			
	Memory size used by the MMAP	≥ 0.0			
Instance Resource Statistics	Instance name	String	60s	Gather resource usage statistics for each instance in a cluster.	85 %
	Read the value (CPU usage %) in postmaster.pid/cm_server.pid/gtm.pid/etcd.pid .	≥ 0.0			
	Read the value (memory usage %) in postmaster.pid/cm_server.pid/gtm.pid/etcd.pid .	≥ 0.0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
Instance Disk Size Statistics	Instance name	String	86400s	Collect statistics on instance disk usage and monitor it.	85 %
	Storage location	String			
	Disk space used by all databases on the current instance	≥ 0			
Node Memory Statistics	Size of all available RAMs, that is, the remaining physical memory minus the reserved bits and kernel usage (unit: KB)	≥ 0	30s	Gather memory usage statistics for the ECS instance where the cluster is located. This metric tracks statistics on the VM OS-level memory, which differs from the instance memory.	70 %
	Unused memory in the system. The value is lowfree+highfree (unit: KB).	≥ 0			
	Size of the cache used for the block device (unit: KB)	≥ 0			
	Size of the file buffer (unit: KB)	≥ 0			
	Total swap space (unit: KB)	≥ 0			
	Size of the RAM memory temporarily stored in the swap file (unit: KB)	≥ 0			
	Memory size of the virus-infected page (unit: KB)	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
Network Status Statistics	NIC name	String	30s	Collect NIC status for each node in the cluster to monitor lost packets on the cluster NIC and network throughput.	N/A
	NIC status (up/down)	up/down			
	NIC speed (1,000 Mbit/s or 100 Mbit/s)	≥ 0			
	Total data received by the NIC (unit: byte)	≥ 0			
	Number of packets received by the NIC	≥ 0			
	Total number of NIC receiving errors	≥ 0			
	Number of received data packets that are dropped by the NIC	≥ 0			
	Number of FIFO buffer errors during reception	≥ 0			
	Number of received packet frame errors	≥ 0			
	Number of received compressed data packets	≥ 0			
	Number of received multicast frames	≥ 0			
	Total data sent by the NIC (unit: byte)	≥ 0			
	Total number of packets sent by the NIC	≥ 0			
	Total number of NIC sending errors	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Total number of data packets discarded by the NIC during transmission	≥ 0			
	Number of FIFO buffer errors during sending	≥ 0			
	Number of collisions detected on the sending interface	≥ 0			
	Number of carrier losses detected by the device driver during transmission	≥ 0			
	Number of sent compressed data packets	≥ 0			
	Specifies whether NIC multi-queue is supported.	Yes/No			
	NIC multi-queue is enabled.	Yes/No			
	Specifies the CPU affinity of a multi-queue NIC.	String			
	Indicates whether the NIC works in duplex mode.	String			
	Network speed	≥ 0.0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
Node SQL Statistics	Node name	String	60s	Use the PGXC_SQL_COUNT view to query the number of running SQL statements on each node and monitor the cluster's QPS.	N/A
	Username	String			
	Number of SELECT statements	≥ 0			
	Number of UPDATE statements	≥ 0			
	Number of INSERT statements	≥ 0			
	Number of DELETE statements	≥ 0			
	Number of MERGEINTO statements	≥ 0			
	Number of DDL statements	≥ 0			
	Number of DML statements	≥ 0			
	Number of DCL statements	≥ 0			
	Total response time of SELECT statements	≥ 0			
	Average response time of SELECT statements	≥ 0			
	Maximum response time of SELECT statements	≥ 0			
	Minimum response time of SELECT statements	≥ 0			
Total response time of UPDATE statements	≥ 0				

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Average response time of UPDATE statements	≥ 0			
	Maximum response time of UPDATE statements	≥ 0			
	Minimum response time of UPDATE statements	≥ 0			
	Total response time of DELETE statements	≥ 0			
	Average response time of DELETE statements	≥ 0			
	Maximum response time of DELETE statements	≥ 0			
	Minimum response time of DELETE statements	≥ 0			
	Total response time of INSERT statements	≥ 0			
	Average response time of INSERT statements	≥ 0			
	Maximum response time of INSERT statements	≥ 0			
	Minimum response time of INSERT statements	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Difference between the numbers of SELECT statements at two time points	≥ 0			
	Difference between the numbers of UPDATE statements at two time points	≥ 0			
	Difference between the numbers of INSERT statements at two time points	≥ 0			
	Difference between the numbers of DELETE statements at two time points	≥ 0			
	Difference between the numbers of MERGE INTO statements at two time points	≥ 0			
	Difference between the numbers of DDL statements at two time points	≥ 0			
	Difference between the numbers of DML statements at two time points	≥ 0			
	Difference between the numbers of DCL statements at two time points	≥ 0			
	Difference between the total SELECT response time at two time points	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Difference between the average SELECT response time at two time points	≥ 0			
	Difference between the maximum SELECT response time at two time points	≥ 0			
	Difference between the minimum SELECT response time at two time points	≥ 0			
	Difference between the total UPDATE response time at two time points	≥ 0			
	Difference between the average UPDATE response time at two time points	≥ 0			
	Difference between the maximum UPDATE response time at two time points	≥ 0			
	Difference between the minimum UPDATE response time at two time points	≥ 0			
	Difference between the total DELETE response time at two time points	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Difference between the average DELETE response time at two time points	≥ 0			
	Difference between the maximum DELETE response time difference at two time points	≥ 0			
	Difference between the minimum DELETE response time difference at two time points	≥ 0			
	Difference between the total INSERT response time at two time points	≥ 0			
	Difference between the average INSERT response time at two time points	≥ 0			
	Difference between the maximum INSERT response time at two time points	≥ 0			
	Difference between the minimum INSERT response time at two time points	≥ 0			
System Status Statistics	TCP protocol stack retransmission rate (%)	≥ 0.0	30s	Collect the TCP protocol, stack protocol, and stack retransmission rate of the ECS instance server.	>0

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
Top SQL Statistics	Database name	String	60s	Collect SQL statements running on each CN in the current cluster to identify deadlock SQL statements, slow SQL statements, or SQL statements with high resource usage. SQL statements can be scanned and killed on the management plane.	N/A
	Instance name	String			
	Thread ID (session ID and session ID)	String			
	Internal query_ID used for statement execution	String			
	Job type, which can be set using the guc parameter query_band . The default value is a null string.	String			
	The value is obtained from the query_band field. The position is 0.	String			
	The value is obtained from the query_band field. The position is 1.	String			
	Username used for connecting to the backend	String			
	Name of the application that is connected to the backend	String			
	IP address of the client connected to the backend. A null value suggests a Unix socket connection or an internal server process, such as autovacuum.	String			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Host name of the connected client, as reported by a reverse DNS lookup of client_addr . This column will only be non-null for IP connections, and only when log_hostname is enabled.	String			
	TCP port number used by the client to communicate with the backend. If the Unix socket is used, the value is -1.	String			
	Whether the backend is currently waiting on a lock. If yes, the value is true .	Yes/No			
	Time when the statement starts to be executed	≥ 0			
	Block time before the statement is executed. The unit is ms.	≥ 0			
	Duration that a statement has been executed. The unit is ms.	≥ 0			
	Estimated execution time of a statement. The unit is ms.	≥ 0			
	Estimated remaining time of statement execution. The unit is ms.	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Resource status in workload management	String			
	Resource pool used by the user	String			
	Priority of a job in the resource pool. The options are as follows: <ul style="list-style-type: none"> • 1: Low • 2: Medium • 4: High • 8: Rush 	≥ 0			
	Cgroup used by the statement	String			
	Minimum memory peak of a statement across all DNs. The unit is MB.	≥ 0			
	Maximum memory peak of a statement across all DNs. The unit is MB.	≥ 0			
	Average memory usage during statement execution. The unit is MB.	≥ 0			
	Memory usage skew of a statement among DNs.	≥ 0			
	Estimated memory used by the statement. The unit is MB.	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Statement spill information on all DNs	String			
	Minimum spilled data among all DNs when a spill occurs. The default value is 0 (unit: MB).	≥ 0			
	Maximum spilled data among all DNs when a spill occurs. The default value is 0 (unit: MB).	≥ 0			
	Average spilled data among all DNs when a spill occurs. The default value is 0 (unit: MB).	≥ 0			
	DN spill skew when a spill occurs	≥ 0			
	Minimum execution time of a statement across all DNs. The unit is ms.	≥ 0			
	Maximum execution time of a statement across all DNs. The unit is ms.	≥ 0			
	Average execution time of a statement across all DNs. The unit is ms.	≥ 0			
	Execution time skew of a statement among DNs.	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Minimum CPU time of a statement across all DNs. The unit is ms.	≥ 0			
	Maximum CPU time of a statement across all DNs. The unit is ms.	≥ 0			
	Total CPU time of a statement across all DNs. The unit is ms.	≥ 0			
	CPU time skew of a statement among DNs.	≥ 0			
	Warning. The following warnings and warnings related to SQL self-diagnosis tuning are displayed:	String			
	Average IOPS peak of a statement across all DNs. It is counted by ones in a column-store table and by ten thousands in a row-store table.	≥ 0			
	I/O skew of a statement among DNs.	≥ 0			
	Maximum IOPS of a statement across all DNs. It is counted by ones in a column-store table and by ten thousands in a row-store table.	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Minimum IOPS of a statement across all DNs. It is counted by ones in a column-store table and by ten thousands in a row-store table.	≥ 0			
	Query statement	String			
	Query plan	String			
	Real-time running status of the current query statement. The value can be active , idle , idleintransaction , idleintransaction (aborted) , fastpathfunctioncall , or disabled .	String			
	Running status of the query statement in the resource pool. The value can be pending , running , finished , aborted , active or unknown .	String			
	Statement attribute (ordinary , simple , complicated , or internal).	String			
	Fast and slow lanes (fast or slow)	String			
	Whether a query is a system query	Yes/No			
	Whether the query is a system query (adapting to monitor)	Yes/No			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Time when this process was started, that is, when the client connected to the server	≥ 0			
	Execution time so far	≥ 0			
	Time when the current transaction was started (NULL if no transactions are active) If the current query is the first of its transaction, the value of this column is the same as that of the query_start column.	≥ 0			
	Time when the status was changed in the previous time	≥ 0			
	Time when the statement starts to be executed	≥ 0			
	Actual execution duration of the statement, in seconds.	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
Historical Queries	Database name	String	180s	Collect historical data from the top SQL view to analyze SQL statements from the past, identify deadlock or slow SQL statements, and optimize cluster performance by rectifying any issues found.	N/A
	Instance name	String			
	Username	String			
	Name of the application that is connected to the backend	String			
	IP address of the client connected to the backend. A null value suggests a Unix socket connection or an internal server process, such as autovacuum.	String			
	Host name of the connected client, as reported by a reverse DNS lookup of client_addr . This column will only be non-null for IP connections, and only when log_hostname is enabled.	String			
	TCP port number used by the client to communicate with the backend. If the Unix socket is used, the value is -1.	String			
	Job type, which can be set using the guc parameter query_band . The default value is a null string.	String			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	The value is obtained from the query_band field. The position is 0.	String			
	The value is obtained from the query_band field. The position is 1.	String			
	Duration that a statement is blocked before being executed, including the statement parsing and optimization duration. The unit is ms.	≥ 0			
	Execution start time of a statement (unit: ms)	≥ 0			
	Execution end time of a statement (unit: ms)	≥ 0			
	Duration that a statement has been executed (unit: ms)	≥ 0			
	Estimated statement execution time (unit: ms)	≥ 0			
	Statement execution end status: <ul style="list-style-type: none"> ● Finished (normal) ● Aborted (Abnormal) 	String			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Exception information displayed if the final statement execution status is aborted .	String			
	Resource pool used by the user	String			
	Priority of a job in the resource pool. The options are as follows: <ul style="list-style-type: none"> ● 8: Rush ● 4: High ● 2: Medium ● 1: Low 	≥ 0			
	Cgroup used by the statement	String			
	Minimum memory peak of a statement across all DNs. The unit is MB.	≥ 0			
	Maximum memory peak of a statement across all DNs. The unit is MB.	≥ 0			
	Average memory usage during statement execution. The unit is MB.	≥ 0			
	Memory usage skew of a statement among DNs	≥ 0			
	Statement spill information on all DNs	String			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Minimum spilled data among all DNs when a spill occurs. The unit is MB. The default value is 0.	≥ 0			
	Maximum spilled data among all DNs when a spill occurs. The unit is MB. The default value is 0.	≥ 0			
	Average spilled data among all DNs when a spill occurs. The unit is MB. The default value is 0.	≥ 0			
	DN spill skew when a spill occurs	≥ 0			
	Minimum execution time of a statement across all DNs. The unit is ms.	≥ 0			
	Maximum execution time of a statement across all DNs. The unit is ms.	≥ 0			
	Average execution time of a statement across all DNs. The unit is ms.	≥ 0			
	Execution time skew of a statement among DNs.	≥ 0			
	Minimum CPU time of a statement across all DNs. The unit is ms.	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Maximum CPU time of a statement across all DNs. The unit is ms.	≥ 0			
	Total CPU time of a statement across all DNs. The unit is ms.	≥ 0			
	CPU time skew of a statement among DNs.	≥ 0			
	Minimum IOPS of a statement across all DNs. It is counted by ones in a column-store table and by ten thousands in a row-store table.	≥ 0			
	Maximum IOPS of a statement across all DNs. It is counted by ones in a column-store table and by ten thousands in a row-store table.	≥ 0			
	Average IOPS peak of a statement across all DNs. It is counted by ones in a column-store table and by ten thousands in a row-store table.	≥ 0			
	I/O skew of a statement among DNs	≥ 0			
	Warning. The following warnings and warnings related to SQL self-diagnosis tuning are displayed:	String			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Query ID	String			
	Statement executed	String			
	Execution plan of a statement	String			
	Logical cluster of the user running the statement	String			
Schema Usage Statistics	Database name	String	3600s	Monitor schema usage in a cluster by collecting usage data for each schema.	85 %
	Schema name	String			
	Used capacity (unit: byte)	≥ 0			
	Total capacity (unit: byte)	≥ 0			
Session Statistics	Database name	String	180s	Gather session information for each CN in a cluster, including statistics on idle sessions and lock holdings. Use the management console to clear any idle sessions.	N/A
	Instance name	String			
	Thread ID (It can be used as a session ID or a connection ID.)	String			
	Database username	String			
	User application name	String			
	Client address	String			
	Host name of the client	String			
	TCP port number used by the client to communicate with the background. If the Unix socket is used, the value is -1.	String			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Time when this process was started, that is, when the client connected to the server	≥ 0			
	Execution time so far.	≥ 0			
	Time when the current transaction was started (NULL if no transactions are active). If the current query is the first of its transaction, the value of this column is the same as that of the query_start column.	≥ 0			
	Time when the status was changed in the previous time	≥ 0			
	Whether the backend is currently waiting on a lock. If yes, the value is true .	Yes/No			
	Current overall state of this backend.	String			
	Resource pool used by the user	String			
	Actual execution duration of the statement, in seconds.	≥ 0			
	ID of a query	String			
	Time when the statement starts to be executed	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	cgroups currently used by the statement.	String			
	Lock type	String			
	Lock mode	String			
	Indicates whether to hold a lock when lock waiting exists. The value is true.	Yes/No			
	Resource that is waiting for the lock	String			
	Statement type	String			
	Query SQL.	String			
	Indicates whether the query is performed by the system.	Yes/No			
	Query plan	String			
SQL Probe Statistics	Query ID of a probe task	String	30s	Monitor cluster performance by collecting statistics on the duration of SQL probe execution, to detect sudden deterioration.	N/A
	Cluster ID	String			
	Cluster project ID	String			
	Task type of the probe SQL	String			
	Time when a probe SQL task is created	≥ 0			
	SQL execution time	≥ 0			
	probe_ID of the associated probe SQL	String			
	Probe	String			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Task status of the SQL probe. The options are as follows: <ul style="list-style-type: none"> • Running • Success • Fail 	String			
Table Dirty Page Rate Statistics	Database name	String	7200s	Gather information on dirty pages for cluster tables, as a high dirty page rate can cause a decline in table query and insertion performance.	50 %
	Schema name	String			
	Table name (full name)	String			
	Table owner	String			
	Table size (unit: byte)	≥ 0			
	Dirty page rate	≥ 0.0			
Table Skew Monitoring Statistics	Database name	String	7200s	Identify tables in the cluster with a skew rate exceeding 5%, as this can negatively affect query performance.	10 %
	Schema name	String			
	Table name (full name)	String			
	Table owner	String			
	Table size (unit: byte)	≥ 0			
	Skew rate	≥ 0.0			
Resource Pool Statistics	Load resource pool	String	120s	Monitor system resource usage and queuing status of the cluster resource pool by collecting relevant information.	N/A
	CPU quota of the resource pool	≥ 0			
	Memory quota of the resource pool	≥ 0			
	Disk quota of the resource pool	≥ 0			

Monitored Object	Metric	Value Range	Monitoring Interval (Raw Data)	Metric Usage	Recommended Threshold
	Maximum number of concurrent simple jobs allowed by the resource pool	≥ 0			
	Maximum number of concurrent queries allowed by the resource pool	≥ 0			
	CPU usage of the resource pool	≥ 0.0			
	Memory usage of the resource pool	≥ 0.0			
	Disk usage of the resource pool	≥ 0.0			
	Number of concurrent simple jobs in the resource pool	≥ 0			
	Current number of concurrent requests in the resource pool	≥ 0			
Resource Pool User Statistics	Load resource pool	String	30s	Gather data on users in the cluster resource pool and track their resource usage.	N/A
	CPU quota of a user.	≥ 0			
	Memory quota of a user	≥ 0			
	Specifies the disk quota of a user	≥ 0			
	CPU usage of a user	≥ 0.0			
	Memory usage of a user	≥ 0.0			
	Disk usage of a user	≥ 0.0			
	User ID	String			
	Username	String			

5.2.3 Cluster Overview

Cluster Overview

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**. The database monitoring page is displayed.
- Step 4** In the navigation pane on the left, click **Cluster Overview**.

On the page that is displayed, you can view the cluster status, real-time resource consumption, cluster resource consumption, and key database metrics.

----End

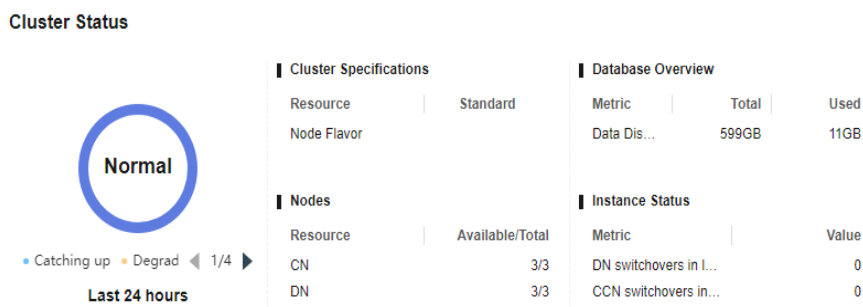
NOTE

Metrics can be collected and displayed on the cluster overview page only if their collection items are enabled. If a collection item is disabled, its metric will not be displayed, and a prompt will be displayed indicating this problem. In this case, you are advised to enable the collection item.

Cluster Status

In the **Cluster Status** area, you can view the statistics about the current cluster status and instance status, including cluster statistics in the last 24 hours, cluster specifications, available/total CNs and DNs, used/total disk capacity, the number of CCN switchovers in the last 24 hours, and the number of primary/standby DN switchovers in the last 24 hours.

Figure 5-9 Cluster Status



NOTE


The OBS usage details are displayed in DWS 3.0 cluster information.

Alarms

In the **Alarms** area, you can view all the uncleared alarms of the current cluster and the alarms generated in the last seven days. You can click **More** in the upper

right corner to view details about the existing cluster alarms. For details, see [Alarms](#).

Figure 5-10 Alarms



Cluster Resources

In the **Cluster Resources** area, you can view the resource usage of the current cluster, including the average CPU usage, disk I/O, disk usage, memory usage, and network I/O. You can click the metric of a resource to view its trend in the last 24 hours and the top five services that are occupying this resource. You can click **More** in the upper right corner of the area to go to the **Node Monitoring** page. Nodes are sorted by the metric value. For details, see [Node Monitoring](#).

Figure 5-11 Cluster Resources



Workloads

In the **Workloads** area, you can view the workload metrics of the current database, including TPS, QPS, stacked SQL queries, and resource pools. You can also click a workload metric to view its trend in the last 24 hours. The **SQL Stack Queries** metric depends on the real-time query monitoring function. If this function is disabled, no data will be displayed for the metrics.

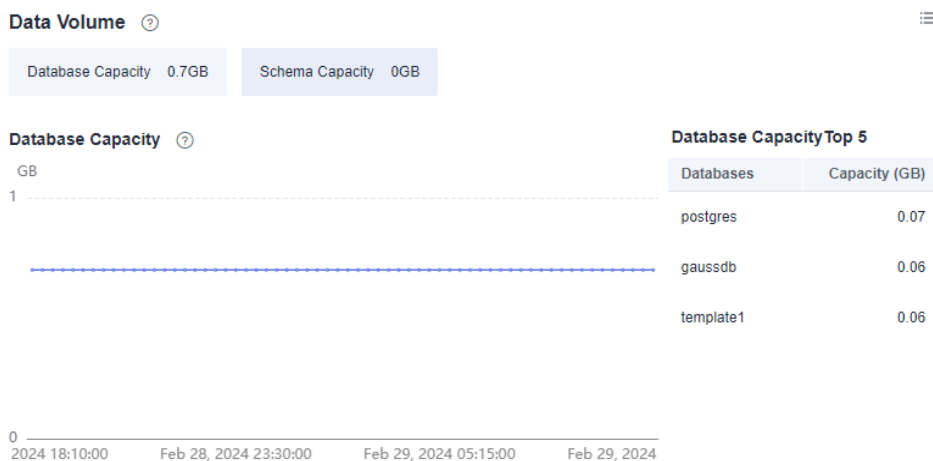
Figure 5-12 Workloads



Data Volume

In the **Data Volume** area, you can view the used capacity of the current database and schema. You can click a capacity metric to view the database or schema capacity trend in the last 24 hours and the top five databases or schemas ranked by capacity usage in the current cluster. You can click **More** in the upper right corner of the area to go to the **Database Monitoring** page. Databases are sorted by used capacity. For details, see [Database Monitoring](#).

Figure 5-13 Data Volume



NOTE

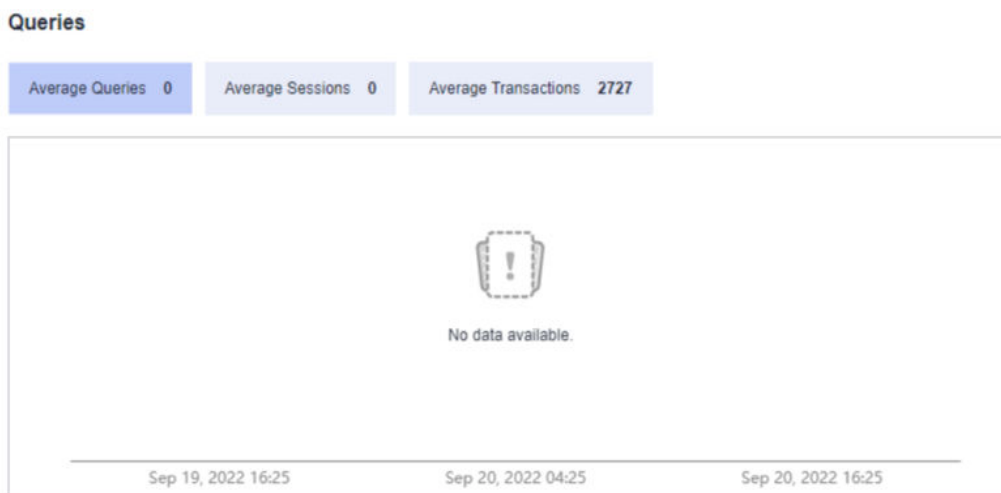
The database capacity data is collected once a day. Therefore, the data volume fluctuates greatly. To view real-time capacity monitoring information, choose **Node Monitoring > Disks**.

Queries

In the **Queries** area, you can check the average number of queries, sessions, and transactions. You can click a metric to view its trend in the last 24 hours. The **Average Queries** and **Average Sessions** metrics depend on the real-time query

monitoring function. If this function is disabled, no data will be displayed for the metrics.

Figure 5-14 Queries



5.2.4 Monitoring

5.2.4.1 Node Monitoring

Node Monitoring

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**. The database monitoring page is displayed.
- Step 4** In the navigation pane on the left, choose **Monitoring > Node Monitoring**.


On the page that is displayed, view the real-time consumption of nodes, memory, disks, disk I/O, and network I/O.


----End

Overview

On the **Overview** tab page, you can view the key resources of a specified node based on the node name, including:

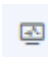
- Node Name
- CPU Usage (%)
- Memory Usage (%)
- Average Disk Usage (%)
- IP Address


- Disk I/O (KB/s)
- TCP Protocol Stack Retransmission Rate (%)
- Network I/O (KB/s)
- Status
- Monitoring: You can click  in the **Monitoring** column to view the overall performance metric topology of the node in the last 1 hour, last 3 hours, last 12 hours, last 24 hours, last 7 days, or last 15 days.



Disks

On the **Disks** tab page, view the real-time disk resource consumption of a node by node name and disk name, including:

- Node Name
- Disk Name
- Disk Type
 - System disk
 - Data disk
 - Log disk
- Disk Capacity (GB)
- Disk Usage (%)
- Disk Read Rate (KB/s)
- Disk Write Rate (KB/s)
- I/O Wait Time (await, ms)
- I/O Service Time (svctm, ms)
- IOPS
- Monitoring: You can click  in the **Monitoring** column to view the disk performance metric topology of the node in the last 1 hour, last 3 hours, last 12 hours, or last 24 hours.



 NOTE

The sum of the used disk space and available disk space is not equal to the total disk space. This is because a small amount of space is reserved in each default partition for system administrators to use. Even if common users have run out of space, system administrators can log in to the system and use their space required for solving problems.

Run the Linux **df** command to collect the disk capacity information, as shown in the following figure.

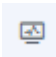
```
[Ruby@host-10-0-16-43 8_1_0]# df -x tmpfs -x devtmpfs
Filesystem      1K-blocks    Used  Available Use% Mounted on
/dev/sda4       569616888 5757444 540228616  2% /
/dev/sda2         999320  107584   822924  12% /boot
/dev/sda1        204580    8368   196212   5% /boot/efi
/dev/sdd       3513495364 390076 3513105288   1% /var/chroot/DWS/data1
/dev/sde       3513495364 274192 3513221172   1% /var/chroot/DWS/data2
/dev/sdb       3513495364  34224 3513461140   1% /var/chroot/DWS/data3
/dev/sdc       3513495364  34224 3513461140   1% /var/chroot/DWS/data4
[Ruby@host-10-0-16-43 8_1_0]#
```

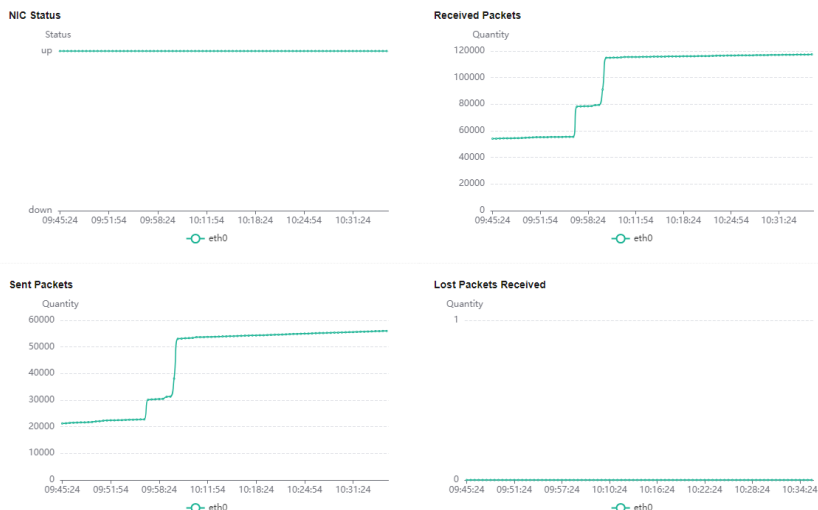
/dev/sda4: Used(5757444) + Available(540228616) != Total(569616888)

- **Filesystem:** path name of the device file corresponding to the file system. Generally, it is a hard disk partition.
- **1K-blocks:** number of data blocks (1024 bytes) in a partition.
- **Used:** number of data blocks used by the disk.
- **Available:** number of available data blocks on the disk.
- **Use%:** percentage of the space used by common users. Even if the space is used up, the partition still reserves the space for system administrators.
- **Mounted on:** mount point of the file system.

Network

On the **Network** tab page, view the real-time network resource consumption of a node by node name and NIC name, including:

- Node Name
- NIC Name
- NIC Status
- NIC Speed (Mbps)
- Received Packets
- Sent Packets
- Lost Packets Received
- Receive Rate (KB/s)
- Transmit Rate (KB/s)
- Monitoring: You can click  in the **Monitoring** column to view the network performance metric topology of the node in the last 1 hour, last 3 hours, last 12 hours, or last 24 hours.



5.2.4.2 Performance Monitoring

Performance Monitoring

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**.
- Step 4** In the navigation pane on the left, choose **Monitoring > Performance Monitoring**. The **Performance Monitoring** page displays the resource consumption trends of clusters, databases, and nodes.

You can select a time range and check the performance trend in this range.

- By default, the monitoring information of the last hour is displayed.
- You can view the monitoring information of the last seven days.

----End

Monitoring Panel

You can configure monitoring views by customizing monitoring panels. Monitoring panels are bound to users. After logging in to the system, you can view the user-defined monitoring panels.

- Creating a monitoring panel: You can click **Create Panel** to customize a monitoring panel.
- Modifying a monitoring panel: You can click **Modify** to change the name of a monitoring panel.
- Deleting a monitoring panel: You can click **Delete** to delete a monitoring panel. The default monitoring panel cannot be deleted.
- Sharing a monitoring panel: You can click **Share** to share a monitoring panel. The recipients can view the panel but cannot modify it.

Adding a Monitoring View

Currently, DMS provides monitoring views for clusters, databases, and nodes. You can click **Add View** to add a monitoring view as required. The monitoring metrics are as follows:

- Cluster: CPU Usage, Memory Usage, Disk Usage, Disk I/O, Network I/O, Status, Abnormal CNs, Read-only, Sessions, Queries, Deadlocks, Abnormal DNs, CPU Usage of DNs, TPS, and QPS
- Database: query waiting queue length, number of sessions, number of queries, number of inserted rows, number of updated rows, number of deleted rows, and capacity.
- Node: CPU usage, memory usage, average disk usage, disk I/O, TCP retransmission rate, network I/O, total disk space, disk usage, disk read rate, disk write rate, disk I/O wait time, disk I/O service time, disk I/O usage, NIC status, number of received packets, number of sent packets, number of lost packets received, receive rate, and transmit rate.

NOTE

- A maximum of 20 views can be added to each panel. Adding too many views will increase the number of page requests and the rendering time.
- A maximum of 20 monitored objects can be selected in the node dimension. This feature is supported only in 8.1.3.310 and later cluster versions.

Exporting Monitoring Data

Performance Monitoring supports data export. You can click **Export Data** to further process data. By default, data in all monitoring views on the current page is exported. The export time range is subject to the selected time range.

NOTE

Performance Monitoring allows data aggregation of different periods. You can aggregate raw data based on the corresponding sampling period to display indicator trends of a longer period.

5.2.4.3 Database Monitoring

Database Monitoring

Step 1 Log in to the GaussDB(DWS) console.

Step 2 On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.

Step 3 In the **Operation** column of the target cluster, click **Monitoring Panel**. The database monitoring page is displayed.

Step 4 In the navigation pane on the left, choose **Monitoring > Database Monitoring**.

The **Database Monitoring** page displays the real-time and historical resource consumption a database.


----End

Database Resource Consumption

You can select a database to view its resource usage. For details about the metrics, see [Monitoring Metrics](#), including:

- Database Name
- Usage (GB)
- Monitoring
- Users
- Applications
- Sessions
- Queries
- Inserted Rows
- Updated Rows
- Deleted Rows
- Deadlocks
- Temporary Files
- Temporary File Capacity

Database Trend Monitoring

In the **Monitoring** column of a database, click  to view the performance indicators of the database, including:

- Capacity
- Sessions
- Queries

5.2.4.4 Real-Time Queries

Going to the Real-time Query Page

Step 1 Log in to the GaussDB(DWS) console.

Step 2 On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.

Step 3 In the **Operation** column of the target cluster, click **Monitoring Panel**.

Step 4 In the navigation pane, choose **Monitoring > Queries**.

You can check the real-time information about all queries and sessions running in the cluster.

----End

NOTICE

- Real-time query is supported only in clusters of version 8.1.2 and later.
- To enable real-time query monitoring, choose **Settings > Monitoring**, click **Monitoring Collection**, and enable **Real-Time Query Monitoring**. For details, see [Monitoring Collection](#). Enabling real-time query may cause a large amount of data. Exercise caution when performing this operation.

Prerequisites


You need to set GUC parameters before viewing data on the monitoring page. If GUC parameters are not set, real-time or historical query may be unavailable. However, if this parameter is set, the cluster performance may deteriorate. Therefore, you need to balance the settings of related parameters. The following table describes recommended settings. For how to modify parameters, see [Modifying Database Parameters](#). [Setting GUC Parameters](#) provides parameter details.

Table 5-2 Recommended GUC parameter settings

GUC Parameter	CN Configuration	DN Configuration
max_active_statements	10	10
enable_resource_track	on	on
resource_track_level	query	query
resource_track_cost	0	0
resource_track_duration	10	10
enable_resource_record	on	on
session_statistics_memory	1000MB	1000MB

Querying Information

You can view the queries statistics, the number of sessions, average session duration (time of all session connections divided by the number of sessions), number of queries, average query duration, and average query waiting time.



Checking Live Sessions

On the **Sessions** tab, you can browse the real-time information about all running queries,

- Session ID
- Username
- Session duration
- Application name
- QueryBand
- Client IP address
- Connected CN
- Session status. It can be:
 - **idle**: The backend is waiting for new client commands.
 - **active**: The backend is executing queries.
 - **idle in transaction**: The backend is in a transaction, but there is no statement being executed in the transaction.
 - **idle in transaction (aborted)**: The backend is in a transaction, but there are statements failed in the transaction.
 - **fastpath function call**: The backend is executing a **fast-path** function.
- Start time
- Lock mode
- Lock holding status
- Locked object
- Query SQL
- Lock wait
- Current query duration
- Current query start time

 **NOTE**

- You can click a session ID to view the queries in the current session. For details, see [Viewing Historical Query Monitoring Details](#).
- To terminate a session, select the session, click **Terminate a Session**, and confirm your operation.
- If you want to terminate all idle sessions, click **Clear Idle Sessions**.
- The fine-grained permission control function is added. Only users with the operate permission are able to terminate sessions. For users with the read-only permission, the **Terminate a Session** button is grayed out.

Checking Real-time Queries

On the **Queries** tab, you can browse all the queries that are running in a specified time period, including:

- Query ID
- Username
- Application name
- Database name
- Resource pool
- Submission time

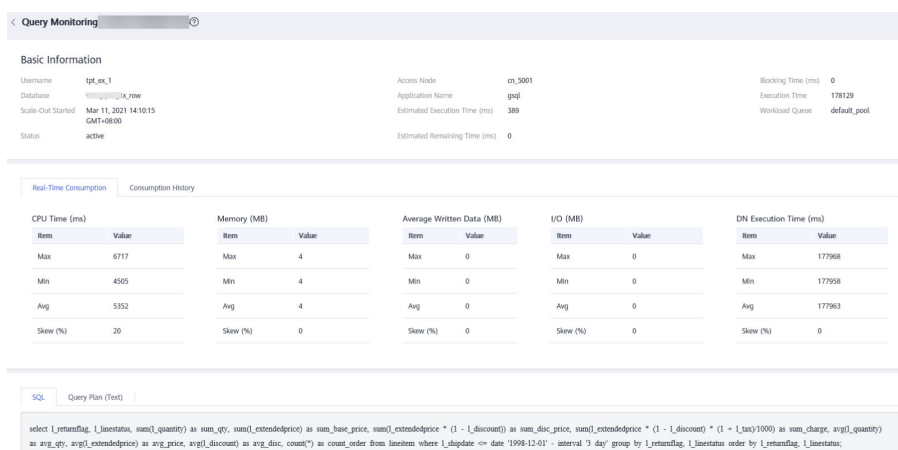
- Blocking time (ms)
- Execution time (ms)
- Statement
- Connected CN
- Client IP address
- Lane
- Query status. It can be:
 - **idle**: The backend is waiting for new client commands.
 - **active**: The backend is executing queries.
 - **idle in transaction**: The backend is in a transaction, but there is no statement being executed in the transaction.
 - **idle in transaction (aborted)**: The backend is in a transaction, but there are statements failed in the transaction.
 - **fastpath function call**: The backend is executing a **fast-path** function.
- Session ID
- Statement status

 **NOTE**

- You can click a query ID to view the monitoring details. However, details cannot be displayed for queries whose ID is **0**. Query **0** indicates that an exception occurs during the query.
- To terminate a query, select the query, click **Terminate Query**, and confirm your operation.
- The fine-grained permission control function is added. Only users with the operate permission are able to terminate queries. For users with the read-only permission, the **Terminate Query** button is grayed out.
- The fast and slow lanes are selected based on the cost in the execution plan. If the optimizer estimates that the memory usage of a statement is greater than 32 MB, the statement enters the slow lane. Otherwise, the statement enters the fast lane.

Viewing Real-time Query Monitoring Details

You can click a query ID to view the query details, including the basic information of query statements, real-time and historical resource consumption, SQL description, and query plan.



The screenshot displays the 'Query Monitoring' interface. It includes a 'Basic Information' section with fields for Username, Database, Scale-Out Started, Status, Access Node, Application Name, Estimated Execution Time, Estimated Remaining Time, Blocking Time, Execution Time, and Workload Queue. Below this is a 'Real-Time Consumption' table with columns for CPU Time (ms), Memory (MB), Average Written Data (MB), I/O (MB), and DN Execution Time (ms), each with sub-columns for Max, Min, Avg, and Skew (%). At the bottom, there is an 'SQL' section showing the query plan text.

CPU Time (ms)		Memory (MB)		Average Written Data (MB)		I/O (MB)		DN Execution Time (ms)	
Item	Value	Item	Value	Item	Value	Item	Value	Item	Value
Max	6717	Max	4	Max	0	Max	0	Max	177968
Min	4505	Min	4	Min	0	Min	0	Min	177958
Avg	5352	Avg	4	Avg	0	Avg	0	Avg	177963
Skew (%)	20	Skew (%)	0	Skew (%)	0	Skew (%)	0	Skew (%)	0

```

select l_returnflag, l_linestatus, sum(qty) as sum_qty, sum(l_extendedprice) as sum_base_price, sum(l_extendedprice * (1 - l_discount)) as sum_disc_price, sum(l_extendedprice * (1 - l_discount) * (1 + l_tax1000)) as sum_charge, avg(l_quantity) as avg_qty, avg(l_extendedprice) as avg_price, avg(l_discount) as avg_disc, count(*) as count_order from lineitem where l_shipdate <= date '1998-12-01' - interval '3 day' group by l_returnflag, l_linestatus
    
```

5.2.4.5 Historical Queries

Going to the Historical Query Page

Step 1 Log in to the GaussDB(DWS) console.

Step 2 On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.

Step 3 In the **Operation** column of the target cluster, click **Monitoring Panel**.

Step 4 In the navigation pane on the left, choose **Monitoring > History**.

All the historical queries in the current cluster will be displayed.

----End

NOTE

- Historical queries can be viewed only in clusters of version 8.1.2 and later.
- To enable historical query monitoring, choose **Settings > Monitoring**, click **Monitoring Collection**, and enable **Historical Query Monitoring**. For details, see [Monitoring Collection](#). Enabling history query may cause a large amount of data. Exercise caution when performing this operation.

Checking Historical Queries

In the **History** area, you can browse all historical query information based on the specified time period, including:

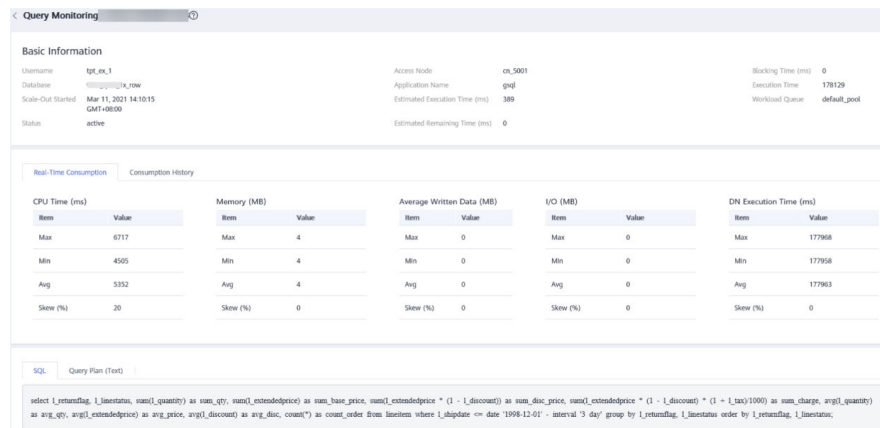
- Query ID
- Username
- Application name
- Database name
- Resource pool
- Submission time
- Blocking time (ms)
- Execution time (ms)
- CPU time (ms)
- CPU time skew (%)
- Statement
- Connected CN
- Client IP address
- Query status
- Completion time
- Estimated execution time (ms)
- Cancellation reason

NOTE

If you do not want to see historical system queries, you can toggle on **Hide System Queries**.

Viewing Historical Query Monitoring Details

You can click a historical query ID to view the query details, such as basic information about query statements, real-time resource consumption during execution, complete SQL statements, and query plans.



5.2.4.6 Instance Monitoring

Instance Monitoring

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**.
- Step 4** In the navigation pane on the left, choose **Monitoring > Instance Monitoring**.


On the **Instance Monitoring** page, you can view the real-time and historical information about detected slow instances.

----End

Slow Instance Detection

DMS can automatically configure and start the slow instance detection script on cluster CNs, periodically collect the cache table of the script, and report the detected slow instance data. You can view the number of slow instances detected within 24 hours and the distribution status in the time dimension on the GUI to quickly locate the slow nodes in the cluster and analyze the root causes.

The **Instance Monitoring** page consists of two parts. The upper part displays the time distribution chart of detected slow instances, that is, the number of slow instances detected in different detection periods. The lower part displays slow instance details. When you select any bar in the time distribution chart, details about the detection time, node name, instance name, and number of detections (within 24 hours) of slow instances are displayed.



NOTE

If the period of an instance exceeds 240 seconds, the instance is reported as a slow instance.

5.2.4.7 Resource Pool Monitoring

Accessing the Resource Monitoring Page

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**.
- Step 4** In the navigation pane on the left, choose **Monitoring > Resource Pool Monitoring**.

You can check the real-time statistics and resource consumption history about resource pools.

----End

Resource Pool

You can check user-defined resource pools, real-time and historical resource consumption, and the resource quotas of resource pools.

- **Resource Pool:** Resource pool name.
- **Monitoring:** You can click the monitoring icon to display the historical consumption trends of resources such as the CPU, memory, and disk.
- **CPU Usage (%):** real-time CPU usage of a resource pool
- **CPU Share (%):** CPU usage share of a resource pool
- **Real-Time Concurrent Short Queries:** number of concurrent simple queries in a resource pool. Concurrent simple queries are not controlled by the resource pool.
- **Concurrent Short Queries:** quota of simple concurrent queries in a resource pool
- **Real-Time Concurrent Queries:** number of concurrent complex queries in a resource pool. Concurrent complex queries are controlled by the resource pool.

- **Query Concurrency:** quota of complex concurrent queries in a resource pool
- **Storage (MB):** storage space of a resource pool
- **Disk Usage (%):** real-time disk usage of a resource pool
- **Memory Resource:** memory quota of a resource pool
- **Memory Usage (%):** percentage of used memory
- **Operation:** resource pool configurations.

User Resource Usage

Click the arrow next to a resource pool name to expand resource usage details.

- **User Name:** name of a user in the current resource pool
- **CPU Usage (%):** real-time CPU usage of a user
- **CPU Resource:** number of CPU cores used
- **Storage Resource (MB):** storage space used by a user
- **Disk Usage (%):** disks used by a user
- **Memory Resource (MB):** memory used by a user
- **Memory Usage (%):** percentage of memory used by a user

Figure 5-15 User Resource Usage

The screenshot shows a 'Resource Pool Monitoring' interface. It has tabs for 'Resource Pools', 'Waiting Queries', and 'Circuit Breaking Queries'. The 'Resource Pools' tab is active, displaying a table with columns: Resource Pools, Monitoring, CPU Usage, Disk Usage, Memory Usage, Real Time Concurrent Query Count, Real Time Concurrent Query Limit, and Operation. Below this, there is a 'User Resource Usage' section with a table showing columns: User Name, CPU Usage, Disk Usage, and Memory Usage. The data shows multiple users (usr_1 to usr_11) with 0% usage across all metrics.

Queries Waiting in a Resource Pool

You can view the queries waiting in a resource pool in real time to check workload status.

- **User:** user name of a query statement
- **Application:** application name of a query statement
- **Database:** name of the database to which a query statement is connected
- **Queuing Status:** queuing status of a query statement in a resource pool
- **Wait Time:** waiting time before a query statement is executed, in ms
- **Resource Pool:** resource pool that the query belongs to
- **Query Statement:** details of a query statement submitted by a user

Checking Circuit Breaking Queries

You can view the status of a triggered circuit breaking query in a resource pool.

- **Query ID:** ID of a circuit breaking query
- **Query Statement:** circuit breaking query statement
- **Blocking Time (ms):** blocking time of a circuit breaking statement, in ms
- **Execution Time (ms):** execution time of a circuit breaking statement, in ms
- **CPU Time (ms):** CPU time consumed by a circuit breaking statement, in ms
- **CPU Skew (%):** CPU skew of a circuit breaking statement on each DN
- **Exception Handling:** exception handling method of a circuit breaking statement
- **Status:** real-time status of a circuit breaking statement

5.2.5 Utilities

5.2.5.1 SQL Diagnosis

Prerequisites

To enable SQL diagnosis, enable monitoring on real-time and historical queries on the **Queries** and **History** tabs, respectively. For details, see [Monitoring Collection](#).

Viewing SQL Diagnosis

Step 1 Log in to the GaussDB(DWS) management console.

Step 2 On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.

Step 3 In the **Operation** column of the target cluster, click **Monitoring Panel**.

Step 4 In the navigation pane on the left, choose **Utilities > SQL Diagnosis**. The metrics include:

- Query ID
- Database
- Schema Name
- User Name
- Client
- Client IP Address
- Running Time (ms)
- CPU Time (ms)
- Scale-Out Started
- Completed
- Details

Step 5 On the **SQL Diagnosis** page, you can view the SQL diagnosis information. In the **Details** column of a specified query ID, click **View** to view the detailed SQL diagnosis result, including:

- Alarm Information
- SQL Statement

- Execution Plan

----End

Setting GUC Parameters

GUC parameters related to SQL diagnosis are as follows. For details, see "GUC Parameters" in the *Data Warehouse Service (DWS) Developer Guide*.

- **enable_resource_track**
 - Value range: boolean
 - Default value: **on**
 - Expected DMS value: **on** (for reference only)
 - Function: Specifies whether to enable the real-time resource monitoring function.

NOTICE

If this parameter is enabled without other GUC-related parameters correctly configured, real-time resource consumption cannot be recorded.

- **resource_track_cost**
 - Value range: an integer ranging from -1 to INT_MAX
 - Default value: **100000**
 - Expected DMS value: **0** (for reference only)
 - Function: Specifies the minimum execution cost of statement resource monitoring for the current session. This parameter is valid only when **enable_resource_track** is **on**.

NOTICE

If this parameter is set to a small value, more statements will be recorded, causing record expansion and affecting cluster performance.

- **resource_track_level**
 - Value range: enumerated type
 - Default value: **query**
 - Expected DMS value: **query** (for reference only)
 - Function: Specifies the resource monitoring level for the current session. This parameter is valid only when **enable_resource_track** is **on**.

NOTICE

If the resource monitoring is set to operator-level, performance will be greatly affected.

- **resource_track_duration**
 - Value range: an integer ranging from 0 to INT_MAX, in seconds
 - Default value: **60**.
 - Expected DMS value: **0** (for reference only)
 - Function: Specifies the minimum statement execution time that determines whether information about jobs of a statement recorded in the real-time view will be dumped to a historical view after the statement is executed. That is, only statements whose execution time exceeds the specified time are recorded in the historical view. This parameter is valid only when **enable_resource_track** is **on**.

NOTICE

If this parameter is set to a small value, the batch processing mechanism for dumping kernel statements becomes invalid, affecting the kernel performance.

- **topsql_retention_time**
 - Value range: an integer ranging from 0 to 3650, in days
 - Default value: **30**
 - Expected DMS value: **14** (for reference only)
 - Function: Specifies the aging time of **pgxc_wlm_session_info** data in the view.

NOTICE

If this parameter is set to **0**, data will not be aged, which will cause storage expansion.

- **enable_resource_record**
 - Value range: boolean
 - Default value: **off**
 - Expected DMS value: **on** (for reference only)
 - Function: Specifies whether to enable the archiving function for resource monitoring records. When this function is enabled, records in the history views (**GS_WLM_SESSION_HISTORY** and **GS_WLM_OPERATOR_HISTORY**) are archived to the info views (**GS_WLM_SESSION_INFO** and **GS_WLM_OPERATOR_INFO**) every 3 minutes. After the archiving, records in the history views are deleted.

NOTICE

When this parameter is enabled, you are advised to set **topsql_retention_time** properly to configure the aging time. Otherwise, data in the **GS_WLM_SESSION_INFO** or **GS_WLM_OPERATOR_INFO** table will expand.

5.2.5.2 SQL Probes

You can upload and verify SQL probes, execute probe tasks in one click, and periodically execute probe tasks. Alarms can be reported for timeout SQL probes. The procedure is as follows:

- [Adding a SQL Probe](#)
- [Enabling or Disabling a SQL Probe](#)
- [Modifying an SQL Probe](#)
- [Deleting a SQL Probe](#)
- [Executing a SQL Probe in One Click](#)

NOTE

- The SQL probe is supported only in 8.1.1.300 and later versions. To use it in earlier versions, contact technical support.
- Only **SELECT** statements can be used as SQL probes.
- Up to 20 SQL probes can be configured.
- To create an SQL probe, you must have the GaussDB(DWS) FullAccess permission.
- To enable the SQL probe function, choose **Monitoring Settings > Monitoring Collection** and enable the **SQL Probe** metric. For details, see [Monitoring Collection](#). The default collection frequency is 30s.

Adding a SQL Probe

Step 1 Log in to the GaussDB(DWS) console.

Step 2 On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.

Step 3 In the **Operation** column of the target cluster, click **Monitoring Panel**.

Step 4 In the navigation pane, choose **Utilities > SQL Probes**. Click **Add SQL Probe**.

Step 5 Configure SQL probe parameters.

- **Probe Name:** Name of a probe.
- **Database:** Database where the probe SQL statement is to be executed.
- **SQL Statement:** Probe SQL statement to be executed. (Only **SELECT** statements are allowed).
- **Probe Threshold (ms):** Timeout threshold of probe SQL execution.
- **Description:** Probe SQL statement description.

Step 6 Confirm the SQL probe information and click **Confirm**.

----End

Enabling or Disabling a SQL Probe

Step 1 Log in to the GaussDB(DWS) console.

Step 2 On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.

Step 3 In the **Operation** column of the cluster, choose **Monitoring Panel**. The database monitoring page is displayed.

- Step 4** In the navigation pane on the left, choose **Utilities > SQL Probes**.
 - Step 5** In the probe list, click **Enable** (or **Disable**) in the **Operation** column of a probe.
 - Step 6** Confirm the information and click **OK**.
- End

Modifying an SQL Probe

- Step 1** Log in to the GaussDB(DWS) console.
 - Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
 - Step 3** In the **Operation** column of the cluster, choose **Monitoring Panel**. The database monitoring page is displayed.
 - Step 4** In the navigation pane on the left, choose **Utilities > SQL Probes**.
 - Step 5** In the probe list, click **Modify** in the **Operation** column of a probe.
 - Step 6** On the **Modify Probe** page, modify the SQL probe parameters as required and click **OK**.
- End

Deleting a SQL Probe

- Step 1** Log in to the GaussDB(DWS) console.
 - Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
 - Step 3** In the **Operation** column of the cluster, choose **Monitoring Panel**. The database monitoring page is displayed.
 - Step 4** In the navigation pane on the left, choose **Utilities > SQL Probes**.
 - Step 5** In the probe list, click **Delete** in the **Operation** column of a probe.
 - Step 6** Confirm the information and click **OK**.
- End

Executing a SQL Probe in One Click

- Step 1** Log in to the GaussDB(DWS) console.
 - Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
 - Step 3** In the **Operation** column of the cluster, choose **Monitoring Panel**. The database monitoring page is displayed.
 - Step 4** In the navigation pane on the left, choose **Utilities > SQL Probes**.
 - Step 5** In the probe list, select a probe and click **Run**. The system will execute the selected probe and update information about the probe.
 - Step 6** Confirm the information and click **OK**.
- End

5.2.5.3 Table Diagnosis

GaussDB(DWS) provides statistics and diagnostic tools for you to learn table status, including:

- **Skew Rate:** monitors and analyzes data table statistics in the cluster, and displays information about the 50 largest tables whose skew rate is higher than 5%.
- **Dirty Page Rate:** monitors and analyzes data table statistics in the cluster, and displays information about the 50 largest tables whose skew rate is higher than 50%.
- **DDL Audit:** DDL review is a type of SQL review. To prevent improper DDL design from affecting services, this tool checks whether DDL metadata is standard, detecting potential table definition problems in advance. The check result can also be used as a reference for locating performance issues.

NOTE

- Only 8.1.1.x and later versions support the table skew rate and dirty page rate features. For earlier versions, contact technical support.
- Only 8.1.1.300 and later versions support the DDL review feature. For earlier versions, contact technical support.
- The data collection period of the table skew and dirty page checks can be configured on the [Monitoring Collection](#) page. Frequent data collection may affect cluster performance. Set a proper period based on your cluster workloads.

Skew Rate

Context

Improper distribution columns can cause severe skew during operator computing or data spill to disk. The workloads will be unevenly distributed on DN, resulting in high disk usage on a single DN and affecting performance. You can query your table size and skew rate, and change the distribution columns of tables with severe skew. In cluster versions 8.1.0 and later, you can use the syntax [ALTER TABLE](#). In other cluster versions, perform the operations described in [How Do I Change Distribution Columns?](#)

Procedure

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**.
- Step 4** In the navigation tree on the left, choose **Utilities > Table Diagnosis** and click the **Skew Rate** tab. The tables that meet the statistics collection conditions in the cluster are displayed.

----End

Dirty Page Rate

Context

DML operations on tables may generate dirty data, which unnecessarily occupies cluster storage. You can query the dirty page rate of tables, and optimize large

tables and tables with high dirty page rate. For details, see [Solution to High Disk Usage and Cluster Read-Only](#).

Procedure

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**.
- Step 4** In the navigation tree on the left, choose **Utilities > Table Diagnosis** and click the **Dirty Page Rate** tab. The tables that meet the statistics collection conditions in the cluster are displayed.

----End

DDL Audit

Viewing and Exporting DDL Audit Results

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**.
- Step 4** In the navigation tree on the left, choose **Utilities > Table Diagnosis**, and click the **DDL Audit** tab. The audit results are displayed.

NOTE

The selected audit items are displayed on the **DDL Audit** tab by default. You can configure the audit items on the [Monitoring Collection](#) tab. For more information, see [Table 5-3](#).

Table 5-3 Audit items

Item	Description
Number of Distribution Keys (disKeyCount)	<p>If there is no data skew, use no more than four distribution keys.</p> <p>Generally, if you use many distribution keys, data can be evenly distributed in a cluster, thus avoid data skew. However, if too many distribution keys are used, the storage performance and joint query performance may deteriorate. You are advised to configure no more than four distribution keys.</p> <ul style="list-style-type: none">• Storage performance issue: When data is added, the hash function calculates the result of each distribution column, aggregates the results, and then determine where to distribute data. A large number of distribution keys require time-consuming, complex calculation.• Union query performance issue: During multi-table join query, if all the columns of the distribution key are involved in the join condition, data does not need to be redistributed in the execution plan. If a large number of distribution keys are used, some of them may not be the columns involved in the join condition, and data redistribution may occur, which consumes many resources and takes long.
Number of Index Columns/PCKs (indexKeyOrPckCount)	<p>It is recommended that the number of partial cluster keys (PCKs)/columns of an index be less than or equal to 4.</p> <ul style="list-style-type: none">• A large number of index columns require many resources to maintain index data, and are likely to contain duplicate indexes.• While column-store data is imported, PCK columns are compared and calculated to determine CU division. A large number of PCKs will consume many resources and much time, affecting performance. To efficiently filter CUs in a query, the prefixes of the columns involved in the query conditions must be PCK columns. (For example, if the PCK columns are a, b, and c, the query criteria must be a>? and b>? and c>?.) Otherwise, all the CUs must be traversed, and data clustering does not contribute to query acceleration.

Item	Description
Invalid PCKs (invalidPck)	<p>Do not create invalid PCK columns.</p> <p>In 8.1.1 and later versions, the cluster can filter and compare data of the char, int8, int2, int4, text, bpchar, varchar, date, time, timestamp and timestampz types. If a column of an unsupported data type is used as a PCK, the column is an invalid PCK column. It does not take effect during CU filtering and will consume resources for its maintenance.</p>
numeric Data Usage (validityOfNumeric)	<p>When numeric data types are used, use integers if possible. If the precision requirement is not high, use the float fixed-length data type. The float fixed-length data type has better computing performance than the numeric variable-length data type.</p> <p>That is, if the numeric type is used, you are advised to specify the scale and precision within 38 bits. When the numeric type is used for calculation, the underlying layer attempts to convert the calculation to the calculation between int and bigint to improve the calculation efficiency. That is, the large integer optimization of the data type is used.</p> <p>In 8.1.1 and later versions, if no precision is specified, a maximum of 131,072 digits can be placed before the decimal point and a maximum of 16,383 digits can be placed after the decimal point. That is, the maximum scale and precision are used. In this case, large integer optimization cannot be performed during calculation, and the calculation efficiency decreases accordingly.</p>
Index Column Width (widthOfIndexKey)	<p>Generally, wide index columns are character string columns, which do not involve compare operations and will lead to large indexes that consume unnecessary space. Specify a value smaller than 64 bytes.</p>
Replication Table Size (sizeOfCopyTable)	<p>Tables that occupy more storage space than the threshold (100 MB) on a single DN will be identified. For such tables, you are advised to use common associated columns as distribution keys (generally with one primary key).</p> <p>The cluster supports replication tables. A replication table maintains a full copy of data on each node and is mainly used to store data of enumerated types. If a table contains too much data, it will occupy a large amount of space. In addition, in a union query, the node traverses all table data, which may take a longer time than the union query after the table type is changed to distribution table. (Although data may be redistributed in the distribution table, the amount of data traversed by each node decreases.)</p>

Item	Description
Skew Detection for Single-Distribution-Key Tables (recognitionOfDataSkew)	Data skew of single-distribution-key tables is detected by statistics. This audit applies only to tables with one distribution key.
Distribution Key Usage (validityOfDiskey)	In a cluster, you are not advised to use a column of the Boolean or date type as a distribution column, because it may cause data skew.
Number of Cached Sequence Values (cacheSizeOfSequence)	Specify a number greater than 100. If a table column uses sequences, its next_value is obtained from the cached value in the local node. If cached sequence values are used up, a request will be sent asking GTM to obtain the value again. If a large amount of data is added but only a few values are cached, GTM will receive many requests, and may get overloaded and even break down. To avoid this problem, you are advised to set the cache value to a value greater than 100 when creating a sequence.
Optimizable Indexes (optimizableIndexKey)	Scenarios where indexes can be optimized: <ul style="list-style-type: none">• The index column of an index is the first <i>N</i> columns of another index.• The index columns of two indexes are the same, but the orders are different.

Step 5 If the review result of an item is **Failed**, click **View** to go to the details page.

Step 6 Click **Export** in the upper left corner to export the audit result.

----End

Manually Auditing DDL Items

Step 1 Log in to the GaussDB(DWS) console.

Step 2 On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.

Step 3 In the **Operation** column of the target cluster, choose **Monitoring Panel**. The database monitoring page is displayed.

Step 4 In the navigation tree on the left, choose **Utilities > Table Diagnosis**, and click the **DDL Audit** tab. On the page that is displayed, select the items to be audited and click **One-Click Audit**.

----End

5.2.5.4 Intelligent Distribution Column Recommendation

Context

Intelligent distribution column recommendation is crucial in distributed database systems. It optimizes data storage and query performance by analyzing past query

logs and suggesting the best distribution column for the data table. This results in efficient data sharding and storage, ultimately improving query performance.

 **NOTE**

If the current console does not support this feature, contact technical support.

Prerequisites

To activate this feature, enable **Historical Query Monitoring** on the **Monitoring Settings > Monitoring Collection** page. Once enabled, the system will generate historical query logs after a period of query operations. For details, see [Monitoring Collection](#).

Creating an Inference Task

- Step 1** Log in to the GaussDB(DWS) console.
 - Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
 - Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**.
 - Step 4** In the navigation tree, choose **Tools > Intelligent Distribution Column Recommendation**.
 - Step 5** To receive column distribution recommendations, click **Create Inference Task** and choose the cluster, database, and time segment you want to analyze.
 - Step 6** Click **OK**.
- End

Viewing the Inference Result

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**.
- Step 4** In the navigation tree, choose **Tools > Intelligent Distribution Column Recommendation**.
- Step 5** View the recommendation job list. The following table describes the related parameters.

Table 5-4 Parameters in the intelligent distribution column recommendation job list

Parameter	Description
Task ID	ID generated after a column distribution recommendation job is created.
Database Name	Name of the database.

Parameter	Description
Log Start Time	Start time for historical query logs used in data analysis. The default value is one week prior to the current time.
Log End Time	End time for historical query logs used in data analysis. The default value is the current time.
Status	Real-time status of a recommendation job based on the current column distribution recommendation job. <ul style="list-style-type: none">● Created: The creation is successful.● Entities extracted: The log data is being processed.● Plans processed: The recommendation result is being generated.● Finish: The recommendation job has been completed.● Failure: The recommendation job fails to be executed.
Operation	<ul style="list-style-type: none">● Click Show Inference Result to view the column distribution recommendation.● Click Delete to delete the record.

----End

5.2.6 Workload Analysis

5.2.6.1 Workload Analysis Overview

The workload analysis tool of GaussDB(DWS) collects and analyzes database performance data. You can create workload snapshots to record cluster workload data in a specified period. A workload diagnosis report can be generated based on two workload information snapshots within a certain time segment. Workload Diagnosis Report (WDR) provides performance data in a specified period and presents the data on HTML web pages. It helps you detect exceptions, diagnose problems, and optimize performance. It is a powerful tool for database performance tuning.

NOTE

- The WDR function is available only in 8.1.1.300 and later cluster versions.
- Workload diagnosis reports can be stored only in OBS.

5.2.6.2 Workload Snapshots

You can check the basic information about the cluster workload snapshots, manually create a snapshot, and configure snapshot parameters.

Checking Workload Snapshots

Step 1 Log in to the GaussDB(DWS) management console.

- Step 2** On the **Cluster > Dedicated Cluster** page, locate the cluster for which you want to perform workload analysis.
 - Step 3** In the **Operation** column of the cluster, choose **Monitoring Panel**. The database monitoring page is displayed.
 - Step 4** In the navigation pane, choose **Workload Analysis > Workload Snapshot**. Workload snapshots will be displayed.
- End

Creating a Workload Snapshot

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** On the **Cluster > Dedicated Cluster** page, locate the cluster for which you want to perform workload analysis.
- Step 3** In the **Operation** column of the cluster, choose **Monitoring Panel**. The database monitoring page is displayed.
- Step 4** In the navigation pane, choose **Workload Analysis > Workload Snapshot**. Workload snapshots will be displayed.
- Step 5** Click **Create Snapshot**. Enter a snapshot name and click **OK**.

 **NOTE**

Before creating a workload snapshot, ensure that the performance view snapshot parameter is enabled. For details, see [Configuring Workload Snapshot Parameters](#).

----End

Configuring Workload Snapshot Parameters

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster for which you want to perform workload analysis.
- Step 3** In the **Operation** column of the cluster, choose **Monitoring Panel**. The database monitoring page is displayed.
- Step 4** In the navigation pane, choose **Workload Analysis > Workload Snapshot**. Workload snapshots will be displayed.
- Step 5** Click **Configure Snapshot** in the upper right corner. In the dialog box that is displayed, check or modify GUC parameters. For details, see [Table 5-5](#).

----End

Table 5-5 Workload snapshot parameters

Name	Default Value	Description
Performance view snapshot (enable_wdr_snapshot)	off	Whether to enable the performance view snapshot function. If this function is enabled, GaussDB(DWS) will periodically create snapshots for certain system performance views and save them to disk. You can also manually create snapshots.
Resource monitoring (enable_resource_track)	on	Whether to enable the resource monitoring function. Resource statistics parameters are valid only if this parameter is enabled.
Logical memory management module (enable_memory_limit)	on	Whether to enable the logical memory management module.
Wait event statistics (enable_track_wait_event)	off	Whether to collect statistics on wait events, including the number of occurrences, number of failures, duration, maximum waiting time, minimum waiting time, and average waiting time.
I/O call time series statistics (track_io_timing)	off	Whether to collect time series statistics on database I/O calls. If this function is enabled, the collector will periodically query the OS time, which may cause heavy overhead on certain platforms.
SQL count (track_sql_count)	The default value is off for versions earlier than 8.1.3 and on for 8.1.3 and later versions.	Whether to collect statistics on the number of the SELECT , INSERT , UPDATE , DELETE , and MERGE INTO statements that are being executed in each session, the response time of the SELECT , INSERT , UPDATE , and DELETE statements, and the number of DDL, DML, and DCL statements. This parameter takes effect only if track_activities is set to on .
Session command statistics (track_activities)	on	Whether to collect statistics on the commands that are being executed in each session.
Unique SQL statistics (instr_unique_sql_count)	0	Whether to collect unique SQL statements and how many statements can be collected.

Name	Default Value	Description
Snapshot creation interval (<code>wdr_snapshot_interval</code>)	60	Interval for automatically creating performance view snapshots. It must be longer than the time needed to create a snapshot. The unit is minute.
Maximum snapshot retention period (<code>wdr_snapshot_retention_days</code>)	8	Maximum retention period of performance snapshots. A large value will require a lot of disk space. The unit is day.

5.2.6.3 Workload Reports

You can create, download, and delete work diagnosis reports, and check historical workload diagnosis reports.

 **NOTE**

To create a workload report, obtain the required OBS bucket permissions first.

Checking Workload Reports

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster for which you want to perform workload analysis.
- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**. The database monitoring page is displayed.
- Step 4** In the navigation pane, choose **Workload Analysis > Workload Reports**. Workload reports will be displayed.

----End

Generating a Workload Report

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster for which you want to perform workload analysis.
- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**. The database monitoring page is displayed.
- Step 4** In the navigation pane, choose **Workload Analysis > Workload Reports**.
- Step 5** Click **Generate Report**. In the displayed dialog box, configure the following parameters and click **OK**:

Table 5-6 Parameters for generating a report

Parameter	Description	Example Value
Report Name	User-defined. Ensure that the name is unique and contains a maximum of 100 characters, including digits, letters, and underscores (_).	test_show
Object Model	The object types are as follows: <ul style="list-style-type: none"> • node: The performance data of a specified node will be provided. • cluster: The performance data of the entire cluster will be provided. 	node
Node Name	User defined	dn_6005_l006
Content Type	The options are as follows: <ul style="list-style-type: none"> • summary: A report contains only brief analysis and calculation results. • detail: A report contains only detailed metric data. • all: A report contains content of both the summary and detail reports. 	all
Starting Snapshot	User defined NOTE The time of the starting snapshot start must be earlier than that of the ending snapshot.	-
Ending Snapshot	User defined	-
OBS Bucket	Bucket name, which is used to store reports.	test123

Parameter	Description	Example Value
OBS Path	Storage directory, which can be customized. Multi-level directories can be separated by slashes (/) and cannot start with slashes (/). Up to 50 characters are allowed.	wdr

----End

Downloading Workload Reports in Batches

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster for which you want to perform workload analysis.
- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**. The database monitoring page is displayed.
- Step 4** In the navigation pane, choose **Workload Analysis > Workload Reports**.
- Step 5** Select reports and click **Download**.

 **NOTE**

Up to 10 report records can be downloaded at a time.

----End

Deleting Workload Reports in Batches

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster for which you want to perform workload analysis.
- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**. The database monitoring page is displayed.
- Step 4** In the navigation pane, choose **Workload Analysis > Workload Reports**.
- Step 5** Select reports and click **Delete**.

----End

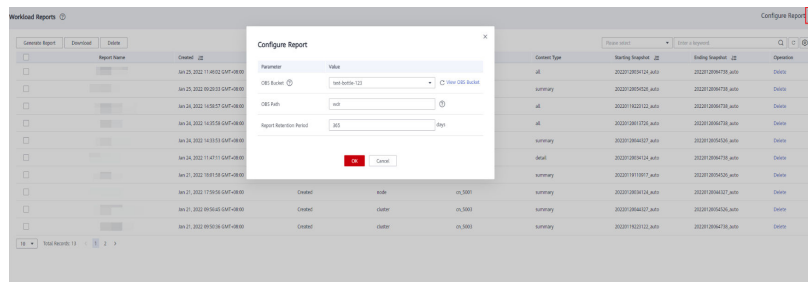
Deleting a Workload Diagnosis Report

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster for which you want to perform workload analysis.

- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**. The database monitoring page is displayed.
- Step 4** In the navigation pane, choose **Workload Analysis > Workload Reports**.
- Step 5** Click **Delete** in the **Operation** column of a report to delete the report record and file.
- End

Configuring Workload Report Parameters

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster for which you want to perform workload analysis.
- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**. The database monitoring page is displayed.
- Step 4** In the navigation pane, choose **Workload Analysis > Workload Reports**.
- Step 5** Click **Configure Report** in the upper right corner. In the displayed dialog box, set the report retention period and OBS parameters.



----End

5.2.7 Settings

The **Monitoring** page displays the collection period of monitoring metrics.

NOTE

- The cluster monitoring function is enabled by default.
- Disable the function if the cluster is being recovered. Enable the function when the fault is rectified.
- When a node in the cluster is powered off or the management IP address of the cluster is unavailable, the cluster monitoring switch and the button for configuring cluster indicator collection are unavailable.

Monitoring Collection

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
- Step 3** In the **Operation** column of the target cluster, click **Monitoring Panel**.

Step 4 In the navigation pane on the left, choose **Monitoring Settings > Monitoring Collection**. You can reconfigure the collection frequency or disable the collection of the monitoring item.

Name	Description	Collection Frequency (s)	Operation
Circuit Breaking Queries	Collect the real-time status of queries that trigger the circuit breaking rule.	30	<input checked="" type="checkbox"/>
Cluster Host Status	Cluster host status indicator collection	60	<input checked="" type="checkbox"/>
Cluster Instance Status	Cluster instance status indicator collection	60	<input checked="" type="checkbox"/>
Slow Instance Detection	Locate and collect the status of slow instances in a cluster.	240	<input checked="" type="checkbox"/>
Cluster Status	Cluster status indicator collection	30	<input checked="" type="checkbox"/>
CPU Status	CPU status indicator collection	30	<input checked="" type="checkbox"/>
Database Active Status	Database active status indicator collection	30	<input checked="" type="checkbox"/>
Database Capacity	Database capacity indicator collection	60,400	<input checked="" type="checkbox"/>
Database Status	Database status indicator collection	60	<input checked="" type="checkbox"/>
DDL Audit	Collection frequency of DDL audit data and related parameters	604,800	<input type="checkbox"/>

----End

5.2.8 Checking Task Details

On the task details page, you can view the status of tasks, such as enabling, disabling, resetting, and modifying cluster monitoring collection items; one-click DDL review; load snapshot generation; load diagnosis report generation; session termination; query termination; and the addition, modification, deletion, and one-click execution of probes.

NOTE

Only 8.1.3.110 and later cluster versions support the task details page.

Prerequisites

Tasks executed by users are related to SQL probes, load analysis, DDL one-click review, or monitoring collection items.

Procedure

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to be monitored.
- Step 3** In the **Operation** column of the cluster, click **Monitoring Panel**. The database overview page is displayed.
- Step 4** In the navigation pane on the left, choose **Tasks** to view the execution details of the commands delivered by the cluster. Task information includes the task name, task execution result, task command, start time, and end time.

Task Name	Execution Result	Command	Started	Ended
No data available.				

----End

5.2.9 Typical Scenarios

5.2.9.1 SQL Diagnosis

Symptom

The execution of SQL statements takes a long time, resulting in great resource consumption.

Troubleshooting Process

If the execution efficiency of SQL statements is low, optimization suggestions are provided after the kernel executes the SQL statements. You can query the execution history to retrieve optimization suggestions and further optimize SQL statements to improve query efficiency.

Troubleshooting Procedure

- Step 1** On the **SQL Diagnosing** page, select a time period that does not seem right.
 - Step 2** Search for SQL statements based on indicators such as the start time, end time, and running duration of the statement.
 - Step 3** Click **Details** to view SQL optimization suggestions.
 - Step 4** Optimize the SQL statement based on suggestions.
- End

5.3 Monitoring Clusters Using Cloud Eye

Function

This section describes how to check cluster metrics on Cloud Eye. By monitoring cluster running metrics, you can identify the time when the database cluster is abnormal and analyze potential activity problems based on the database logs, improving database performance. This section describes the metrics that can be monitored by Cloud Eye as well as their namespaces and dimensions. You can use the management console or **APIs** provided by Cloud Eye to query the monitoring metrics and alarms generated by GaussDB(DWS).

Namespace

SYS.DWS

Cluster Monitoring Metrics

With the GaussDB(DWS) monitoring metrics provided by Cloud Eye, you can obtain information about the cluster running status and performance. This information will provide a better understanding of the node-level information.

Table 5-7 describes GaussDB(DWS) monitoring metrics.

Table 5-7 GaussDB(DWS) monitoring metrics

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
dws001_shared_buffer_hit_ratio	Cache Hit Ratio	Ratio of requested data that already exists in the cache. It is the ratio of the amount of data that already exists in the cache to the total amount of requested data. A higher cache hit ratio means higher cache usage of the system, fewer times that data needs to be read from the disk or network, and faster system response speed. Unit: Percent	0% to 100%	Data warehouse cluster	4 minutes

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
dws002_in_memory_sort_ratio	In-memory Sort Ratio	Ratio of the extra memory space used by the sorting algorithm to the memory space occupied by the sorted data. In a merge sort, for example, the size of the merge buffer is often proportional to the size of the sorted data, so the in-memory ratio is usually between 10% and 50%. Unit: Percent	0% to 100%	Data warehouse cluster	4 minutes
dws003_physical_reads	File Reads	Total number of database file reads	> 0	Data warehouse cluster	4 minutes
dws004_physical_writes	File Writes	Total number of database file writes	> 0	Data warehouse cluster	4 minutes
dws005_physical_reads_per_second	File Reads per Second	Number of database file reads per second	≥ 0	Data warehouse cluster	4 minutes
dws006_physical_writes_per_second	File Writes per Second	Number of database file writes per second	≥ 0	Data warehouse cluster	4 minutes
dws007_db_size	Data Volume	Total size of data in the database, in MB	≥ 0 MB	Data warehouse cluster	4 minutes

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
dws008_active_sql_count	Active SQL Count	Number of active SQLs in the database	≥ 0	Data warehouse cluster	4 minutes
dws009_session_count	Session Count	Number of sessions that access the database	≥ 0	Data warehouse cluster	4 minutes
dws010_cpu_usage	CPU Usage	CPU usage of each node in a cluster, in percentage	0% to 100%	Data warehouse node	1 minute
dws011_mem_usage	Memory Usage	<p>Memory usage of each node in a cluster, in percentage</p> <p>NOTE After the console is upgraded to 8.3.0.202, the memory usage includes the memory occupied by the cache. Therefore, the value of this metric increases compared with that before the upgrade.</p>	0% to 100%	Data warehouse node	1 minute
dws012_iops	IOPS	Number of I/O requests processed by each node in the cluster per second	≥ 0	Data warehouse node	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
dws013_bytes_in	Network Input Throughput	Data input to each node in the cluster per second over the network Unit: byte/s	≥ 0 bytes/s	Data warehouse node	1 minute
dws014_bytes_out	Network Output Throughput	Data sent to the network per second from each node in the cluster Unit: byte/s	≥ 0 bytes/s	Data warehouse node	1 minute
dws015_disk_usage	Disk Usage	Disk usage of each node in a cluster, in percentage	0% to 100%	Data warehouse node	1 minute
dws016_disk_total_size	Total Disk Size	Total disk space of each node in the cluster Unit: GB	100 to 2000 GB	Data warehouse node	1 minute
dws017_disk_used_size	Used Disk Space	Used disk space of each node in the cluster Unit: GB	0 to 3600 GB	Data warehouse node	1 minute
dws018_disk_read_throughput	Disk Read Throughput	Data volume read from each disk in the cluster per second Unit: byte/s	≥ 0 bytes/s	Data warehouse node	1 minute
dws019_disk_write_throughput	Disk Write Throughput	Data volume written to each disk in the cluster per second Unit: byte/s	≥ 0 bytes/s	Data warehouse node	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
dws020_avg_disk_sec_per_read	Average Time per Disk Read	Average time used each time when a disk reads data Unit: second	> 0s	Data warehouse node	1 minute
dws021_avg_disk_sec_per_write	Average Time per Disk Write	Average time used each time when data is written to a disk Unit: second	> 0s	Data warehouse node	1 minute
dws022_avg_disk_queue_length	Average Disk Queue Length	Average I/O queue length of a disk	≥ 0	Data warehouse node	1 minute
dws_024_dn_diskio_util	DN I/O usage	Average disk I/O usage of DNs in a cluster	0% to 100%	Data warehouse instance	1 minute

Dimensions


Key	Value
datastore_id	Data warehouse cluster ID
dws_instance_id	Data warehouse node ID

Cluster and Node Monitoring Information

Step 1 Log in to the GaussDB(DWS) console and choose **Clusters > Dedicated Clusters**.

Step 2 View the cluster information. In the cluster list, click **View Metric** in the **Operation** column where a specific cluster resides. The Cloud Eye management console is displayed. By default, the cluster monitoring information on the Cloud Eye management console is displayed.

Additionally, you can specify a specific monitoring metric and the time range to view the performance curve.

Step 3 View the node information. Click  to return to the Cloud Eye management console. On the **Data Warehouse Nodes** tab page in the right pane, you can view metrics of each node in the cluster.

Additionally, you can specify a specific monitoring metric and the time range to view the performance curve.

Cloud Eye also supports the ability to compare the monitoring metrics of multiple nodes. For details, see [Comparing the Monitoring Metrics of Multiple Nodes](#).

----End

Comparing the Monitoring Metrics of Multiple Nodes

Step 1 In the navigation pane of the Cloud Eye management console, choose **Dashboards > My Dashboards**. Click the name of the dashboard for which you want to add a graph. On the **My Dashboards** page that is displayed, click **Add Graph**.

Step 2 On the **Add Graph** page, you can select **Line Chart** or **Bar Chart** to display the graph. After confirming that the information is correct, click **OK**.

For example, select **Line Chart** and **One View for Multiple Metrics** to compare the CPU usage of three GaussDB(DWS) nodes. The following table describes the parameters.

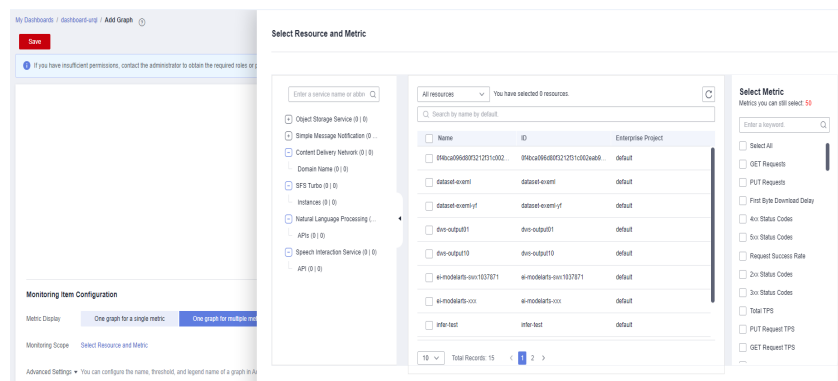


Table 5-8 Configuration example

Parameter	Example Value
Resource Type	DWS
Dimension	Data Warehouse Node
Monitored Object	dws-demo-dws-cn-cn-2-1 dws-demo-dws-cn-cn-1-1 dws-demo-dws-dn-1-1
Metric	CPU Usage

Step 3 Click **OK**.

On the selected **My Dashboards** page, you can view the metric trend on the newly added monitoring graph. You can click the zoom in button to zoom in and view detailed metric comparison data.

----End

Creating Alarm Rules

GaussDB(DWS) enables you to customize alarm rules for monitoring specific objects and notification policies, ensuring you stay informed about its running status in a timely manner.

A GaussDB(DWS) alarm rule includes the alarm rule name, monitored object, metric, threshold, monitoring interval, and whether to send a notification. This section describes how to set GaussDB(DWS) alarm rules.

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane, choose **Clusters > Dedicated Clusters**.
- Step 3** Locate the row containing the target cluster, click **View Metric** in the **Operation** column to enter the Cloud Eye management console and view the GaussDB(DWS) monitoring information.

The status of the target cluster must be **Available**. Otherwise, you cannot create alarm rules.

- Step 4** In the left navigation pane of the Cloud Eye management console, choose **Alarm Management > Alarm Rules**.
- Step 5** On the **Alarm Rules** page, click **Create Alarm Rule** in the upper right corner.
- Step 6** On the **Create Alarm Rule** page, set parameters as prompted.
1. Configure the rule name and description.
 2. Configure the alarm parameters as prompted.

Table 5-9 Configuring alarm parameters

Parameter	Description	Example Value
Resource Type	Name of the cloud service resource for which the alarm rule is configured.	Data Warehouse Service
Dimension	Metric dimension of the alarm rule. You can select Data Warehouse Nodes or Data Warehouses .	Data Warehouse Node
Monitoring Scope	Resource scope to which an alarm rule applies. Select Specific resources and select one or more monitoring objects. For GaussDB(DWS), select the cluster ID or node ID in the dialog box that is displayed.	Specific resources

Parameter	Description	Example Value
Trigger Rule	You can select an associated template, use an existing template or create a custom template as required.	Create manually
Template	This parameter is valid only when Use template is selected. Select the template to be imported. If no alarm template is available, click Create Custom Template to create one that meets your requirements.	-
Alarm Policy	This parameter is valid only when Create manually is selected. Set the policy that triggers an alarm. For example, trigger an alarm if the CPU usage equals to or is greater than 80% for 3 consecutive periods. Table 5-7 lists the GaussDB(DWS) monitoring metrics.	-
Alarm Severity	Severity of an alarm. Valid values are Critical, Major, Minor , and Informational .	Major

3. Configure the alarm notification parameters as prompted.

Table 5-10 Configuring alarm notifications

Parameter	Description	Example Value
Alarm Notification	Whether to notify users when alarms are triggered. Notifications can be sent as emails or text messages, or HTTP/HTTPS requests sent to the servers. You can enable (recommended) or disable Alarm Notification .	Enable
Validity Period	Cloud Eye sends notifications only within the validity period specified in the alarm rule. For example, if Validity Period is set to 00:00-8:00 , Cloud Eye sends notifications only within 00:00-8:00.	-

Parameter	Description	Example Value
Notification Object	Name of the topic to which the alarm notification is sent. If you enable Alarm Notification , you need to select a topic. If no desired topics are available, create one first, whereupon the SMN service is invoked. For details about how to create a topic, see the <i>Simple Message Notification User Guide</i> . For details about how to create a topic, see the Simple Message Notification User Guide .	-
Trigger Condition	Condition for triggering the alarm. You can select Generated alarm , Cleared alarm , or both.	-

4. After the configuration is complete, click **Next**.

After the alarm rule is created, if the metric data reaches the specified threshold, Cloud Eye will immediately inform you that an exception has occurred.

----End

Transferring Data to OBS

Raw data of metrics is kept for two days on Cloud Eye. You can enable OBS and save the raw data to OBS so that it can be saved for a longer time.

For details about how to configure OBS storage transfer, see "Viewing Alarm History > Configuring OBS Data Storage" in the [Cloud Eye User Guide](#).

5.4 Alarms

5.4.1 Alarm Management

Overview

Alarm management includes viewing and configuring alarm rules and subscribing to alarm information. Alarm rules display alarm statistics and details of the past week for users to view tenant alarms. In addition to providing a set of default GaussDB(DWS) alarm rules, this feature allows you to modify alarm thresholds based on your own services. GaussDB(DWS) alarm notifications are sent using the SMN service.

NOTE

- This feature is supported only in cluster version 8.1.1.200 and later.
- Currently, alarms cannot be categorized and managed by enterprise project.

Visiting the Alarms Page

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane on the left, click **Alarms**.

Step 3 On the page that is displayed:

- **Existing Alarm Statistics**

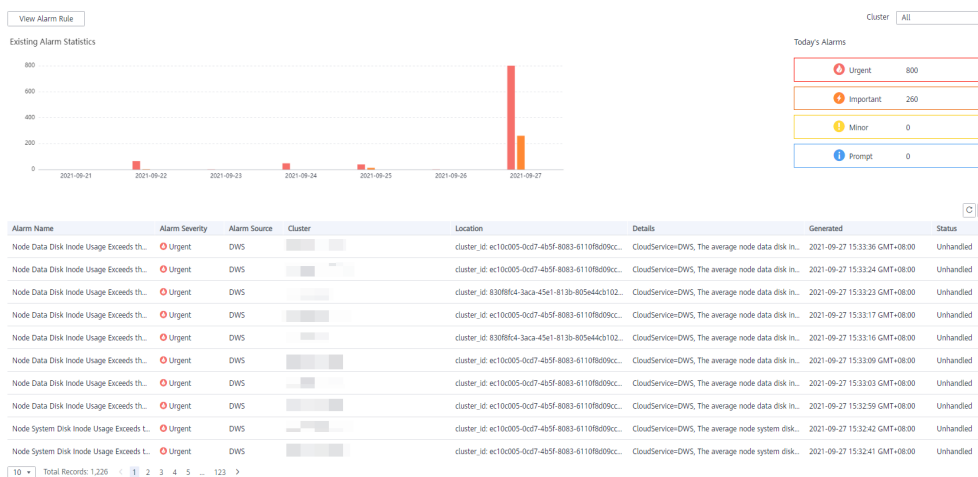
Statistics of the existing alarms in the past seven days are displayed by alarm severity in a bar chart. In this way, you can see clearly the number and category of the alarms generated in the past week.

- **Today's Alarms**

Statistics of the existing alarms on the current day are displayed by alarm severity in a list. In this way, you can see clearly the number and category of the unhandled alarms generated on the day.

- **Alarm details**

Details about all alarms, handled and unhandled, in the past seven days are displayed in a table for you to quickly locate faults, including the alarm name, alarm severity, cluster name, location, description, generation date, and status.



NOTE

The alarm data displayed (a maximum of 30 days) is supported by the Event Service microservice.

----End

Alarm Types and Alarms

NOTE

The alarm policy is triggered based on the current configuration.

Table 5-11 Threshold alarms of DMS alarm sources

Type	Name	Severity	Description
Default	Node CPU Usage Exceeds the Threshold	Urgent	This alarm is generated if the threshold of CPU usage (system + user) of any node in the cluster is exceeded within the specified period and the constraint is not met. The alarm will be cleared when the CPU usage (system + user) is lower than the threshold and the constraint is not met.
Default	Node Data Disk Usage Exceeds the Threshold	Urgent: > 85%; Important: > 80%	This alarm is generated if the threshold of data disk (<code>/var/chroot/DWS/data[n]</code>) usage of any node in the cluster is exceeded within the specified period and the constraint is not met. The alarm will be cleared when the data disk (<code>/var/chroot/DWS/data[n]</code>) usage is lower than the threshold and the constraint is not met.
Default	Node Data Disk I/O Usage Exceeds the Threshold	Urgent	This alarm is generated if the threshold of data disk (<code>/var/chroot/DWS/data[n]</code>) I/O usage (<code>util</code>) of any node in the cluster is exceeded within the specified period and the constraint is not met. The alarm will be cleared when the data disk (<code>/var/chroot/DWS/data[n]</code>) I/O usage (<code>util</code>) is lower than the threshold and the constraint is not met.
Default	Node Data Disk Latency Exceeds the Threshold	Important	This alarm is generated if the threshold of data disk (<code>/var/chroot/DWS/data[n]</code>) I/O latency (<code>await</code>) of any node in the cluster is exceeded within the specified period and the constraint is not met. The alarm will be cleared when the data disk (<code>/var/chroot/DWS/data[n]</code>) I/O latency (<code>await</code>) is lower than the threshold and the constraint is not met.
Default	Data Flushed to Disks of the Query Statement Exceeds the Threshold	Urgent	This alarm is generated if the threshold of data flushed to disks of the SQL statement in the cluster is exceeded within the specified period and the constraint is not met. The alarm can be cleared only after you handle the SQL statement.

Type	Name	Severity	Description
Default	Number of Queuing Query Statements Exceeds the Threshold	Urgent	This alarm is generated if the threshold of the number of queuing SQL statements is exceeded within the specified period. The alarm will be cleared when the number of queuing SQL statements is less than the threshold.
Default	Queue congestion in the cluster default resource pool	Urgent	This alarm is generated if the queue in the default resource pool of a cluster is congested and no alarm suppression conditions are met. This alarm will be cleared if the queue is not congested.
Default	Long SQL Probe Execution Duration in a Cluster	Urgent	This alarm is generated if the DMS alarm module detects a SQL probe execution duration on a server and no alarm suppression conditions are met. If no execution duration exceeds the threshold, the alarm will be automatically cleared. NOTE The alarm is supported only in 8.1.1.300 and later cluster versions. For earlier versions, contact technical support.
Default	A vacuum full operation that holds a table lock for a long time exists in the cluster.	Important	In a specified period, the DMS alarm module detects that VACUUM FULL has been running for a long time in the cluster and blocks other operations. This alarm is generated if there are other SQL statements in the lock wait state and no suppression conditions are met. This alarm will be cleared if VACUUM FULL in the cluster did not cause lock wait. NOTE If this alarm is generated, contact technical support engineers.
Custom	<i>Name of the user-defined threshold alarm</i>	<i>User-defined alarm severity</i>	<i>Alarm description</i>

5.4.2 Alarm Rules

Overview

- Concepts related to threshold alarms
 - Alarm rule: consists of the alarm rule name, rule description, clusters associated with the rule, alarm policy triggering relationship, and alarm policy. An alarm rule can apply to one or all clusters, and can consist of one or more policies. The relationship between alarm policies can be selected in **Triggered Policies**. Each alarm policy consists of the triggers and constraint of each alarm rule.
 - Alarm policy: consists of the triggers, constraint, and alarm severity for an alarm metric.
 - Alarm metric: indicates a database cluster metric, which is generally time series data, for example, node CPU usage and amount of data flushed to disks.
- Alarm rule types
 - Default rule: best practices of GaussDB(DWS) threshold alarms.
 - User-defined rule: personalized alarm rules by configuring or combining monitoring metrics. (The current version supports only user-defined alarm rules of schema usage.)
- Alarm rule operations
 - Modify: modifies an alarm rule. All alarm rules apply (all items of user-defined alarm rules but only some items of the default alarm rules).
 - Enable/Disable: enables or disables an alarm rule. All alarm rules apply. When an alarm rule is enabled, it is added to the check list of the alarm engine and can be triggered normally. Disabled rules are not in the check list.
 - Delete: deletes an alarm rule. You can delete only user-defined rules. Default alarm rules cannot be deleted.

Precautions

After a cluster is migrated, to monitor alarms of the new cluster, change the cluster bound to the alarm rule to the new cluster. You can also create an alarm rule for the new cluster.

Viewing Alarm Rules

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** In the navigation pane on the left, choose **Alarms**.
- Step 3** Click **View Alarm Rule** in the upper left corner. On the page that is displayed, you can see the threshold alarm rules of database cluster monitoring metrics, as shown in the following figure.

Alarm Rule	Rule Status	Associated Cluster	Rule Type	Rule Description	Operation
TCP Retransmission after packet loss	Enable		Default	This alarm is generated if the DMS alarm module detects a high retransmission rate on a server and no alarm suppress...	Modify Disable Delete
File Handle Usage Exceeds Threshold	Enable	All	Default	This alarm is generated if the DMS alarm module detects high file handle usage on a server and no alarm suppress...	Modify Disable Delete
Number of Queuing Query Statements Exceeds the Threshold	Enable	All	Default	This alarm is generated if the threshold of the number of queuing SQL statements is exceeded within the specified p...	Modify Disable Delete
Data Flushed to Disk of the Query Statement Exceeds the T...	Enable	All	Default	This alarm is generated if the threshold of data flushed to disk of the SQL statement in the cluster is exceeded withi...	Modify Disable Delete
Node CPU Usage Exceeds the Threshold	Enable	All	Default	This alarm is generated if the threshold of CPU usage (system + user) of any node in the cluster is exceeded within 1...	Modify Disable Delete
Node Data Disk I/O Usage Exceeds the Threshold	Enable	All	Default	This alarm is generated if the threshold of data disk (vartchrood(DMSData)[] I/O usage (all)) of any node in the clust...	Modify Disable Delete
Node Data Disk Inode Usage Exceeds the Threshold	Enable	All	Default	This alarm is generated if the threshold of data disk (vartchrood(DMSData)[] inode usage of any node in the cluster...	Modify Disable Delete
Node Data Disk Usage Exceeds the Threshold	Enable	All	Default	This alarm is generated if the threshold of data disk (vartchrood(DMSData)[] usage of any node in the cluster is exc...	Modify Disable Delete
Node Data Disk Latency Exceeds the Threshold	Enable		Default	This alarm is generated if the threshold of data disk (vartchrood(DMSData)[] I/O latency (await)) of any node in the c...	Modify Disable Delete
Node Log Disk I/O Usage Exceeds the Threshold	Enable	All	Default	This alarm is generated if the threshold of log disk (vartchrood(DMSManager) I/O usage (all)) of any node in the clus...	Modify Disable Delete

----End

Modifying an Alarm Rule

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** In the navigation pane on the left, click **Alarms**.
- Step 3** Click **View Alarm Rule** in the upper left corner.
- Step 4** On the **Alarm Rules** page that is displayed, click **Modify** in the **Operation** column of the target alarm rule.

NOTE

You can modify only some items of the default rules (associated cluster, alarm policy threshold, time period, and alarm constraint). User-defined rules support modification of all items.

Table 5-12 Alarm rule parameters

Parameter	Description	Example Value
Alarm Rule	The rule name contains 6 to 64 characters and must start with a non-digit character.	-
Description	User-defined description, which contains a maximum of 490 characters.	-
Associated Cluster	You can select a cluster of the current tenant from the drop-down list box as the monitoring cluster of the alarm module.	All
Triggered Policies	Policy triggering relationships are as follows: <ul style="list-style-type: none"> ● Independent: Alarm policies are triggered independently of each other. ● Priority: Alarm policies are triggered by priority. Policies of a lower priority will be automatically triggered after those of a higher priority. 	Independent

Parameter	Description	Example Value
Alarm Policy	The alarm policies are as follows: <ul style="list-style-type: none">• Metric: GaussDB(DWS) monitoring metric, which is the data source used by the alarm engine for threshold determination.• Alarm Object: databases in the selected cluster and schemas in the selected databases.• Trigger: calculation rule for threshold determination of a monitoring metric. Select the average value within a period of time of a metric to reduce the probability of alarm oscillation.• Constraint: suppresses the repeated triggering and clearance of alarms of the same type within the specified period.• Alarm Severity: includes Urgent, Important, Minor, and Prompt.	-

Step 5 Confirm the information and click **OK**.

----End

Creating an Alarm Rule

Step 1 Log in to the GaussDB(DWS) management console.

Step 2 In the navigation pane on the left, click **Alarms**.

Step 3 Click **View Alarm Rule** in the upper left corner.

Step 4 Click **Create Alarm Rule** in the upper right corner. You can configure items, such as the alarm rule name, rule description, associated cluster, and alarm policy. For details, see [Table 5-12](#).

NOTE

Currently, only alarm rules of schema usage metrics can be created on GaussDB(DWS).

----End

5.4.3 Alarm Subscriptions

You can subscribe to GaussDB(DWS) alarm notifications to receive notifications by SMS message, email, or application when an alarm of a specified severity is generated.

Creating a Subscription

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** In the navigation pane on the left, choose **Alarms > Subscriptions**.
- Step 3** Click **Create Subscription** in the upper left corner of the page.
- Step 4** In the **Subscription Settings** area, configure the basic information and alarm severity of the subscription.

Subscription Settings
Edit subscription information and select alarm severities

* Status ?

* Subscription Name ?

Alarm Severity ▼

Table 5-13 Subscription parameters

Parameter	Description
Status	Whether to enable the alarm subscription. When you disable a subscription, you will not receive the corresponding alarm notifications, but the subscription will not be deleted.
Subscription Name	Name of the alarm subscription: <ul style="list-style-type: none"> • Contains only letters, digits, hyphens (-), and underscores (_), and must start with a letter or digit. • Contains 1 to 256 characters.
Alarm Severity	Severity of the alarm you want to subscribe to: Urgent, Important, Minor, or Prompt

- Step 5** The **Subscribed Alarms** area displays the subscribed alarms by subscription settings. Select an SMN topic from the drop-down list.

To create a topic, click **Create Topic** to switch to the SMN console page. For details, .

 **NOTE**

The selected topic must have granted GaussDB(DWS) the permission for publishing messages to the topic. To grant permissions, configure topic policies on the SMN management console. When configuring the topic policy, select **DWS** as the service that can publish messages to this topic.

- Step 6** Confirm the information and click **OK**.

----End

Modifying a Subscription

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** In the navigation pane on the left, choose **Alarms > Subscriptions**.
- Step 3** In the **Operation** column of the target subscription, click **Edit**.
- Step 4** On the **Edit Subscription** page displayed, modify the parameters. For details, see [Step 4](#) to [5](#).
- Step 5** Click **OK**.
- End

Deleting a Subscription

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** In the navigation pane on the left, choose **Alarms > Subscriptions**.
- Step 3** In the **Operation** column of the target subscription, click **Delete**. A confirmation dialog box is displayed.
- Step 4** Click **Yes** to delete the subscription.
- End

5.4.4 Alarm Handling

5.4.4.1 DWS_200000001 Node CPU Usage Exceeds the Threshold

Description

GaussDB(DWS) collects the CPU usage of each node in a cluster every 30 seconds. If the average CPU usage of a node in the last 10 minutes (configurable) exceeds 90% (configurable), an alarm is reported indicating that the node CPU usage exceeds the threshold. If the average usage is lower than 85% (that is, the reporting threshold minus 5%), the alarm is cleared.

NOTE

If the average CPU usage of a node is always greater than the alarm threshold, the alarm is generated again 24 hours (configurable).

Attributes

Alarm ID	Alarm Category	Alarm Severity	Alarm Type	Service Type	Auto Cleared
DWS_200000001	Management plane alarm	Urgent: > 90%	Operation alarm	GaussDB(DWS)	Yes

Parameters

Category	Name	Description
Location information	Name	Node CPU Usage Exceeds the Threshold
	Type	Operation alarm
	Generation time	Time when the alarm is generated
Other information	Cluster ID	Cluster details such as resourceId and domain_id

Impact on the System


If the CPU usage is high for a long time, service processes may respond slowly or become unavailable.

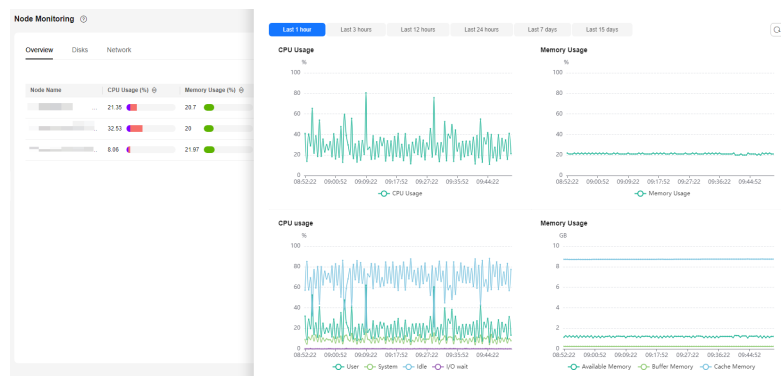
Possible Causes

- Complex services occupy a large number of CPU resources.
- The CPU configuration of the cluster is too low to meet service requirements.

Handling Procedure

Step 1 Check the CPU usage of each node.

1. Log in to the GaussDB(DWS) console.
2. Choose **Monitoring > Alarm**, select the cluster for which the alarm is generated in the cluster selection drop-down list in the upper right corner, view the alarm information of the cluster in the last seven days, and locate the name of the node for which the alarm is generated based on the location information.
3. Choose **Dedicated Clusters > Clusters**, locate the row that contains the cluster for which the alarm is generated, and click **Monitoring Panel** in the **Operation** column.
4. Choose **Monitoring > Node Monitoring > Overview** to view the CPU usage of each node in the current cluster. Click  on the right to view the CPU performance metrics in the last 1, 3, 12, or 24 hours and see whether there is a sharp increase in the CPU usage.



- If the CPU usage frequently increases and then returns to normal in a short period of time, it indicates that the CPU usage temporarily spikes during service execution. In this case, you can adjust the alarm threshold through **Step 2** to reduce the number of reported alarms.
- If the CPU usage remains high for a long time, it indicates that the cluster is overloaded. In this case, check cluster services by referring to **Step 3** or enhance the cluster flavor. For details, see **Changing the Node Flavor**.

Step 2 Check whether the CPU usage alarm configuration is proper.

1. Choose **Monitoring > Alarm** and click **View Alarm Rule**.
2. Locate the row that contains the **Node CPU Usage Exceeds the Threshold**, and click **Modify** in the **Operation** column. The **Modifying an Alarm Rule page** is displayed.
3. Adjust the alarm threshold and detection period. A higher alarm threshold and a longer detection period indicate a lower alarm sensitivity. For details about the GUI configuration, see **Alarm Rules**.

Step 3 Check whether the CPU usage of the current cluster service is too high.




1. Choose **Dedicated Clusters > Clusters**, locate the row that contains the target cluster, and click **Monitoring Panel** in the **Operation** column.
2. On the monitoring page, choose **Monitoring > Real-Time Queries**, click , and select **CPU Time (ms)** to view the query with the longest CPU time.

Figure 5-16 Viewing CPU time information



3. After confirming with the service side, select the query ID to be stopped and click **Stop Query**.

Figure 5-17 Terminating a query



Step 4 See advanced tuning operations in "Troubleshooting" > "Cluster Performance".

----End

Alarm Clearance

After the CPU usage decreases, the alarm is automatically cleared.

5.4.4.2 DWS_200000006 Node Data Disk Usage Exceeds the Threshold

Description

GaussDB(DWS) collects the usage of all disks on each node in a cluster every 30 seconds.

- If the maximum disk usage in the last 10 minutes (configurable) exceeds 80% (configurable), a major alarm is reported. If the average disk usage is lower than 75% (that is, the alarm threshold minus 5%), this major alarm is cleared.
- If the maximum disk usage in the last 10 minutes (configurable) exceeds 85% (configurable), a critical alarm is reported. If the average disk usage is lower than 85% (that is, the alarm threshold minus 5%), this critical alarm is cleared.

NOTE

If the maximum disk usage is always greater than the alarm threshold, the system generates an alarm again 24 hours later (configurable).

Attributes

Alarm ID	Alarm Category	Alarm Severity	Alarm Type	Service Type	Auto Cleared
DWS_200000006	Management plane alarm	Urgent: > 85%; important: > 80%	Operation alarm	GaussDB(DWS)	Yes

Parameters

Category	Name	Description
Location information	Name	Node Data Disk Usage Exceeds the Threshold
	Type	Operation alarm
	Generation time	Time when the alarm is generated
Other information	Cluster ID	Cluster details such as resourceId and domain_id

Impact on the System


If the cluster data volume or temporary data spill size increases and the usage of any single disk exceeds 90%, the cluster becomes read-only, affecting customer services.

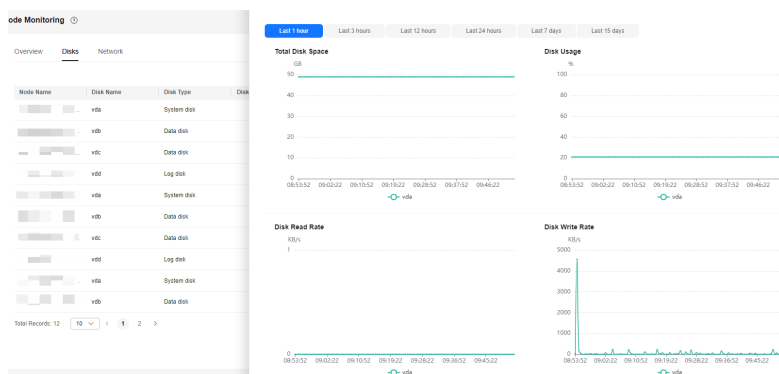
Possible Causes

- The service data volume increases rapidly, and the cluster disk capacity configuration cannot meet service requirements.
- Dirty data is not cleared in a timely manner.
- There are skew tables.

Handling Procedure

Step 1 Check the disk usage of each node.

1. Log in to the GaussDB(DWS) console.
2. On the **Alarms** page, select the current cluster from the cluster selection drop-down list in the upper right corner and view the alarm information of the cluster in the last seven days. Locate the name of the node for which the alarm is generated and the disk information based on the location information.
3. Choose **Dedicated Clusters > Clusters**, locate the row that contains the cluster for which the alarm is generated, and click **Monitoring Panel** in the **Operation** column.
4. Choose **Monitoring > Node Monitoring > Disks** to view the usage of each disk on the current cluster node. If you want to view the historical monitoring information about a disk on a node, click  on the right to view the disk performance metrics in the last 1, 3, 12, or 24 hours.



- If the data disk usage frequently increases and then returns to normal in a short period of time, it indicates that the disk usage temporarily spikes due to service execution. In this case, you can adjust the alarm threshold through [Step 2](#) to reduce the number of reported alarms.
- If the usage of a data disk exceeds 90%, read-only is triggered and error **cannot execute INSERT in a read-only transaction** is reported for write-related services. In this case, you can refer to [Step 3](#) to delete unnecessary data.
- If the usage of more than half of the data disks in the cluster exceeds 70%, the data volume in the cluster is large. In this case, refer to [Step 4](#) to clear data or perform [Disk Capacity Expansion](#).
- If the difference between the highest and lowest data disk usage in the cluster exceeds 10%, refer to [Step 5](#) to handle data skew.

Step 2 Check whether the alarm configuration is proper.

1. Return to the GaussDB(DWS) console, choose **Monitoring > Alarm** and click **View Alarm Rule**.
2. Locate the row that contains **Node Data Disk Usage Exceeds the Threshold** and click **Modify** in the **Operation** column. On the **Modifying an Alarm Rule** page, view the configuration parameters of the current alarm.
3. Adjust the alarm threshold and detection period. A higher alarm threshold and a longer detection period indicate a lower alarm sensitivity. For details about the GUI configuration, see [Alarm Rules](#).
4. If the data disk specification is high, you are advised to increase the threshold based on historical disk monitoring metrics. Otherwise, perform other steps. If the problem persists, you are advised to perform [Disk Capacity Expansion](#).

Step 3 Check whether the cluster is in the read-only state.

1. When a cluster is in read-only state, stop the write tasks to prevent data loss caused by disk space exhaustion.
2. Return to the GaussDB(DWS) console and choose **Dedicated Clusters > Clusters**. In the row of the abnormal cluster whose cluster status is **Read-only**, click **Cancel Read-only**.
3. In the displayed dialog box, confirm the information and click **OK** to cancel the read-only state for the cluster. For details, see [Removing the Read-only Status](#).
4. After the read-only mode is disabled, use the client to connect to the database and run the **DROP/TRUNCATE** command to delete unnecessary data.

NOTE

You are advised to lower the disk usage to below 70%. Check whether there are other tables that need to be rectified by referring to [Step 4](#) and [Step 5](#).

Step 4 Check whether the usage of more than half of the data disks in the cluster exceeds 70%.

1. Run the **VACUUM FULL** command to clear data. For details, see [Solution to High Disk Usage and Cluster Read-Only](#). Connect to the database, run the following SQL statement to query tables whose dirty page rate exceeds 30%, and sort the tables by size in descending order:

```
SELECT schemaname AS schema, relname AS table_name, n_live_tup AS analyze_count,
pg_size_pretty(pg_table_size(relid)) as table_size, dirty_page_rate
FROM PGXC_GET_STAT_ALL_TABLES
WHERE schemaName NOT IN ('pg_toast', 'pg_catalog', 'information_schema', 'cstore', 'pmk')
AND dirty_page_rate > 30
ORDER BY table_size DESC, dirty_page_rate DESC;
```

The following is an example of the possible execution result of the SQL statement (the dirty page rate of a table is high):

```
schema | table_name | analyze_count | table_size | dirty_page_rate
-----+-----+-----+-----+-----
public | test_table |          4333 | 656 KB   |          71.11
(1 row)
```

2. If any result is displayed in the command output, clear the tables with a high dirty page rate in serial mode.
`VACUUM FULL ANALYZE schema.table_name`

NOTICE

The **VACUUM FULL** operation occupies extra defragmentation space, which is Table size x (1 – Dirty page rate). As a result, the disk usage temporarily increases and then decreases. Ensure that the remaining space of the cluster is sufficient and will not trigger read-only when the **VACUUM FULL** operation is performed. You are advised to start from small tables. In addition, the **VACUUM FULL** operation holds an exclusive lock, during which access to the operated table is blocked. You need to properly arrange the execution time to avoid affecting services.

3. If no command output is displayed, no table with a high dirty page rate exists. You can expand the node or disk capacity of the cluster based on the following data warehouse types to prevent service interruption caused by read-only triggered by further disk usage increase.
 - a. To scale out a storage-compute coupled data warehouse with cloud SSDs, see [Disk Capacity Expansion of an EVS Cluster](#).
 - b. To scale out a storage-compute coupled data warehouse with local SSDs or a standalone system, see [Scaling Out a Cluster](#).

Step 5 Check whether the difference between the highest and lowest data disk usages in the cluster exceeds 10%.

1. If the data disk usage differs greatly, connect to the database and run the following SQL statement to check there are skew tables in the cluster:

```
SELECT schemaname, tablename, pg_size_pretty(totalsize), skewratio FROM pgxc_get_table_skewness WHERE skewratio > 0.05 ORDER BY totalsize desc;
```

The following is an example of the possible execution result of the SQL statement:

```
schemaname | tablename | pg_size_pretty | skewratio
-----+-----+-----+-----
scheduler | workload_collection | 428 MB | .500
public | test_table | 672 KB | .429
public | tbl_col | 104 KB | .154
scheduler | scheduler_storage | 32 KB | .250
(4 rows)
```

2. If the SQL statement output is displayed, select another distribution column for the table with severe skew based on the table size and skew rate. For 8.1.0 and later versions, use the [ALTER TABLE](#) syntax to adjust the distribution column. For other versions, see [How Do I Adjust Distribution Columns?](#)

----End

Alarm Clearance

After the disk usage decreases, the alarm is automatically cleared.

5.4.4.3 DWS_200000009 Node Data Disk I/O Usage Exceeds the Threshold

Description

GaussDB(DWS) collects the data disk I/O usage of each cluster node every 30 seconds. This alarm is generated when the average usage of a data disk on a node exceeds 90% (configurable) in the last 10 minutes (configurable), and is

automatically cleared when the average usage drops below 85% (alarm threshold minus 5%).

 **NOTE**

- If the data disk I/O usage of a node is always greater than the alarm threshold, the alarm is generated again 24 hours later (configurable).
- When using an SSD storage-based cluster, disk I/O can surpass 100% as the service volume grows. But, this does not always mean there is a performance bottleneck. To confirm the alarm's validity, you should evaluate the service's actual running status.

Alarm Attributes

Alarm ID	Alarm Category	Alarm Severity	Alarm Type	Service Type	Auto Cleared
DWS_200000009	Management plane alarm	Urgent: > 90%	Operation alarm	GaussDB(DWS)	Yes

Alarm Parameters

Category	Name	Description
Location information	Name	Node Data Disk I/O Usage Exceeds the Threshold
	Type	Operation alarm
	Generation time	Time when the alarm is generated
Other information	Cluster ID	Cluster details such as resourceId and domain_id

Impact on the System

- High disk I/O usage affects data read and write performance, thereby affecting cluster performance.
- A large number of disk writes occupy the disk capacity. If the disk capacity exceeds 90%, the cluster becomes read-only.

Possible Causes

- A large number of read or write operations are performed during peak hours.
- A large amount of data spills to disks due to the execution of complex statements.
- Data is scanned by the Scan operator.

Handling Procedure

Step 1 Choose **Dedicated Clusters > Clusters**, locate the row that contains the target cluster, and click **Monitoring** in the **Operation** column.


Step 2 In the navigation pane on the left, choose **Monitoring > Node Monitoring**. On the **Node Monitoring** page, click the **Disks** tab to view the data disk I/O usage and disk I/O rate.

If the disk I/O rate is high and the data disk usage keeps increasing, it indicates that services are writing data to disks. This may be caused by complex queries.

Step 3 Click **Queries** in the navigation pane to view the real-time queries.

If the execution time of a statement exceeds the expected time, stop the query and check the disk I/O usage again. For details, see [2](#).

Figure 5-18 Terminating a query



----End

Alarm Clearance

This alarm is automatically cleared when the data disk I/O usage drops to a certain value.

5.4.4.4 DWS_200000012 Node Data Disk Latency Exceeds the Threshold

Description

GaussDB(DWS) collects the data disk latency of each node in the cluster every 30 seconds. This alarm is generated when the average latency of a data disk on a node exceeds 400 ms (configurable) in the last 10 minutes (configurable), and is automatically cleared when the average latency drops below 400 ms.

NOTE

If the data disk latency of a node is always greater than the alarm threshold, this alarm is generated again after 24 hours (configurable).

Alarm Attributes

Alarm ID	Alarm Category	Alarm Severity	Alarm Type	Service Type	Auto Cleared
DWS_200000012	Management plane alarm	Important: > 400 ms	Operation alarm	GaussDB(DWS)	Yes

Alarm Parameters

Category	Name	Description
Location information	Name	Node Data Disk Latency Exceeds the Threshold
	Type	Operation alarm
	Generation time	Time when the alarm is generated
Other information	Cluster ID	Cluster details such as resourceId and domain_id

Impact on the System

High disk latency will slow down the data read/write speed, causing the cluster performance to deteriorate.

Possible Causes

The database is in peak hours and there are a large number of read and write requests.

Handling Procedure

- Step 1** Choose **Dedicated Clusters > Clusters**, locate the row that contains the target cluster, and click **Monitoring** in the **Operation** column.
- Step 2** In the navigation pane on the left, choose **Monitoring > Node Monitoring**. On the **Node Monitoring** page, view the CPU usage, disk usage, and memory usage.

If the CPU usage and disk I/O rate are high, the cluster is in peak hours. You can adjust the latency threshold based on service requirements. For details, see [3](#).
- Step 3** Return to the console home page. In the navigation pane on the left, choose **Management > Alarms**. On the displayed page, click **View Alarm Rule** in the upper left corner.
- Step 4** Locate the row that contains **Node Data Disk Latency Exceeds the Threshold**, and click **Modify** in the **Operation** column. On the displayed page, change the threshold.

----End

Alarm Clearance

This alarm is automatically cleared when the data disk latency drops to a certain value.

5.4.4.5 DWS_200000016 Data Spilled to Disks of the Query Statement Exceeds the Threshold

Description

During the execution of service queries, the database may choose to store the temporary result to the disk, which is called **operator spilling**.

GaussDB(DWS) checks the load management records of jobs being executed on CNs through the **GS_WLM_SESSION_STATISTICS** view every 60 seconds and calculates the maximum amount of data spilled to DNs.

If the number of SQL statements spilled to disks exceeds 5 GB (configurable) within 10 minutes (configurable), an alarm is reported indicating that a query statement triggers the data spill threshold. This alarm is automatically cleared when the data spill drop below the alarm conditions. For details about how to modify alarm configurations, see **Modifying Alarm Rules**.

NOTE

If blocked SQL statements that can trigger the alarm persists, the alarm is generated again after 24 hours (configurable).

Attributes

Alarm ID	Alarm Category	Alarm Severity	Alarm Type	Service Type	Auto Cleared
DWS_200000016	Management plane alarm	Urgent: > 5 GB	Operation alarm	GaussDB(DWS)	Yes

Parameters

Category	Name	Description
Location information	Name	Data Spilled to Disks of the Query Statement Exceeds the Threshold
	Type	Operation alarm
	Generation time	Time when the alarm is generated
Other information	Cluster ID	Cluster details such as resourceid and domain_id

NOTE

You can connect to the database and run the **SELECT * FROM GS_WLM_SESSION_STATISTICS** command to view the **max_spill_size** column in the view.

Impact on the System

If a large amount of data spills to disks, a large number of system I/O resources are occupied. As a result, the data disk space may be insufficient or exhausted, triggering the database to become read-only and interrupting services.

Possible Causes

- The amount of service data spilled to disks exceeds the alarm threshold.
- The performance of the SQL query plan is poor, causing a large amount of data to be imported to the memory and spilled to disks.
- Expired data is not cleared in a timely manner. As a result, too much invalid data is scanned and spilled to disks.

Handling Procedure

Step 1 Check whether the execution plan is poor in performance.

1. Obtain the SQL statement from the additional information of the alarm, run the **ANALYZE** statement on the involved tables. Run the SQL statement again and check whether the amount of data spilled to disks decreases.
2. If there is no obvious effect, run the **EXPLAIN PERFORMANCE** command to view the actual execution information of the alarm SQL statement. For details, see [SQL Execution Plan](#). If the estimated (**operator memory**) and peak memory usage (**Peak Memory**) are high, exceeding 20% of `max_process_memory`, optimization of the query is necessary based on execution information. For details, see [SQL Tuning Process](#).

Step 2 Check whether the alarm configuration is proper.


1. Log in to the GaussDB(DWS) console, choose **Monitoring > Alarm** and click **View Alarm Rule**.
2. Click **Modify** in the **Operation** column of the row that contains **Data Flushed to Disks of the Query Statement Exceeds the Threshold**. The **Modifying an Alarm Rule** page is displayed.
3. If the cluster disk capacity is high, you can increase the alarm reporting threshold. It is recommended that the alarm reporting threshold be less than or equal to 5% of the capacity of a single data disk.

CAUTION

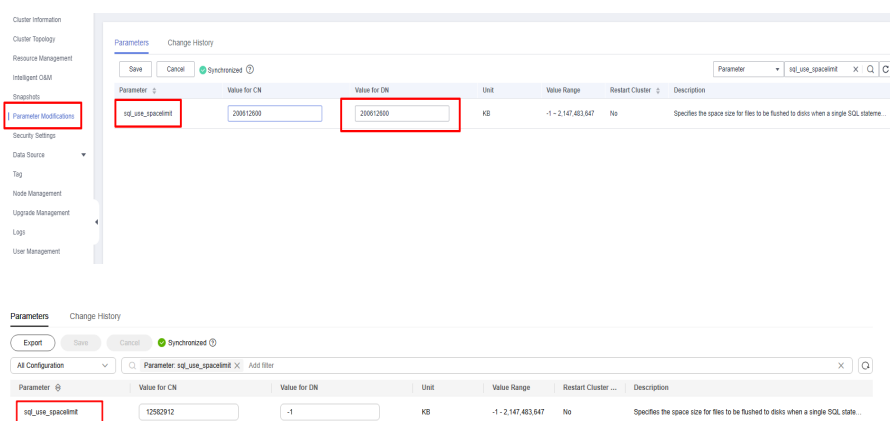
If the threshold is too large, data spilled to a disk may cause disk usage alarms or even the cluster to be read-only. If the data disk usage is close to or exceeds 80%, you are advised to clear unnecessary data when adjusting the threshold. For details about the GUI configuration, see [Alarm Rules](#).

Step 3 Kill the SQL statements that cause large data spills.

1. Return to the GaussDB(DWS) console.
2. Choose **Dedicated Clusters > Clusters**, locate the row that contains the cluster for which the alarm is generated, and click **Monitoring Panel** in the **Operation** column.

3. Choose **Monitoring > Queries**. Click  to see the data spill in the **Max. DN Data Spill (MB)** column.
4. After confirming with the service side, select the query ID of the query to be stopped and click **Stop Query**.
5. Adjust the database parameters for controlling the disk space of service statements. For details about the parameters, see [Statement Disk Space Control](#). For details about the procedure, see [Modifying Database Parameters](#).

For example, the default value of `sql_use_spacelimit` is 10% of the total storage space of the DB instance. If the storage space is sufficient, you can increase the value. If the disk write volume of a single DN exceeds the value, GaussDB(DWS) stops the query and displays a message indicating that the disk write volume of a single DN exceeds the threshold.



----End

Alarm Clearance

This alarm is automatically cleared when data spill drops down to a low level.

5.4.4.6 DWS_200000017 Number of Queuing Query Statements Exceeds the Threshold

Description

When real-time query monitoring is enabled, GaussDB(DWS) checks the queuing status of jobs on CNs through the `GS_WLM_SESSION_STATISTICS` view every 60 seconds by default.

This alarm is generated when the number of queuing SQL statements in the cluster exceeds 10 (configurable) within 10 minutes (configurable), and is automatically cleared when the number of queuing SQL statements drops below 10.

NOTE

If there continues to be queuing query statements more than the alarm threshold, the alarm is generated again 24 hours later (configurable).

Alarm Attributes

Alarm ID	Alarm Category	Alarm Severity	Alarm Type	Service Type	Auto Cleared
DWS_2000000017	Management plane alarm	Urgent: >10	Operation alarm	GaussDB(DWS)	Yes

Alarm Parameters

Category	Name	Description
Location information	Name	Number of Queuing Query Statements Exceeds the Threshold
	Type	Operation alarm
	Generation time	Time when the alarm is generated
Other information	Cluster ID	Cluster details such as resourceId and domain_id

Impact on the System


SQL queries are blocked. As a result, the execution time is too long.


Possible Causes

The number of queuing query statements during service execution exceeds the alarm threshold.

Handling Procedure

Check whether the current queuing jobs in the cluster are normal.

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Monitoring > Alarm** in the navigation pane on the left, select the current cluster from the cluster selection drop-down list in the upper right corner and view the alarm information of the cluster in the last seven days. Locate the name of the cluster that triggers the alarm based on the location information.
- Step 3** Choose **Dedicated Clusters > Clusters**, locate the row that contains the cluster for which the alarm is generated, and click **Monitoring Panel** in the **Operation** column.
- Step 4** Choose **Monitoring > Queries** to view the real-time sessions and queries of the current cluster. Select the **Queries** tab to view the status of jobs being executed in the current cluster. Click  and select **Blocking Time (ms)** and **Waiting** status.

Click  to sort the values of **Blocking Time (ms)**. You can view the information about the waiting SQL statements with the longest blocking time. If a query job is in the waiting state and the blocking time is abnormal, you can terminate the query.

Current queuing status of the statements, including:

- **Global**: global queuing.
- **Respool**: resource pool queuing.
- **CentralQueue**: queuing on the CCN
- **Transaction**: being in a transaction block
- **StoredProc**: being in a stored procedure
- **None**: not in a queue
- **Forced None**: being forcibly executed (transaction block statement or stored procedure statement are) because the statement waiting time exceeds the specified value

----End

Alarm Clearance

This alarm is automatically cleared when the number of queuing statements drops below the threshold.

5.4.4.7 DWS_2000000018 Queue Congestion in the Default Cluster Resource Pool

Description

GaussDB(DWS) uses resource pools to control memory, I/O, and CPU resources, manages and allocates resources based on task priorities, and manages user services loads. For details about resource pools, see [Resource Pool](#). When resources are insufficient, some SQL statements have to queue up to wait for other statements to be executed. For details, see [CCN Queuing Under Dynamic Load Management](#).

GaussDB(DWS) checks the queue in the default resource pool **default_pool** every 5 minutes. This alarm is generated when there are SQL statements that are queued up for a long time (20 minutes by default and configurable). This alarm is automatically cleared when the alarm threshold is no longer met.

NOTE

If blocked SQL statements that can trigger the alarm persists, the alarm is generated again after 24 hours (configurable).

Attributes

Alarm ID	Alarm Category	Alarm Severity	Alarm Type	Service Type	Auto Cleared
DWS_2000000018	Management plane alarm	Critical	Operation alarm	GaussDB(DWS)	Yes

Parameters

Category	Name	Description
Location information	Name	Queue Congestion in the Default Cluster Resource Pool
	Type	Operation alarm
	Generation time	Time when the alarm is generated
Other information	Cluster ID	Cluster details such as resourceid and domain_id

Impact on the System

When the default resource pool is blocked, all complex queries (estimated memory greater than or equal to 32 MB) associated with the default resource pool in the cluster may also be blocked. The queries in the queue are woken up only when the running queries are complete.

Possible Causes

- The estimated query memory usage is too large. As a result, the accumulated estimated memory usage exceeds the upper limit of the dynamic available memory, causing CCN queuing.
- Competition for public resources such as CPU and I/O deteriorates the performance of running queries.

Handling Procedure


Step 1 Check whether the queue is caused by too large estimated memory.

Rectify the fault by referring to [CCN Queuing Under Dynamic Load Management](#).

Step 2 Check whether the available memory of the cluster is normal.

1. Log in to the GaussDB(DWS) console.
2. Choose **Monitoring** > **Alarm** in the navigation pane on the left, select the current cluster from the cluster selection drop-down list in the upper right corner and view the alarm information of the cluster in the last seven days.

Locate the name of the cluster that triggers the alarm based on the location information.

3. Choose **Dedicated Clusters** > **Clusters**, locate the row that contains the cluster for which the alarm is generated, and click **Monitoring Panel** in the **Operation** column.
4. Choose **Monitoring** > **Node Monitoring** > **Overview** to view the memory usage of each node in the current cluster. If you want to view the historical monitoring information about the memory usage of a node, click  on the right to view the memory usage in the last 1, 3, 12, or 24 hours.
If the cluster memory usage is low (for example, lower than 50%), the alarm may be generated because the estimated memory usage of queries is too large. In this case, perform the **Analyze** operation on related tables.

Step 3 Check the competition of other resources.

1. Check the CPU, I/O, and network usage of the cluster by referring to section **Step 2**.
2. If the database is fully loaded, query **Real-Time Top SQL** and kill the statements that occupy a large number of resources.

Step 4 Check whether too many queries are submitted in a short period of time.

1. Run the following SQL statement to query the task execution status:

```
SELECT
  s.resource_pool AS rpname, s.node_group,
  count(1) AS session_cnt,
  SUM(CASE WHEN a.enqueue = 'waiting in global queue' THEN 1 ELSE 0 END) AS global_wait,
  SUM(CASE WHEN s.lane= 'fast' AND a.state = 'active' AND (a.enqueue IS NULL OR a.enqueue = 'no
waiting queue') THEN 1 ELSE 0 END) AS fast_run,
  SUM(CASE WHEN s.lane= 'fast' AND a.enqueue = 'waiting in respool queue' THEN 1 ELSE 0 END)
AS fast_wait,
  SUM(CASE WHEN s.lane= 'slow' AND a.state = 'active' AND (a.enqueue IS NULL OR a.enqueue =
'no waiting queue') THEN 1 ELSE 0 END) AS slow_run,
  SUM(CASE WHEN s.lane= 'slow' AND (a.enqueue = 'waiting in ccn queue' OR a.enqueue = 'waiting
in respool queue') THEN 1 ELSE 0 END) AS slow_wait,
  SUM(CASE WHEN (a.enqueue IS NULL OR a.enqueue = 'no waiting queue') AND a.state = 'active'
THEN statement_mem ELSE 0 END) AS est_mem
FROM pgxc_session_wlmstat s,pgxc_stat_activity a
WHERE s.threadid=a.pid(+) AND s.attribute != 'Internal'
GROUP BY 1,2;
```

The following is an example of the possible execution result of the SQL statement:

rpname	node_group	session_cnt	global_wait	fast_run	fast_wait	slow_run	slow_wait	est_mem
default_pool	installation	6	0	0	0	0	0	0
root	installation	1	0	0	0	0	0	0

(2 rows)

- In the query result, if the value of **slow_wait** corresponding to **default_pool** is not 0, the cluster is fully loaded due to too many jobs. As a result, an alarm is generated. In this case, you can locate the row that contains the specified cluster on the console, choose **Monitoring Panel** in the **Operation** column. On the displayed page, choose **Monitoring** > **Real-Time Queries** to query the task with the longest execution time, and kill the task.

- If the alarm is frequently generated, you are advised to schedule services in off-peak hours or create new resource pools to manage system resources in a more refined manner. For details, see [Creating a Resource Pool](#).

----End

Alarm Clearance

This alarm is automatically cleared when the resource pool blocking is relieved.

NOTE

To view historical blocked SQL statements, locate the row that contains the target cluster on the console, choose **Monitoring Panel** in the **Operation** column. On the displayed page, choose **Monitoring > History** to query the execution time of historical SQL statements.

5.4.4.8 DWS_200000020 Long SQL Probe Execution Duration in a Cluster

Alarm Description

GaussDB(DWS) collects the execution status of the SQL probe on each node in the cluster every 30 seconds. If the execution duration of an SQL probe on a server in a cluster exceeds twice the threshold (or another user-defined value), a critical alarm is generated. If the execution duration of all SQL probes falls below the threshold, the critical alarm is cleared.

NOTE

If the SQL probe duration remains higher than the alarm reporting threshold, the alarm is generated again in 24 hours(or another user-defined value).

Attributes

Alarm ID	Alarm Category	Alarm Severity	Alarm Type	Service Type	Auto Cleared
DWS_200000020	Management plane alarm	Important	Operation alarm	GaussDB(DWS)	Yes

Alarm Parameters

Category	Name	Description
Location information	Name	Long SQL Probe Execution Duration in a Cluster
	Type	Operation alarm
	Generation time	Time when the alarm is generated
Other information	Cluster ID	Cluster details such as resourceId and domain_id

Impact on the System

The cluster performance deteriorates or the cluster is faulty.

Possible Causes

The service load of the cluster is high or the cluster is faulty. As a result, the execution of the SQL probe becomes slow.

Handling Procedure

- Step 1** In the navigation pane of the monitoring panel, choose **Utilities > SQL Probes**. Check SQL probe execution.
- Step 2** In the navigation pane, choose **Monitoring > Performance Monitoring**. Check the monitoring metrics such as the CPU usage, disk usage, and memory usage to determine whether the workloads are high or any metric is abnormal.
- Step 3** In the navigation pane, choose **Monitoring > Real-Time Queries**. Check whether there are queries or sessions that have been running for a long time and affect cluster running. You can terminate abnormal sessions or queries.

----End

Alarm Clearance

This alarm is automatically cleared when the time consumed by an SQL probe on all servers in all clusters falls below the threshold.

5.4.4.9 DWS_200000023 A Vacuum Full Operation That Holds a Lock for A Long Time Exists in the Cluster

Alarm Description

VACUUM FULL holds a level-8 lock on a table. If it holds the lock on a table for longer than 20 minutes (or another user-defined value), a major alarm is reported, indicating that the VACUUM FULL operation holds a lock for too long in the cluster. This major alarm is cleared when VACUUM FULL is complete.

Attributes

Alarm ID	Alarm Category	Alarm Severity	Alarm Type	Service Type	Auto Cleared
DWS_200000023	Management plane alarm	Important	Operation alarm	GaussDB(DWS)	Yes

Alarm Parameters

Category	Name	Description
Location information	Name	A Vacuum Full Operation That Holds a Table Lock for A Long Time Exists in the Cluster
	Type	Operation alarm
	Generation time	Time when the alarm is generated
Other information	Cluster ID	Cluster details such as resourceId and domain_id

Impact on the System


Other operations cannot the table. As a result, workloads cannot be executed.

Possible Causes


There is a VACUUM FULL operation that holds a table lock for a long time in the cluster.

Handling Procedure

Step 1 In the navigation pane of the monitoring panel, choose **Monitoring > Real-Time Queries > Sessions**. In the session list, set the search criteria to **LIKE** and search for the keyword **vacuum full**.



Step 2 Check whether there is a table lock waiting for VACUUM FULL to complete by querying the locked object.



Step 3 Check whether the VACUUM FULL operation needs to be handled.

1. Check whether VACUUM FULL is a system behavior and whether it affects system functions. If VACUUM FULL does not affect other service queries, wait until it is complete. The alarm will be automatically cleared.
2. If VACUUM FULL affects normal service execution, you can find and kill related sessions on the **Real-Time Queries** tab and re-execute VACUUM FULL later.

----End

Alarm Clearance

This alarm is automatically cleared when the VACUUM FULL operation is complete.

5.5 Event Notifications

5.5.1 Event Notifications Overview

Overview

GaussDB(DWS) uses the Simple Message Notification (SMN) service to send notifications of GaussDB(DWS) events. The SMN function is only available by subscription. In a subscription, you need to specify one or more event filtering conditions. When an event that matches all filtering conditions occurs, GaussDB(DWS) sends a notification based on the subscription. The filter conditions include the **Event Type** (for example, **Management**, **Monitoring**, or **Security**), **Event Severity** (for example, **Normal** or **Warning**), and **Event Source Category** (for example, **Cluster** or **Snapshot**).

Supported Event Types and Events

Events are records of changes in the user's cluster status. Events can be triggered by user operations (such as audit events), or may be caused by cluster service status changes (for example, cluster repaired successfully or failed to repair the cluster). The following tables list the events and event types supported by GaussDB(DWS).

- The following table lists the events whose **Event Source Category** is **Cluster**.

Table 5-14 Events whose **Event Source Category** is **Cluster**

Event Type	Event Name	Event Severity	Event
Management	createClusterFail	Warning	Failed to create the cluster.
Management	createClusterSuccess	Normal	Cluster created successfully.
Management	createCluster	Normal	Cluster creation started.
Management	extendCluster	Normal	Cluster scale-out started.
Management	extendClusterSuccess	Normal	Cluster scaled out successfully.
Management	extendClusterFail	Warning	Failed to scale out the cluster.

Event Type	Event Name	Event Severity	Event
Management	deleteClusterFail	Warning	Failed to delete the cluster.
Management	deleteClusterSuccess	Normal	Cluster deleted successfully.
Management	deleteCluster	Normal	Cluster deletion started.
Management	restoreClusterFail	Warning	Failed to restore the cluster.
Management	restoreClusterSuccess	Normal	Cluster restored successfully.
Management	restoreCluster	Normal	Cluster restoration started.
Management	restartClusterFail	Warning	Failed to restart the cluster.
Management	restartClusterSuccess	Normal	Cluster restarted successfully.
Management	restartCluster	Normal	Cluster restarted.
Management	configureMRSExtDataSources	Normal	Configuration of MRS external data source for the cluster started.
Management	configureMRSExtDataSourcesFail	Warning	Failed to configure the MRS external data source for the cluster.
Management	configureMRSExtDataSourcesSuccess	Normal	MRS external data source configured successfully for the cluster.
Management	deleteMRSExtDataSources	Normal	Deletion of MRS external data source for the cluster started.
Management	deleteMRSExtDataSourcesFail	Warning	Failed to delete the MRS external data source for the cluster.
Management	deletedMRSExtDataSourcesSuccess	Normal	MRS external data source deleted successfully for the cluster.

Event Type	Event Name	Event Severity	Event
Management	bindEipToCluster	Normal	Bound an EIP to the cluster.
Management	bindEipToClusterFail	Warning	Failed to bind an EIP to the cluster.
Management	unbindEipToCluster	Normal	Unbound an EIP from the cluster.
Management	unbindEipToCluster-Fail	Warning	Failed to unbind an EIP from the cluster.
Management	refreshEipToCluster	Normal	Refreshed the cluster's EIP.
Management	refreshEipToCluster-Fail	Warning	Failed to refresh the cluster's EIP.
Management	dmsCreateWDRSuccessfully	Normal	Generating the workload report...
Management	failedToCreateWDR	Warning	Failed to generate the workload report.
Management	dmsDeleteWDRSuccessfully	Normal	Workload report deleted.
Management	failedToDeleteWDR	Warning	Failed to delete a workload report.
Management	dmsUpdateWDRConfigSuccessfully	Normal	Workload report parameters updated.
Management	failedToUpdateWDR-Config	Warning	Failed to update workload report parameters.
Management	dmsCreateWorkloadSnapshotSuccessfully	Normal	Creating the workload snapshot...
Management	failedToCreateWorkloadSnapshot	Warning	Failed to create the workload snapshot.
Security	resetPasswordFail	Warning	Failed to reset the password.
Security	resetPasswordSuccess	Normal	Password of the cluster reset successfully.
Security	updateConfiguration	Normal	Updating security parameters of the cluster started.

Event Type	Event Name	Event Severity	Event
Security	updateConfiguration-Fail	Warning	Failed to update security parameters of the cluster.
Security	updateConfiguration-Success	Normal	Security parameters of the cluster updated successfully.
Monitoring	repairCluster	Normal	The node is faulty. Repairing the cluster starts.
Monitoring	repairClusterFail	Warning	Failed to repair the cluster.
Monitoring	repairClusterSuccess	Normal	Cluster repaired successfully.

- The following table lists the events whose **Event Source Category** is **Snapshot**.

Table 5-15 Events whose **Event Source Category** is **Snapshot**

Event Type	Event Name	Event Severity	Event
Management	deleteBackup	Normal	Snapshot deleted successfully.
Management	deleteBackupFail	Warning	Failed to delete the snapshot.
Management	createBackup	Normal	Snapshot creation started.
Management	createBackupSuccess	Normal	Snapshot created successfully.
Management	createBackupFail	Warning	Failed to create the snapshot.

5.5.2 Subscribing to Event Notifications

After subscribing to GaussDB(DWS) event notification, you will receive notifications by text message, email, or application when management, monitoring, or security events occur in a specific cluster or snapshot.

Creating a Subscription

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation tree on the left, click **Event Management**.
- Step 3** On the **Event Management** page, choose **Subscription > Create Subscription**.
- Step 4** In the **Subscription Settings** area, set basic subscription information and event filtering.

The **Subscribed Event List** area displays the events filtered by the system based on the subscription settings.

Table 5-16 Subscription parameters

Parameter	Description
Notification	Enable or disable event subscription. After notification is disabled, the system stops sending notifications of subscribed events but does not delete the subscription.
Subscription Name	Enter the name of a subscription. <ul style="list-style-type: none">The name can contain letters (upper or lower case), digits, hyphens (-), and underscores (_) and must start with a letter or digit.The name must be between 1 and 256 characters in length.
Event Type	Select the type of the event to be subscribed. Possible values are Management, Monitoring, and Security .
Event Severity	Select the alarm severity of the event. Possible values are Normal and Warning .
Event Source Category	Select the event source category: cluster, snapshot,.

- Step 5** Select a message notification topic from the **Message Notification Topic** drop-down list.
 - The selected topic must have granted GaussDB(DWS) the permission to publish messages to the topic.
If GaussDB(DWS) has not been authorized to publish messages to the selected topic, go to the topic management page of the SMN console to configure topic authorization. For details, see **Topic Management > Configuring Topic Policies** in *Simple Message Notification User Guide*. When

configuring the topic policy, select **GaussDB(DWS)** for **Services that can publish messages to this topic**.

- To create a topic, click **Create Topic**. The SMN console is displayed. For details, see **Topic Management > Creating a Topic** in the *Simple Message Notification User Guide*.

Step 6 Click **OK** to complete the subscription.

----End

Modifying the Subscription

Step 1 On the **Event Management** page of the GaussDB(DWS) console, click **Subscription**.

Step 2 In the **Operation** column of the row containing the specified subscription, click **Edit** to enter the **Edit Subscription** page.

Step 3 On the **Edit Subscription** page, set the parameters to be modified. For details, see [Step 4](#) to [Step 6](#) in section "Creating a Subscription".

----End

Deleting the Subscription

Step 1 On the **Event Management** page of the GaussDB(DWS) console, click **Subscription**.

Step 2 In the **Operation** column of the row containing the specified subscription, click **Delete**. The **Delete Subscription** dialog box is displayed.

Step 3 Click **Yes** to delete the subscription.

----End


5.5.3 Viewing Events

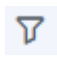
This section describes how to search for events that occur in a cluster or snapshot.

Step 1 Log in to the GaussDB(DWS) management console.

Step 2 In the navigation tree on the left, click **Events**.

On the **Events** tab page, all events that occur in the clusters or snapshots are displayed by default.

You can sort the events in descending or ascending order by clicking  next to **Time**.

You can filter the events by clicking  next to a field (except **Time**) and selecting the criteria.

----End

6 Specifications Change and Scaling

6.1 Viewing Inspection Results

Context

GaussDB(DWS) allows you to inspect the cluster before making any changes like scaling, changing specifications, or upgrading. Simply click **Inspect** on the relevant page, and the system will check if the cluster's health status and metrics meet the requirements for the change. Once the inspection is passed, you can proceed with the change. If the inspection fails, you can view the inspection details page to see which items did not pass the inspection. From there, you can handle the inspection items based on the details provided. For details about the inspection standards, see [Table 6-1](#).

NOTE

This feature is supported only in cluster version 8.1.1 or later.

Precautions

- The inspection plug-in 8.3.1.100 or later has been installed in the cluster.
- The inspection result is valid for 24 hours, during which you can make the change operation. Once the 24-hour validity period expires, you will need to perform the inspection again.
- If the cluster has not been inspected within 24 hours before the change, inspect it before making any changes like scaling, changing specifications, or upgrading. Ensure that the inspection is passed before proceeding with the change.

Viewing Inspection Details

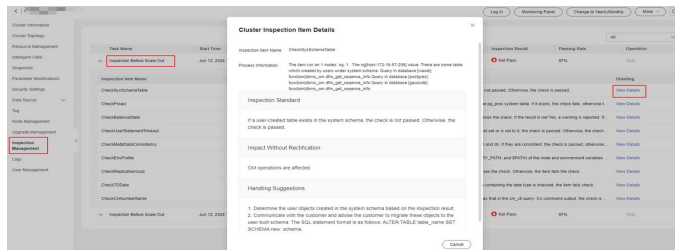
Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the cluster list, click the name of the target cluster.

Step 3 On the cluster details page, click **Inspection Management**.

- Step 4** Click the drop-down button next to its name to check the inspection status, execution progress, inspection result, and pass rate. For more details about these inspection items, click **View Details** in the task's row.

Figure 6-1 Viewing inspection details



NOTE

After creating an inspection task on the configuration change page, you can keep track of its progress and view details. You can even stop the inspection from that same page.

----End

Stopping an Inspection Task

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the cluster list, click the name of the target cluster.
- Step 3** On the cluster details page, click **Inspection Management**.
- Step 4** Locate the row that contains the inspection task and click **Stop** in the **Operation** column to stop the inspection task.

----End

Inspection Criteria

Table 6-1 Inspection criteria

Change Operation	Item	Check Criteria
Scaling operations and typical modifications	CheckTimeZone	The inspection passes if all nodes in the cluster use the same time zone, and fails if they do not.
	CheckSpaceUsage	When usage goes beyond the warning threshold (set at 70% by default), a warning is issued. If it goes beyond the NG threshold (set at 90% by default), the inspection fails. The inspection fails if the available space in the GAUSSHOME/PGHOST/GPHOME/GAUSSLOG/tmp/data directory is less than the threshold.

Change Operation	Item	Check Criteria
	CheckClusterState	It verifies the CM process, fenced UDF, and cluster status. If the CM process is missing, the inspection fails. A warning is issued if the fenced UDF status is down . The inspection passes for a normal cluster status, but fails otherwise.
	CheckEnvProfile	It checks the environment variables (\$GAUSSHOME , \$LD_LIBRARY_PATH , and \$PATH) of a node. If these variables exist and are properly configured, the inspection passes. Otherwise, the inspection fails.
	CheckReadonly Mode	It checks the value of default_transaction_read_only on all nodes that contain CNs in the cluster. If the value is off , the inspection passes. Otherwise, the inspection fails.
	CheckCatchup	It checks whether the CatchupMain function can be found in the gaussdb process stack. If it cannot be found, the inspection passes. Otherwise, the inspection fails.
	CheckCollector	It checks whether the information collection is successful. If it is, the inspection passes. Otherwise, the inspection fails.
	CheckTrust	If any node is not trusted, the inspection fails. Otherwise, the inspection passes.
	CheckBalanceState	It checks the Balanced attribute of the cluster. If it is Yes , the inspection passes. A warning is displayed if it is not. If the query fails, the inspection also fails.
	CheckCnNumber Same	It checks if the number of CNs queried by running the /opt/dws/xml/cluster.xml command is different from that queried by running the cm_ctl query -Cv command. If they are different, the inspection passes. If not, the inspection fails.

Change Operation	Item	Check Criteria
	CheckCMPParam	It checks if the value of enable_transaction_read_only is on and if the value of coordinator_heartbeat_timeout is consistent on each node. If both are true, the inspection passes.
	CheckUtilslib	It checks for the existence of the \$GAUSSHOME/utilslib directory. If it exists, the inspection fails. If it does not exist, the inspection passes.
	CheckPgxcgroup	It checks the number of records in the pgxc_group table where in_redistribution is Y . If the number is 0 , the inspection passes. If the number is greater than 0, the inspection fails.
	CheckLockState	It checks whether the cluster is locked. If the cluster is not locked, the inspection passes. If it is locked, the inspection fails.
	CheckDBConnection	It checks whether the database can be connected. If it can, the inspection passes. Otherwise, the inspection fails.
	CheckGUConsistent	It checks whether the GUC parameters of CNs and DNs are consistent. If they are consistent, the inspection passes. Otherwise, the inspection fails.
	CheckTDDate	The inspection fails if the ORC table in the TD database exists and has columns of the date type.
	CheckPgxcRedistrib	If any temporary table remains in the database after data redistribution, the inspection fails.
	CheckMetaData	If metadata in the system table is consistent, the inspection passes. Otherwise, the inspection fails.
	CheckGUCSetting	If the GUC parameters in postgresql.conf are consistent with those in pg_settings , the inspection passes. Otherwise, the inspection fails.
	CheckProacl	The inspection fails if proacl in the pg_proc system table has usernames with only digits. Otherwise, it passes.

Change Operation	Item	Check Criteria
	CheckMetaData Consistency	If the data in the system table between CN and DN is consistent, the inspection passes. Otherwise, the inspection fails.
	CheckReturnTyp e	If the return value is invalid, the inspection fails. Otherwise, the inspection passes.
	CheckUltraWide Table	If a table with more than 996 columns exists, the inspection fails. Otherwise, the inspection passes.
	CheckDataRedis Schema	If a data_redis schema exists in the database and the owner name is not redisuser , the inspection fails. Otherwise, the inspection passes.
	CheckDiskSpace Limited	If the disk space of the user is limited, the inspection fails. Otherwise, the inspection passes.
	CheckTableColla te	The inspection fails if the database has a PCK table or a column-store partitioned table with a collate field. Otherwise, it passes.
	CheckDefaultOri entation	It checks the GUC parameter. If both the database and default_orientation are set to row storage, the inspection passes. Otherwise, the inspection fails.
	CheckReplicatio- nUuid	The inspection passes if there is no replicated table that uses the default UUID. Otherwise, it fails.
	CheckUserState mentTimeout	If statement_timeout is not set or is set to 0 , the inspection passes. Otherwise, the inspection fails.
	CheckJsonb	Run the select attrelid::regclass from pg_attribute a join pg_type t on a.atttypid = t.oid and t.typname = 'jsonb' group by 1 SQL statement. If the jsonb type is used, the inspection fails. Otherwise, the inspection passes.
	CheckLengthOfI ndex	Run the SELECT length(pg_get_indexdef(indexrelid)) FROM pg_index order by 1 desc limit 1 SQL statement. If the result is greater than 192 x 1024, the inspection fails. Otherwise, the inspection passes.

Change Operation	Item	Check Criteria
	CheckLengthOfTable	Run the select c.oid from pg_class c,pg_namespace n where c.relnamespace=n.oid and relkind='r' and n.nspname not in ('cstore') and length(n.nspname) ' c.relname')>=64; SQL statement. If the result is not empty, the inspection fails. Otherwise, the inspection passes.
	CheckUseWorkloadManager	Run the show use_workload_manager SQL statement. If the result is on , the inspection passes. Otherwise, the inspection fails. This inspection item is not included in version 8.1.3.320 and beyond. Meaning, it has been verified and passed in later versions.
Pre-upgrade health check	CheckClusterParams	The cluster configuration parameters (IP address, port, and path parameters) specified in postgresql.conf or pgxc_node should match those in the static configuration file. Otherwise, the inspection fails.
	CheckCNNum	It checks the number of CNs in the cluster. If the number is greater than 2 and no more 10, the inspection passes. Otherwise, the inspection fails.
	CheckDDL	Start a transaction to delete schemas and tables. If the transaction can be committed, the inspection passes. Otherwise, the inspection fails.
	CheckTimeZone	The inspection passes if all nodes in the cluster use the same time zone, and fails if they do not.
	CheckXidEpoch	It checks the XID consumption. If the value is greater than or equal to 2 to the power of 32, the inspection fails.
	CheckCnNumberSame	It checks if the number of CNs queried by running the /opt/dws/xml/cluster.xml command is different from that queried by running the cm_ctl query -Cv command. If they are different, the inspection passes. If not, the inspection fails.

Change Operation	Item	Check Criteria
	CheckGaussVer	The inspection fails if the binary files in the \$GAUSSHOME/bin directory on each node have identical versions.
	CheckPsort	The inspection fails if the psort index exists.
	CheckCatchup	It checks whether the CatchupMain function can be found in the gaussdb process stack. If it cannot be found, the inspection passes. Otherwise, the inspection fails.
	CheckClusterState	It verifies the CM process, fenced UDF, and cluster status. If the CM process is missing, the inspection fails. A warning is issued if the fenced UDF status is down . The inspection passes for a normal cluster status, but fails otherwise.
	CheckMetaData Consistency	If the data in the system table between CN and DN is consistent, the inspection passes. Otherwise, the inspection fails.
	CheckDependSystemObj	If self-created objects depend on system objects, the inspection fails. Otherwise, the inspection passes.
	CheckPgKeywords	If names of tables, columns, functions, or data types are new reserved keywords, the inspection fails. Otherwise, the inspection passes.
	CheckReadOnly Mode	It checks the value of default_transaction_read_only on all nodes that contain CNs in the cluster. If the value is off , the inspection passes. Otherwise, the inspection fails.
	CheckMetaData	If metadata in the system table is consistent, the inspection passes. Otherwise, the inspection fails.
	CheckGUCSetting	If the GUC parameters in postgresql.conf are consistent with those in pg_settings , the inspection passes. Otherwise, the inspection fails.

Change Operation	Item	Check Criteria
	CheckPgxcgroup	It checks the number of records in the pgxc_group table where in_redistribution is Y . If the number is 0 , the inspection passes. If the number is greater than 0, the inspection fails.
	CheckCmserverStandby	If the value of cm_server in the cluster is standby , the inspection passes. Otherwise, a warning is displayed.
	CheckSpaceUsage	When usage goes beyond the warning threshold (set at 70% by default), a warning is issued. If it goes beyond the NG threshold (set at 90% by default), the inspection fails. The inspection fails if the available space in the GAUSSHOME/PGHOST/GPHOME/GAUSSLOG/tmp/data directory is less than the threshold.
	CheckEnvProfile	It checks the environment variables (\$GAUSSHOME , \$LD_LIBRARY_PATH , and \$PATH) of a node. If these variables exist and are properly configured, the inspection passes. Otherwise, the inspection fails.
	CheckBalanceState	It checks the Balanced attribute of the cluster. If it is Yes , the inspection passes. A warning is displayed if it is not. If the query fails, the inspection also fails.
	CheckTDDate	The inspection fails if the ORC table in the TD database exists and has columns of the date type.
	CheckCatalog	If there is a custom database object in pg_catalog , the inspection fails. Otherwise, the inspection passes.
	CheckPgauthid	If the most significant bit of oid in pg_authid is 1 , the inspection fails. Otherwise, the inspection passes.
	CheckSysdate	If the sysdate view is used in tables, views, and stored procedures, the inspection fails. Otherwise, the inspection passes.

Change Operation	Item	Check Criteria
	CheckFilesNumber	If the number of tmp files in the GAUSSHOME/PGHOST/GPHOME directory is greater than 10,000, the inspection fails. Otherwise, the inspection passes.
	CheckKeyFilesExist	If the upgrade_version , conf , control , and data files exist in key directories, the inspection passes. Otherwise, the inspection fails.
	CheckReturnTyp	If the return value is invalid, the inspection fails. Otherwise, the inspection passes.
	CheckTrust	If any node is not trusted, the inspection fails. Otherwise, the inspection passes.

6.2 Managing Nodes

Overview

On the **Nodes** tab page, you can view the node list of the current cluster, add new nodes to or remove nodes from it, and view the node usage, status, flavor, and AZ.


To modify an alias, click  next to it.

Figure 6-2 Managing Nodes



Node Name	Node Alias Name	AZ	Resource Status	Node Status	Node Flavor
[Redacted]	-	AZ3	In Use	Available	dnx.vlarge.4
[Redacted]	-	AZ3	In Use	Available	dnx.vlarge.4
[Redacted]	-	AZ3	In Use	Available	dnx.vlarge.4

NOTE

- This feature is supported only in 8.1.1.200 or later cluster versions.
- The hybrid data warehouse (standalone) does not support node management.

Adding Nodes

This function is more suited for large-scale scale-out. Nodes can be added in batches in advance without interrupting services. To add 180 nodes, add them in three batches of 60 nodes each. If any nodes fail to be added, retry adding them. Once all 180 nodes are added, use them for scaling out.

Precautions

- Nodes can be added only when no other task is running on the management side.

- The storage size of a new node must be the same as that of each of the existing nodes in the cluster.
- An added node, typically for scaling out, is referred to as an idle node. It starts incurring charges once added. You are advised to only add nodes when needed and promptly use them for scaling out.
- The anti-affinity rule dictates that the number of nodes to be added at a time must be an integer multiple of the cluster ring size. For example, if the cluster ring size is 3, the number of nodes to be added must be an integer multiple of 3.
- In the anti-affinity deployment mode, when a node is idle and fails due to power-off or other causes, it makes other nodes in its server group unavailable. In this case, you should remove and re-add the failed node.
- The anti-affinity rule dictates that, if a node fails to be added and is rolled back, other nodes that are being added in the same server group will also be rolled back.

Procedure

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Clusters > Dedicated Clusters**. All clusters are displayed by default.
- Step 3** Click the name of the target cluster. On the **Cluster Information** page that is displayed, choose **Nodes**.
- Step 4** Click **Add**, enter the number of nodes to be added, and click **Next: Confirm**. If you create a BMS cluster, contact technical support to add the cluster to the whitelist for cross-flavor scale-out. After the cluster is whitelisted, **Resource Flavor** will be displayed on the **Add Node** page. If there are not enough IP addresses in the original subnet, you can add idle nodes from other subnets.
- Step 5** After confirming that the information is correct, click **Submit**. The **Nodes** page is displayed. On this page, you can start adding nodes. Nodes that fail to be added are automatically rolled back and recorded in the failure list.

----End

Removing Nodes

Precautions

- Nodes can be removed only when no other task is running on the management side.
- Only nodes whose resource status is **Idle** can be removed. Nodes that are in use cannot be removed.
- In anti-affinity deployment, nodes are removed by cluster ring. For example, when you remove a node, other nodes in the same ring will be automatically selected and displayed.

Procedure

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Clusters > Dedicated Clusters**. All clusters are displayed by default.

- Step 3** Click the name of the target cluster. On the **Cluster Information** page that is displayed, choose **Nodes**.
- Step 4** On the **Nodes** page, select the node to be deleted and click **Delete**.
- Step 5** Confirm the information and click **OK**. After the deletion is successful, the node is no longer displayed on the **Nodes** page.

----End

6.3 Scaling Nodes

6.3.1 Scaling Out a Cluster

When you need more compute and storage resources, add more nodes for cluster scale-out on the management console.

NOTE

- When you scale out the standard data warehouse cluster, use the same storage specifications as the cluster.
- Nodes cannot be added to a hybrid data warehouse (standalone).
- If you create a BMS cluster, contact technical support to add you to the whitelist for cross-flavor scale-out. After you are whitelisted, **Resource Flavor** will be displayed on the **Add Node** page.
- If the number of subnet IP addresses is insufficient, cross-subnet scale-out is allowed.

After the data in a data warehouse is deleted, the occupied disk space may not be released, resulting in dirty data and disk waste. Therefore, if you need to scale out your cluster due to insufficient storage capacity, run the **VACUUM** command to reclaim the storage space first. If the used storage capacity is still high after you run the **VACUUM** command, you can scale out your cluster.

Impact on the System

- Before the scale-out, disable the client connections that have created temporary tables because temporary tables created before or during the scale-out will become invalid and operations performed on these temporary tables will fail. Temporary tables created after the scale-out will not be affected.
- After you start a scale-out task, the cluster automatically takes a snapshot before the task begins.
- During the scale-out, functions such as cluster restart, scale-out, snapshot creation, database administrator password resetting, and cluster deletion are disabled.
- During an offline scale-out, the cluster automatically restarts. Therefore, the cluster changes to **Unavailable** for a period of time. After the cluster is restarted, the status becomes **Available**. After scale-out, the system dynamically redistributes user data among all nodes in the cluster.
- During offline scale-out, stop all services or run only a few query statements. During table redistribution, a shared lock is added to tables. All insert, update, and delete operations as well as DDL operations on the tables are blocked for

a long time, which may cause a lock wait timeout. After a table is redistributed, you can access the table. Do not perform queries that take more than 20 minutes during the redistribution (the default time for applying for the write lock during redistribution is 20 minutes). Otherwise, data redistribution may fail due to lock wait timeout.

- In an online scale-out, during node addition, the cluster is locked. Database objects are checked when the cluster is locked. To ensure that the cluster is successfully locked, avoid executing statements that create or delete databases and tablespaces while adding nodes to the cluster.
- During online scale-out, you can perform insert, update, and delete operations on tables, but data updates are still blocked for a short period of time. Redistribution consumes lots of CPU and I/O resources, which will greatly impact job performance. Therefore, perform redistribution when services are stopped or during periods of light load. Phase-based scale-out is also recommended: Perform high-concurrency redistribution during periods of light load, and stop redistribution or perform low-concurrency redistribution during periods of heavy load.
- If a new snapshot is created for the cluster after the scale-out, the new snapshot contains data on the newly added nodes.
- If the cluster scale-out fails, the database automatically performs the rollback operation in the background so that the number of nodes in the cluster can be restored to that before the scale-out.
 - If the rollback is successful and the cluster can be normally used, you can perform **Scale Out** again. If the scale-out still fails, contact the technical support.
 - If the rollback fails due to some exceptions, the cluster may become **Unavailable**. In this case, you cannot perform **Scale Out** or restart the cluster. Contact the technical support.
- In the DWS 3.0 scale-out scenario, if the number of buckets allocated to each DN is not between [3, 20], automatic scaling is triggered. You can view the number of buckets using the GUC parameter **table_buckets**.
 - Currently, buckets can only be scaled offline. The procedure is the same as that of the existing scaling procedure. The system automatically determines and executes the bucket scaling process.
 - During scaling, the cluster will be restarted and all connections will be closed. The restart takes several minutes.
 - After the restart is complete, the database can be read but cannot be written until data redistribution is complete.

For example, if the number of buckets on the current node is 32 and the number of DNs in the logical cluster is 9, and the number of DNs needs to be increased to 15, as $32/15=2$ (rounded down) does not fall within the range [3,20], bucket scale-out will be triggered.

Prerequisites

- The cluster to be scaled out is in the **Available** or **Unbalanced** state.
- The number of nodes to be added must be less than or equal to the available nodes. Otherwise, system scale-out is not allowed.
- To scale out a cluster as an IAM user, ensure that the IAM user has permissions for VPC, EVC, and BMS.

Scaling Out a Cluster

NOTE

- A cluster becomes read-only during scale-out. Exercise caution when performing this operation.
- To ensure data security, you are advised to create a snapshot before the scale-out. For details about how to create a snapshot, see [Manual Snapshots](#).
- After you start a scale-out, the system first checks for scale-out prerequisites. If your cluster fails the check, modify configurations as prompted and try again. For details, see [What Do I Do If the Scale-out Check Fails?](#)

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters** > **Dedicated Clusters**.

All clusters are displayed by default.

Step 3 In the **Operation** column of the target cluster, choose **More** > **Scale Node** > **Scale Out**. The scale-out page is displayed. If you create a BMS cluster, contact technical support to add the cluster to the whitelist for cross-flavor scale-out. After the cluster is whitelisted, **Resource Flavor** will be displayed on the **Add Node** page. If the IP addresses of the original subnet are insufficient, idle nodes can be added across subnets.

NOTE

- Before scaling out the cluster, it is crucial to verify if it meets the inspection conditions. Click **Immediate Inspection** to complete the inspection and proceed to the next step only if it passes. For more information, see [Viewing Inspection Results](#).

Step 4 Specify the number of nodes to be added.

- DNs are added during scale-out. For details about how to add CNs, see [Managing CNs](#).
- The number of nodes after scale-out must be at least three nodes more than the original number. The maximum number of nodes that can be added depends on the available quota. In addition, the number of nodes after the scale-out cannot exceed 256.

If the node quota is insufficient, click **Increase quota** to submit a service ticket and apply for higher node quota.

- Flavor of the new nodes must be the same as that of existing nodes in the cluster.
- The VPC, subnet, and security group of the cluster with new nodes added are the same as those of the original cluster.
- The number of nodes to be added to a multi-AZ cluster must be a multiple of 3.

Step 5 Configure advanced parameters.

- If you choose **Default**, **Auto Redistribution** will be enabled, **Scale Online** will be disabled, and **Redistribution Mode** will be **Offline** by default.
- If you choose **Custom**, you can configure the following advanced configuration parameters for online scale-out:
 - **Scale Online**: Online scale-out can be enabled. During online scale-out, data can be added, deleted, modified, and queried in the database; and

some DDL syntaxes are supported. Errors will be reported for unsupported syntaxes.

- **Terminate Blocked Job:** If you enable online scale-out, you can configure automatic job termination.
- **Time Before Blocked Job Termination (s):** If job termination is enabled and congestion occurs during online scale-out, the system waits for the duration you specified and then terminates congested jobs. The value can be an integer in the range 30 to 1200.

 **NOTE**

Clusters of version 8.2.1.100 and later support job termination.

- **Auto Redistribution:** Automatic redistribution can be enabled. If automatic redistribution is enabled, data will be redistributed immediately after the scale-out is complete. If this function is disabled, only the scale-out is performed. In this case, to redistribute data, select a cluster and choose **More > Scale Node > Redistribute**.
- **Redistribution Concurrency:** If automatic redistribution is enabled, you can set the number of concurrent redistribution tasks. The value range is 1 to 32. The default value is 4.
- **Redistribution Mode:** It can be set to **Online** or **Offline**. After confirming that the information is correct, click **OK** in the displayed dialog box.

Step 6 Click **Next: Confirm**.

Step 7 Click **Submit**.

- After you submit the scale-out application, task information of the cluster changes to **Scaling out** and the process will take several minutes. During the scale-out, the cluster automatically restarts. Therefore, the cluster status will stay **Unavailable** for a while. After the cluster is restarted, the status will change to **Available**. After the scale-out is complete, the system dynamically redistributes user data in the cluster, during which the cluster is in the **Read-only** state.
- A cluster is successfully scaled out only when the cluster is in the **Available** state and task information **Scaling out** is not displayed. Then you can use the cluster.
- If **Scale-out failed** is displayed, the cluster fails to be scaled out.

----End

Scaling Out with Idle Nodes

To ensure reliability, prepare ECS first by referring to [Adding Nodes](#) for a large-scale cluster, and scale out the cluster using idle nodes.

 **NOTE**

- Disable automatic redistribution when you scale out a large-scale cluster to facilitate retries upon failures for improved reliability.
- After the scale-out is complete, manually perform [redistribution](#) to ensure that multiple retries can be performed in this phase.

Precautions

- A number of available nodes must be added to the cluster in advance so that idle nodes can be created and added for scale-out.
- The anti-affinity rule dictates that the number of idle nodes to be added must be an integer multiple of the cluster ring size.
- Make sure to configure the scale-out task before submitting it. This involves completing the scale-out preparation. Once done, wait for a moment.
- After you start a scale-out, the system first checks for scale-out prerequisites. If your cluster fails the check, modify configurations as prompted and try again. For details, see [What Do I Do If the Scale-out Check Fails?](#)

Procedure

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters > Dedicated Clusters**. All clusters are displayed by default.

Step 3 In the **Operation** column of the target cluster, choose **More > Scale Node > Scale Out**.

If there are idle nodes in the cluster, the system displays a message asking you whether to add nodes.

Step 4 Click the corresponding button to make scale-out preparations and wait until the preparation is complete.

Step 5 Configure the parameters as required. For details, see [Scaling Out a Cluster](#).

Then click **Next: Confirm**.

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

----End

Viewing Scaling Details

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters > Dedicated Clusters**.

Step 3 In the **Task Information** column of a cluster, click **View Details**.

Step 4 Check the scale-out status of the cluster on the scaling details page.

Figure 6-3 Viewing scale-out details

Tasks ×

Cluster Name: test_1206 Task Name: Scaling out Running Status: ▶ Running

	Task Information	Estimate	Scale-Out Started	Completed	Status
Create Node	Prepare Create Node [0/3]	1min	2022-12-28 11:53:02	--	▶ Running
	Create VM [0/3]	5min	--	--	○ Waiting
	Deliver Configurations [0/3]	2min	--	--	○ Waiting
	Install Software [0/3]	2min	--	--	○ Waiting
	Post Create Node [0/3]	1min	--	--	○ Waiting
Build Node	Waiting for Instance Created	--	2022-12-28 11:53:02	--	▶ Running
	Build Node	13min	--	--	○ Waiting
	Post Build Node	1min	--	--	○ Waiting

----End

6.3.2 Cluster Redistribution

6.3.2.1 Redistributing Data

Data redistribution, where data in existing nodes is evenly allocated to the new nodes after you scale out a cluster, is a time-consuming yet crucial task that accelerates service response.

By default, redistribution is automatically started after cluster scale-out. For enhanced reliability, disable the automatic redistribution function and manually start a redistribution task after the scale-out is successful. In this way, both scale-out and redistribution can be retried upon failures.

Currently, **offline redistribution** and **online redistribution** are supported. The default mode is offline redistribution.

Before redistribution starts or when redistribution is paused, you can set redistribution priorities for the tables that have not been redistributed by schema or table.

NOTICE

- The cluster redistribution function is supported in 8.1.1.200 or later cluster versions.
- This function can be manually enabled only when the cluster task information displays **To be redistributed** after scale-out.
- You can also select the redistribution mode when you configure cluster scale-out (see **Configure advanced parameters**).
- Redistribution queues are sorted based on the relpage size of tables. To ensure that the relpage size is correct, you are advised to perform the **ANALYZE** operation on the tables to be redistributed.

Offline Redistribution

Precautions

- In offline redistribution mode, the database does not support DDL and DCL operations. Tables that are being redistributed support only simple DQL operations.
- During table redistribution, a shared lock is added to tables. All insert, update, and delete operations as well as DDL operations on the tables are blocked for a long time, which may cause a lock wait timeout. Do not perform queries that take more than 20 minutes during the redistribution (the default time for applying for the write lock during redistribution is 20 minutes). Otherwise, data redistribution may fail due to lock wait timeout.

Procedure

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters > Dedicated Clusters**. All clusters are displayed by default.

Step 3 In the **Operation** column of the target cluster, choose **More > Scale Node > Redistribute**, as shown in the following figure.

The **Redistribution** page is displayed.

Step 4 On the **Redistribute** page that is displayed, keep the default **offline** redistribution mode and click **Next: Confirm** to submit the task.

----End

Online Redistribution

Precautions

In online redistribution mode, the database supports partial DDL and DCL operations.

- Local tables that are being redistributed support insert, delete and update operations and some DDL operations:
 - **INSERT, DELETE, UPDATE, MERGE INTO, OVERWRITE, UPSERT**
 - Join queries across node groups
 - Local table renaming, schema modification, **DROP, TRUNCATE, TRUNCATE-PARTITION**
- The following operations cannot be performed on tables that are being redistributed:
 - Run **ALTER TABLE** statements (except for **TRUNCATE PARTITION**), including adding or deleting columns or partitions.
 - Create, modify, or delete indexes.
 - Perform **VACUUM FULL** and **CLUSTER** operations on tables.
 - Modify the sequence objects on which a column depends, including creating and modifying them. Typical statements are **CREATE** and **ALTER SEQUENCE ... OWNED BY**.

- During the redistribution of a table with more than 996 columns, **UPDATE** and **DELETE** statements cannot be executed. **SELECT** and **INSERT** statements are allowed.
- Database and tablespace objects cannot be created, deleted, or modified during redistribution.
- A partition swap can be performed only if the redistribution is complete for both of the tables to be swapped. The two tables belong to different node groups and do not allow partition swap if either of them is being redistributed.

Procedure

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Clusters > Dedicated Clusters**. All clusters are displayed by default.
- Step 3** In the **Operation** column of the target cluster, choose **More > Scale Node > Redistribute**, as shown in the following figure.
- Step 4** On the **Redistribute** page that is displayed, set **Advanced** to **Custom**, set the redistribution mode to **Online mode**, and click **Next: Confirm** to submit the task.

----End

6.3.2.2 Viewing Redistribution Details

On the **View Redistribution Details** page, you can check the monitoring information, including the redistribution mode, redistribution progress, and table redistribution details of the current cluster. You can pause and resume redistribution, set the redistribution priority, and change the number of concurrent redistribution tasks.

NOTE

The function of viewing redistribution details is supported by 8.1.1.200 and later cluster versions. Details about the data table redistribution progress are supported only by 8.2.1 and later cluster versions.

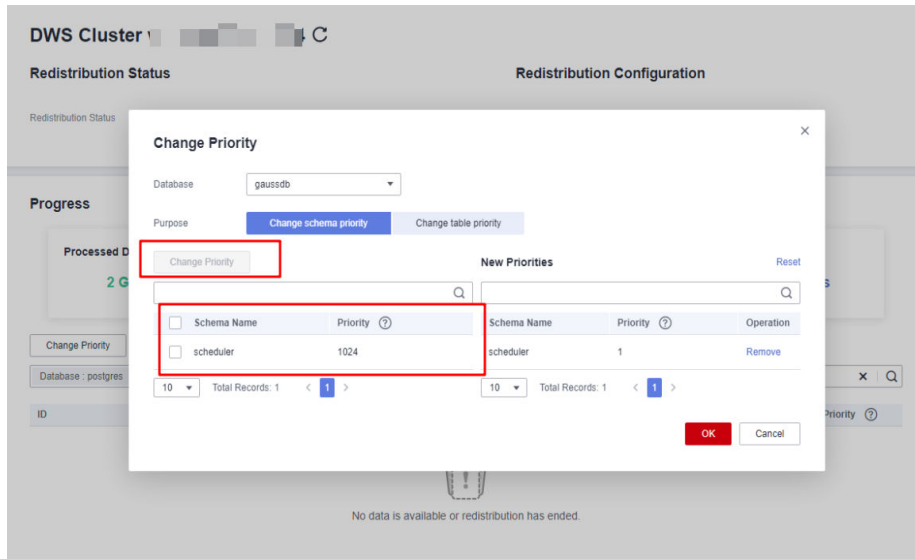
Precautions

You can check redistribution details only if the cluster is being redistributed, failed to be redistributed, or is suspended. There may be a delay in the statistics update.

Procedure

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Clusters > Dedicated Clusters**. All clusters are displayed by default.
- Step 3** In the **Task Information** column of a cluster, click **View Details**.
- Step 4** Check the redistribution status, configuration, progress, and redistribution details of all the tables in a specified database. Specify a database that can be searched by table redistribution status and table name. If all the tables in a database have completed redistribution, no data will be displayed for the database.
- Step 5** When redistribution is paused, you can set the redistribution priority (in schema or table dimension), and redistribution will be performed based on the configured

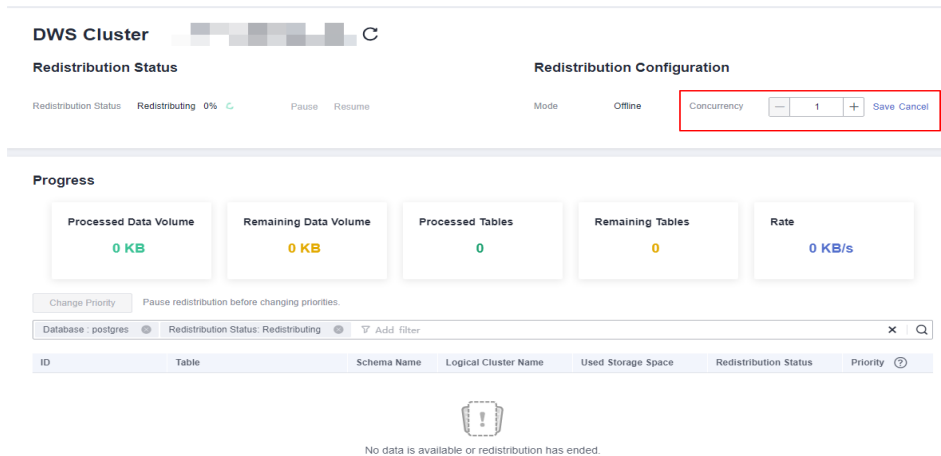
redistribution sequence. You can also set the redistribution priority before the redistribution starts.



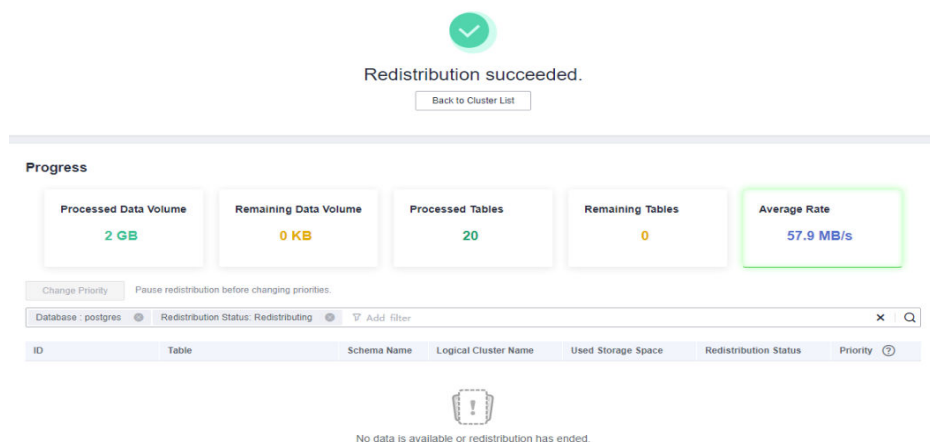
Step 6 The number of concurrent redistribution tasks can be adjusted during redistribution.

NOTE

Cluster 8.1.0 and earlier versions do not support dynamic adjustment. To change redistribution concurrency, suspend redistribution first.



Step 7 Check the redistribution progress. After the redistribution is complete, the amount of completed data, amount of remaining data, number of completed tables, number of remaining tables, and average rate during redistribution are displayed.



----End

6.3.3 Scaling In a Cluster

You can scale in your clusters on the console to release unnecessary compute and storage resources provided by GaussDB(DWS).

NOTE

- Scale-in is supported only by pay-per-use clusters of version 8.1.1.300 and later. For clusters billed in yearly/monthly mode, the function is supported only in version 8.2.1 and later.
- By default, scaled in nodes are charged by quantity.
- When you scale in a standard data warehouse cluster, you can only modify the same storage specifications as used by the cluster.
- A hybrid data warehouse (cluster mode) cannot be scaled in to a standalone cluster.

Impact on the System

- Before the scale-in, close the client connections that have created temporary tables, because temporary tables created before or during the scale-in will become invalid and operations performed on these temporary tables will fail. Temporary tables created after the scale-in will not be affected.
- If you start a scale-in, an automatic snapshot will be created for the cluster before scale-in. If you do not need the snapshot, you can disable the automated backup function on the scale-in page.
- Ensure that the skew rate is below 10% before scale-in and there is no specific guideline for the dirty page rate. However, it is advisable to maintain a skew rate of 20-30% for large tables exceeding 50 GB in size.
- When scaling in a cluster, several functions are disabled, including cluster restart, cluster scale-out, snapshot creation, node management, intelligent O&M, resource management, parameter modification, security configurations, log service, database administrator password resetting, and cluster deletion.
- During offline scale-in, stop all services or run only a few query statements. During table redistribution, a shared lock is added to tables. All insert, update, and delete operations as well as DDL operations on the tables are blocked for a long time, which may cause a lock wait timeout. After a table is

redistributed, you can access the table. Do not perform queries that take more than 20 minutes during the redistribution (the default time for applying for the write lock during redistribution is 20 minutes). Otherwise, data redistribution may fail due to lock wait timeout.

- During online scale-in, you can perform insert, update, and delete operations on tables, but data updates may still be blocked for a short period of time. Redistribution consumes lots of CPU and I/O resources, which will greatly impact job performance. Therefore, perform redistribution when services are stopped or during periods of light load.
- If a node is deleted while DDL statements are being executed (to create a schema or function) during online scale-in, errors may occur because the DN cannot be found. To resolve this issue, you can simply retry the statements.
- If a cluster scale-in fails, the database does not automatically roll back the scale-in operation, and no O&M operations can be performed. In this case, you need to click the **Scale In** on the console to try again.

Prerequisites

- The cluster is in **Available** state, is not read-only, and there is no data being redistributed in the cluster.
- A cluster configuration file has been generated, and configuration information is consistent with the current cluster configuration.
- Before the scale-in operation starts, the value of **default_storage_nodegroup** is **installation**.
- The cluster is configured in the ring mode. A ring is the smallest unit for scale-in. Four or five hosts form a ring. The primary, standby, and secondary DNs are deployed in this ring.
- The scale-in host does not contain the GTM, ETCD, or CM Server component.
- There are no CNs on the nodes to be scaled in.
- Scale-in does not support rollback but supports retry. A data redistribution failure after a scale-in does not affect services. You can complete scale-in at other appropriate time. Otherwise, unbalanced data distribution will persist for a long time.
- Before redistribution, ensure that the **data_redis** schema in the corresponding database is reserved for redistribution and that no user operation on it or its tables is allowed. During redistribution, **data_redis** is used. After the operation is complete, the schema will be deleted. User tables (if any) in the schema will also be deleted.
- **gs_cgroup** cannot be used during scale-in.
- Before the scale-in, check the remaining capacity of the cluster. The nodes remaining in a scale-in must have sufficient space to store the data of the entire cluster. Otherwise, the scale-in cannot be properly performed.
 - The used physical disk space on each node is less than 80%.
 - All the users and roles use less than 80% of resource quota in total.
 - The estimated space usage after scale-in must be less than 80%.
 - The available space is 1.5 times larger than the maximum size of a single table.
- Automatic removal of faulty CNs is disabled during the scale-in and is enabled after the scale-in is complete.

Procedure

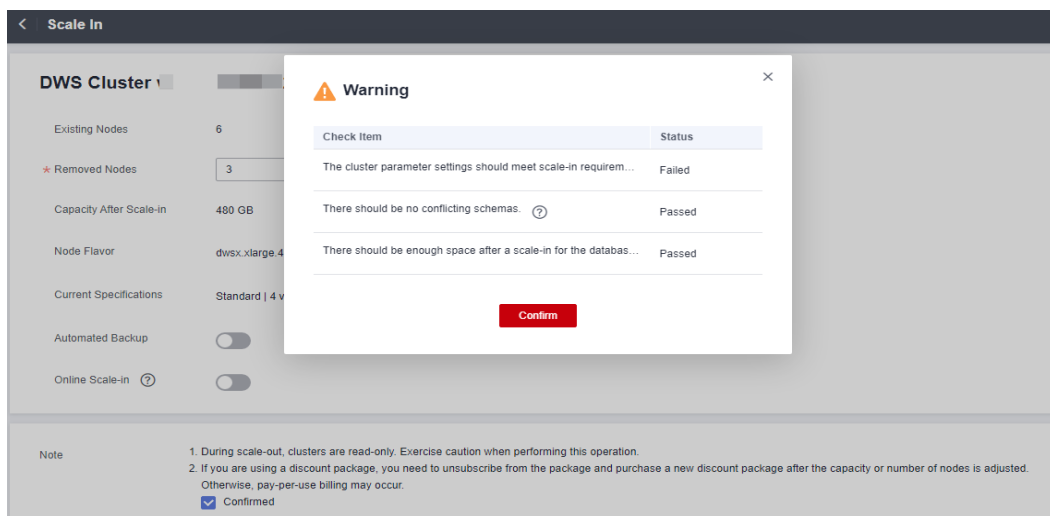
- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Clusters > Dedicated Clusters**.
- Step 3** In the **Operation** column of the target cluster, choose **More > Scale Node > Scale In**.
- Step 4** The scale-in page is displayed. You can select the number of nodes to be scaled in. The automated backup function is enabled by default.



NOTE

Before scaling in the cluster, it is crucial to verify if it meets the inspection conditions. Click **Immediate Inspection** to complete the inspection and proceed to the next step only if it passes. For more information, see [Viewing Inspection Results](#).

- Step 5** Click **Next: Confirm**. The system will check the cluster status before scale-in. If your cluster fails the check, an error message will be displayed.



Step 6 After the check is passed, click **Confirm** to return to the cluster list. The cluster status is **Scaling in**. Wait for a while.

----End

 **NOTE**

- If the cluster parameters fail the check, the scale-in will fail. To avoid this problem, ensure your parameter settings are correct.
- If schemas fail the check, the scale-in will fail. To avoid this problem, check whether any schema that conflicts with the scale-in exists.
- If the disk space fails the check, the scale-in may fail or the cluster may become read-only after the scale-in. To avoid this problem, increase your cluster disk capacity.

6.4 Changing Specifications

6.4.1 Using the Elastic Specification Change

Overview


Heavy service traffic requires additional resources (such as CPU, memory, and disk resources) to support it. If the current cluster resources are insufficient, creating a new cluster with more resources may be necessary. However, this can be costly and time-consuming. Moreover, creating a cluster with many resources but low service volume can result in resource redundancy and high costs.

The elastic flavor change function is introduced to tackle this problem. It is ideal for scenarios where computing capabilities (CPU and memory) need to be adjusted during peak hours or when only computing capabilities need to be changed. By using elastic flavor change before peak hours, the cluster's computing capability can be quickly increased. After peak hours, the cluster configuration can be reduced to minimize costs. For more information, see [Supported node flavors](#).

You can modify the CPU and memory configurations of the VM nodes in the target cluster by utilizing the underlying ECS capabilities. The following figure illustrates this process.

- To prevent service disruptions, it is crucial to schedule the elastic flavor change time window properly since the cluster must be stopped during the entire process.
- Changing all nodes concurrently ensures that the process will not take longer due to the number of nodes. Typically, the entire process takes around 5 to 10 minutes.

Figure 6-4 Principle of elastic flavor change



 **NOTE**

- Only cluster versions 8.1.1.300 and later support elastic flavor change. For an earlier version, contact technical support to upgrade it first.
- Elastic specification change is only supported by storage-compute coupled and decoupled clusters that use ECSs and EVS disks. Clusters with local ECSs do not have this capability.

Precautions

- Choosing a lower target flavor when decreasing a cluster's flavor can impact its performance, so it is crucial to assess the potential impact on services before proceeding with the operation.
- Make sure to check if there are enough ECS resources and tenant CPU quotas in the current region before modifying the flavors.
- You can change the flavors again if needed. In case the flavors of some nodes fail to change, you can resubmit the change task to execute the process.

Constraints and Limitations

- You can upgrade or downgrade ECS flavors of the same type. For instance, you can change from **dwsx2.2xlarge.m7** to **dwsx2.4xlarge.m7**, but not to **dwsx2.4xlarge.m6**.
- Stop the VM before changing the flavor. The flavor change can only be done offline and it takes 5 to 10 minutes.

Procedure

Step 1 Log in to the GaussDB(DWS) console.

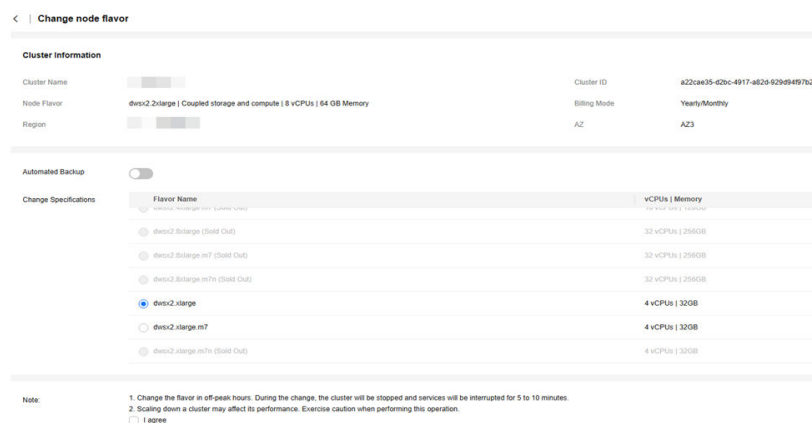
Step 2 Click **Dedicated Clusters** > **Clusters**. All clusters are displayed by default.

Step 3 In the row of a cluster, choose **More** > **Change Flavor** in the **Operation** column and click **Change Node Flavor**.

Step 4 Configure the flavor. Enable automatic backup as needed.

- (Default) You can switch the flavor to another flavor of the same type. For details about the supported flavors, see [Table 6-2](#).
- If you want to choose a different flavor, select one from the available options displayed on the page. In case you encounter any issues while changing the flavor, contact technical support for assistance.

Figure 6-5 Changing flavors



Step 5 Confirm the settings, select the confirmation check box, and click **Next: Confirm**.

Step 6 Click **Submit**.

Step 7 Return to the cluster list. The cluster status will change to **Changing node flavor**. Wait for about 5 to 10 minutes.

----End

Supported Flavors

Table 6-2 Supported node flavors

Current Flavor Name	Target Flavor Name
dwsk2.xlarge	dwsk2.2xlarge, dwsk2.4xlarge, dwsk2.12xlarge, dwsk2.8xlarge
dwsk2.2xlarge	dwsk2.12xlarge, dwsk2.8xlarge, dwsk2.4xlarge
dwsk2.4xlarge	dwsk2.2xlarge, dwsk2.8xlarge, dwsk2.12xlarge
dwsk2.8xlarge	dwsk2.2xlarge, dwsk2.4xlarge, dwsk2.12xlarge
dwsk2.12xlarge	dwsk2.2xlarge, dwsk2.4xlarge, dwsk2.8xlarge
dwsk2.h.12xlarge.4.kc1	dwsk2.h.xlarge.4.kc1, dwsk2.h.2xlarge.4.kc1, dwsk2.h.4xlarge.4.kc1, dwsk2.h.8xlarge.4.kc1
dwsk2.h.2xlarge.4.kc1	dwsk2.h.8xlarge.4.kc1, dwsk2.h.12xlarge.4.kc1, dwsk2.h.xlarge.4.kc1, dwsk2.h.4xlarge.4.kc1
dwsk2.h.4xlarge.4.kc1	dwsk2.h.8xlarge.4.kc1, dwsk2.h.12xlarge.4.kc1, dwsk2.h.xlarge.4.kc1, dwsk2.h.2xlarge.4.kc1
dwsk2.h.8xlarge.4.kc1	dwsk2.h.xlarge.4.kc1, dwsk2.h.2xlarge.4.kc1, dwsk2.h.4xlarge.4.kc1, dwsk2.h.12xlarge.4.kc1
dwsk2.h.xlarge.4.kc1	dwsk2.h.2xlarge.4.kc1, dwsk2.h.4xlarge.4.kc1, dwsk2.h.8xlarge.4.kc1, dwsk2.h.12xlarge.4.kc1

Current Flavor Name	Target Flavor Name
dwsk2.h1.12xlarge.4.kc1	dwsk2.h1.4xlarge.4.kc1, dwsk2.h1.8xlarge.4.kc1, dwsk2.h1.2xlarge.4.kc1
dwsk2.h1.2xlarge.4.kc1	dwsk2.h1.4xlarge.4.kc1, dwsk2.h1.8xlarge.4.kc1, dwsk2.h1.12xlarge.4.kc1
dwsk2.h1.4xlarge.4.kc1	dwsk2.h1.8xlarge.4.kc1, dwsk2.h1.12xlarge.4.kc1, dwsk2.h1.2xlarge.4.kc1
dwsk2.h1.8xlarge.4.kc1	dwsk2.h1.4xlarge.4.kc1, dwsk2.h1.12xlarge.4.kc1, dwsk2.h1.2xlarge.4.kc1
dwsk2.h1.xlarge.2.kc1	dwsk2.h1.2xlarge.4.kc1, dwsk2.h1.4xlarge.4.kc1, dwsk2.h1.8xlarge.4.kc1, dwsk2.h1.12xlarge.4.kc1
dwsx2.xlarge	dwsx2.2xlarge, dwsx2.4xlarge, dwsx2.8xlarge, dwsx2.16xlarge
dwsx2.2xlarge	dwsx2.4xlarge, dwsx2.8xlarge, dwsx2.16xlarge
dwsx2.4xlarge	dwsx2.2xlarge, dwsx2.8xlarge, dwsx2.16xlarge
dwsx2.8xlarge	dwsx2.2xlarge, dwsx2.4xlarge, dwsx2.16xlarge
dwsx2.16xlarge	dwsx2.2xlarge, dwsx2.4xlarge, dwsx2.8xlarge
dwsx2.xlarge.m7	dwsx2.2xlarge.m7, dwsx2.4xlarge.m7, dwsx2.8xlarge.m7, dwsx2.16xlarge.m7
dwsx2.2xlarge.m7	dwsx2.4xlarge.m7, dwsx2.8xlarge.m7, dwsx2.16xlarge.m7
dwsx2.4xlarge.m7	dwsx2.2xlarge.m7, dwsx2.8xlarge.m7, dwsx2.16xlarge.m7
dwsx2.8xlarge.m7	dwsx2.2xlarge.m7, dwsx2.4xlarge.m7, dwsx2.16xlarge.m7
dwsx2.16xlarge.m7	dwsx2.2xlarge.m7, dwsx2.4xlarge.m7, dwsx2.8xlarge.m7
dwsx2.xlarge.m7n	dwsx2.2xlarge.m7n, dwsx2.8xlarge.m7n, dwsx2.16xlarge.m7n
dwsx2.2xlarge.m7n	dwsx2.8xlarge.m7n, dwsx2.16xlarge.m7n
dwsx2.8xlarge.m7n	dwsx2.2xlarge.m7n, dwsx2.16xlarge.m7n
dwsx2.16xlarge.m7n	dwsx2.2xlarge.m7n, dwsx2.8xlarge.m7n
dwsx2.h.xlarge.4.c6	dwsx2.h.2xlarge.4.c6, dwsx2.h.4xlarge.4.c6, dwsx2.h.8xlarge.4.c6, dwsx2.h.16xlarge.4.c6
dwsx2.h.2xlarge.4.c6	dwsx2.h.4xlarge.4.c6, dwsx2.h.8xlarge.4.c6, dwsx2.h.16xlarge.4.c6

Current Flavor Name	Target Flavor Name
dwsx2.h.4xlarge.4.c6	dwsx2.h.8xlarge.4.c6, dwsx2.h.16xlarge.4.c6, dwsx2.h.2xlarge.4.c6
dwsx2.h.8xlarge.4.c6	dwsx2.h.4xlarge.4.c6, dwsx2.h.16xlarge.4.c6, dwsx2.h.2xlarge.4.c6
dwsx2.h.16xlarge.4.c6	dwsx2.h.2xlarge.4.c6, dwsx2.h.4xlarge.4.c6, dwsx2.h.8xlarge.4.c6
dwsx2.h.xlarge.4.c7	dwsx2.h.4xlarge.4.c7, dwsx2.h.8xlarge.4.c7, dwsx2.h.16xlarge.4.c7, dwsx2.h.2xlarge.4.c7
dwsx2.h.2xlarge.4.c7	dwsx2.h.4xlarge.4.c7, dwsx2.h.8xlarge.4.c7, dwsx2.h.16xlarge.4.c7
dwsx2.h.4xlarge.4.c7	dwsx2.h.2xlarge.4.c7, dwsx2.h.8xlarge.4.c7, dwsx2.h.16xlarge.4.c7
dwsx2.h.8xlarge.4.c7	dwsx2.h.16xlarge.4.c7, dwsx2.h.2xlarge.4.c7, dwsx2.h.4xlarge.4.c7
dwsx2.h.16xlarge.4.c7	dwsx2.h.8xlarge.4.c7, dwsx2.h.xlarge.4.c7, dwsx2.h.2xlarge.4.c7, dwsx2.h.4xlarge.4.c7
dwsx2.h.xlarge.4.c7n	dwsx2.h.2xlarge.4.c7n, dwsx2.h.4xlarge.4.c7n, dwsx2.h.8xlarge.4.c7n, dwsx2.h.16xlarge.4.c7n
dwsx2.h.2xlarge.4.c7n	dwsx2.h.4xlarge.4.c7n, dwsx2.h.8xlarge.4.c7n, dwsx2.h.16xlarge.4.c7n
dwsx2.h.4xlarge.4.c7n	dwsx2.h.2xlarge.4.c7n, dwsx2.h.8xlarge.4.c7n, dwsx2.h.16xlarge.4.c7n
dwsx2.h.8xlarge.4.c7n	dwsx2.h.16xlarge.4.c7n, dwsx2.h.2xlarge.4.c7n, dwsx2.h.4xlarge.4.c7n
dwsx2.h.16xlarge.4.c7n	dwsx2.h.4xlarge.4.c7n, dwsx2.h.8xlarge.4.c7n, dwsx2.h.2xlarge.4.c7n
dwsx2.h1.xlarge.2.c6	dwsx2.h1.8xlarge.4.c6, dwsx2.h1.16xlarge.4.c6, dwsx2.h1.2xlarge.4.c6, dwsx2.h1.4xlarge.4.c6
dwsx2.h1.2xlarge.4.c6	dwsx2.h1.4xlarge.4.c6, dwsx2.h1.8xlarge.4.c6, dwsx2.h1.16xlarge.4.c6
dwsx2.h1.4xlarge.4.c6	dwsx2.h1.2xlarge.4.c6, dwsx2.h1.8xlarge.4.c6, dwsx2.h1.16xlarge.4.c6
dwsx2.h1.8xlarge.4.c6	dwsx2.h1.16xlarge.4.c6, dwsx2.h1.4xlarge.4.c6, dwsx2.h1.2xlarge.4.c6
dwsx2.h1.16xlarge.4.c6	dwsx2.h1.4xlarge.4.c6, dwsx2.h1.2xlarge.4.c6, dwsx2.h1.8xlarge.4.c6
dwsx2.h1.xlarge.2.c7	dwsx2.h1.4xlarge.4.c7, dwsx2.h1.8xlarge.4.c7, dwsx2.h1.16xlarge.4.c7, dwsx2.h1.2xlarge.4.c7

Current Flavor Name	Target Flavor Name
dwsx2.h1.16xlarge.4.c7	dwsx2.h1.4xlarge.4.c7, dwsx2.h1.8xlarge.4.c7, dwsx2.h1.2xlarge.4.c7
dwsx2.h1.2xlarge.4.c7	dwsx2.h1.4xlarge.4.c7, dwsx2.h1.8xlarge.4.c7, dwsx2.h1.16xlarge.4.c7
dwsx2.h1.4xlarge.4.c7	dwsx2.h1.2xlarge.4.c7, dwsx2.h1.8xlarge.4.c7, dwsx2.h1.16xlarge.4.c7
dwsx2.h1.8xlarge.4.c7	dwsx2.h1.4xlarge.4.c7, dwsx2.h1.2xlarge.4.c7, dwsx2.h1.16xlarge.4.c7
dwsx2.h1.xlarge.2.c7n	dwsx2.h1.2xlarge.4.c7n, dwsx2.h1.4xlarge.4.c7n, dwsx2.h1.8xlarge.4.c7n, dwsx2.h1.16xlarge.4.c7n
dwsx2.h1.2xlarge.4.c7n	dwsx2.h1.16xlarge.4.c7n, dwsx2.h1.4xlarge.4.c7n, dwsx2.h1.8xlarge.4.c7n
dwsx2.h1.4xlarge.4.c7n	dwsx2.h1.8xlarge.4.c7n, dwsx2.h1.16xlarge.4.c7n, dwsx2.h1.2xlarge.4.c7n
dwsx2.h1.8xlarge.4.c7n	dwsx2.h1.4xlarge.4.c7n, dwsx2.h1.16xlarge.4.c7n, dwsx2.h1.2xlarge.4.c7n
dwsx2.h1.16xlarge.4.c7n	dwsx2.h1.2xlarge.4.c7n, dwsx2.h1.4xlarge.4.c7n, dwsx2.h1.8xlarge.4.c7n
dwsx2.rt.xlarge.m7	dwsx2.rt.2xlarge.m7, dwsx2.rt.4xlarge.m7, dwsx2.rt.8xlarge.m7, dwsx2.rt.16xlarge.m7
dwsx2.rt.2xlarge.m7	dwsx2.rt.4xlarge.m7, dwsx2.rt.8xlarge.m7, dwsx2.rt.16xlarge.m7
dwsx2.rt.4xlarge.m7	dwsx2.rt.2xlarge.m7, dwsx2.rt.8xlarge.m7, dwsx2.rt.16xlarge.m7
dwsx2.rt.8xlarge.m7	dwsx2.rt.2xlarge.m7, dwsx2.rt.4xlarge.m7, dwsx2.rt.16xlarge.m7
dwsx2.rt.16xlarge.m7	dwsx2.rt.2xlarge.m7, dwsx2.rt.4xlarge.m7, dwsx2.rt.8xlarge.m7
dwsk2.rt.xlarge.km1	dwsk2.rt.2xlarge.km1, dwsk2.rt.4xlarge.km1, dwsk2.rt.8xlarge.km1, dwsk2.rt.12xlarge.km1
dwsk2.rt.2xlarge.km1	dwsk2.rt.4xlarge.km1, dwsk2.rt.8xlarge.km1, dwsk2.rt.12xlarge.km1
dwsk2.rt.4xlarge.km1	dwsk2.rt.2xlarge.km1, dwsk2.rt.8xlarge.km1, dwsk2.rt.12xlarge.km1
dwsk2.rt.8xlarge.km1	dwsk2.rt.2xlarge.km1, dwsk2.rt.4xlarge.km1, dwsk2.rt.12xlarge.km1
dwsk2.rt.12xlarge.km1	dwsk2.rt.2xlarge.km1, dwsk2.rt.4xlarge.km1, dwsk2.rt.8xlarge.km1

Current Flavor Name	Target Flavor Name
dwsx2.rt.xlarge.m7n	dwsx2.rt.2xlarge.m7n, dwsx2.rt.8xlarge.m7n, dwsx2.rt.16xlarge.m7n
dwsx2.rt.2xlarge.m7n	dwsx2.rt.8xlarge.m7n, dwsx2.rt.16xlarge.m7n
dwsx2.rt.8xlarge.m7n	dwsx2.rt.2xlarge.m7n, dwsx2.rt.16xlarge.m7n
dwsx2.rt.16xlarge.m7n	dwsx2.rt.2xlarge.m7n, dwsx2.rt.8xlarge.m7n
4U32G.4DPU	8U64G.8DPU, 16U128G.16DPU, 32U256G.32DPU, 64U512G.64DPU, 96U768G.96DPU, and 128U1024G.128DPU
8U64G.8DPU	16U128G.16DPU, 32U256G.32DPU, 64U512G.64DPU, 96U768G.96DPU, and 128U1024G.128DPU
16U128G.16DPU	8U64G.8DPU, 32U256G.32DPU, 64U512G.64DPU, 96U768G.96DPU, and 128U1024G.128DPU
32U256G.32DPU	8U64G.8DPU, 16U128G.16DPU, 64U512G.64DPU, 96U768G.96DPU, and 128U1024G.128DPU
64U512G.64DPU	8U64G.8DPU, 16U128G.16DPU, 32U256G.32DPU, 96U768G.96DPU, and 128U1024G.128DPU
96U768G.96DPU	8U64G.8DPU, 16U128G.16DPU, 32U256G.32DPU, 64U512G.64DPU, 96U768G.96DPU, and 128U1024G.128DPU
128U1024G.128DPU	8U64G.8DPU, 16U128G.16DPU, 32U256G.32DPU, 64U512G.64DPU, 96U768G.96DPU, and 128U1024G.128DPU

6.4.2 Changing All Specifications

If you want to change your cluster topology or capacity but the **Change node flavor** option is grayed out, you can select **Change all specifications** to increase or decrease the nodes and their capacities on the GaussDB(DWS) console. First, you need to configure the new specifications you want, and a cluster with these specifications will be created. Then, data will be migrated from the old cluster to the new one. In case you need to restore data, a full snapshot will be taken for the old cluster, and the old cluster will be retained for a period of time.

 NOTE

- To use this feature, contact technical support engineers to upgrade your version first.
- The stream data warehouse cluster and the hybrid data warehouse cluster deployed in stand-alone mode do not support changing all specifications.
- Currently, DWS 3.0 clusters do not support changing all specifications.
- The new cluster does not incur charges before the change completes. The old cluster enters the retention period and will not incur charges after the resizing completes.
- A cluster can have up to 240 nodes. The old and new clusters can have up to 480 nodes in total.
- The **Change all specifications** option do not support logical clusters.

Impact of Changing All Specifications

- Before the change, you need to exit the client connections that have created temporary tables, because temporary tables created before or during the change will become invalid and operations performed on these temporary tables will fail. The temporary tables created after the change are not affected.
- When all specifications are changed, the cluster status changes to **Read-only** during data redistribution. During the change, services may be blocked for a long time. You are advised to perform the change with the assistance of engineers to prevent services from being affected.
- After the specifications are changed, the private IP address changes, which should be updated for connection.
- After the specifications are changed, the domain name remains unchanged, and the IP address bound to the domain name is switched. During the switchover, the connection is interrupted for a short period of time. Therefore, avoid writing service statements in the switchover. If the service side uses a domain name for connection, you need to update the cache information corresponding to the domain name to prevent connection failure after the change.
- If an ELB is bound to the cluster, the connection address on the service side remains unchanged after the specifications are changed, while the internal server address of the ELB is changed to the new connection address.
- In case you need to restore data, a full snapshot will be taken for the old cluster (on condition that your cluster support snapshot creation). You can check it in the snapshot list and manually delete it if it is no longer necessary.
- During the change, the cluster is read-only, affecting intelligent O&M tasks. You are advised to start these tasks after the change or pause them before the change.

Prerequisites

- The cluster to be changed is in the **Available**, **Read-only**, or **Unbalanced** state.
- The number of nodes after resizing must be smaller than or equals the available node quotas, or the resizing will fail.
- The total capacity of the new cluster after the change must be at least 1.2 times greater than the used capacity of the old cluster.

- To perform the change in a cluster as an IAM user, ensure that the IAM user has permissions for VPC, EVC, and BMS.

Changing All Specifications

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters > Dedicated Clusters**. All clusters are displayed by default.

Step 3 In the row of a cluster, choose **More > Change Specifications** in the **Operation** column and click **Change all specifications**. The Change Classic page is displayed.

NOTE

Before the change, it is crucial to verify if it meets the inspection conditions. Click **Inspect** to complete the inspection and proceed to the next step only if it passes. For more information, see [Viewing Inspection Results](#).

- For the **Node Flavor** parameter, select a flavor. The VPC, subnet, and security group of the new cluster are the same as those of the original cluster.
- For the **Set to** parameter, set the number of nodes you want for the new cluster.

Step 4 (Optional) If the cluster storage can be modified, you can set the storage type and the available storage for each node.

NOTE

Extreme SSD and **Extreme SSD V2** can only be selected for clusters that use ECS for computing and EVS for storage.

Step 5 Read the nodes and select **Confirmed**. Click **Resize Cluster Now**.

Step 6 Click **Submit**.

- After the change request is submitted, **Task Status** of the cluster changes to **Changing all specifications**. The process will take several minutes.
- During the change, the cluster automatically restarts, and **Cluster Status** is **Unavailable** for a period of time. After the restart is complete, **Cluster Status** changes to **Available**. Data is redistributed during resizing. During the redistribution, **Cluster Status** is **Read-only**.
- The resizing succeeds only when **Cluster Status** is **Available** and the **Change all specifications** task in **Task Information** is complete. Then the cluster begins providing services.
- If **Change all specifications failed** is displayed, the cluster failed to be changed.
- If change fails, and a message requiring retry is displayed when you click **Resize**, the failure is probably caused by abnormal cluster status or network problems. In this case, contact technical support to troubleshoot the problem and try again.

----End

6.4.3 Disk Capacity Expansion of an EVS Cluster

Context

As customer services evolve, disk space often becomes the initial bottleneck. In scenarios where other resources are ample, the conventional scale-out process is not only time-consuming but also resource-inefficient. Disk capacity expansion can quickly increase storage without service interruption. You can expand the disk capacity when no other services are running. If the disk space is insufficient after the expansion, you can continue to expand the disk capacity. If the expansion fails, you can expand the disk capacity again.

NOTE

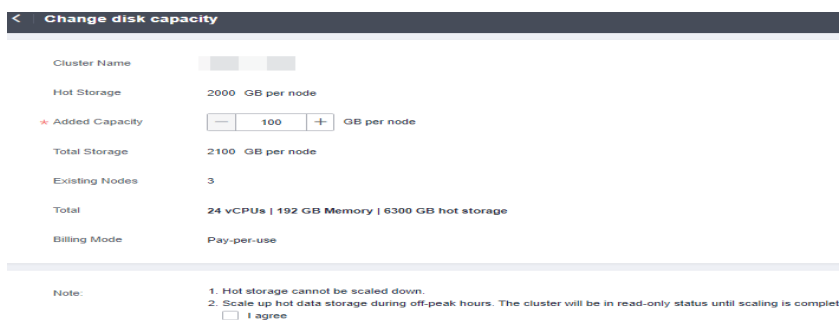
- Disk capacity expansion can be performed only for standard data warehouses using SSD, hybrid data warehouses, or stream data warehouses. Only version 8.1.1.203 and later are supported.
- Disk capacity can be expanded only if the cluster is in **Available**, **To be restarted**, **Read-only**, or **Node fault**, **Unbalanced** state.

Precautions

- Hot storage disks cannot be scaled down.
- Scale up hot data storage during off-peak hours.
- If the cluster is in the read-only state, a message will be displayed after you click **Expand Disk Capacity**. After you start expansion, wait until it is completed and the cluster changes to the available state.

Procedure

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Clusters > Dedicated Clusters**. All clusters are displayed by default.
- Step 3** In the **Operation** column of the target cluster, choose **More > Change Specifications** and click **Change disk capacity**. The **Expand Disk Capacity** page is displayed.
- Step 4** Set the capacity and click **Resize Cluster Now**.



Change disk capacity

Cluster Name

Hot Storage 2000 GB per node

Added Capacity GB per node

Total Storage 2100 GB per node

Existing Nodes 3

Total 24 vCPUs | 192 GB Memory | 6300 GB hot storage

Billing Mode Pay-per-use

Note:

- Hot storage cannot be scaled down.
- Scale up hot data storage during off-peak hours. The cluster will be in read-only status until scaling is complete.

I agree

NOTE

Hot Storage is changed to **Hot Storage (with Cache)** for DWS 3.0 clusters.

Step 5 Confirm the settings and click **Submit**.

Step 6 Return to the cluster list and check the disk capacity expansion progress.

----**End**

7 Backup and Disaster Recovery

7.1 Snapshots

7.1.1 Overview

A snapshot is a full or incremental backup of a GaussDB(DWS) cluster at a specific point in time. It records the current database data and cluster information, including the number of nodes, node specifications, and database administrator name. Snapshots can be created manually or automatically. For details, see [Manual Snapshots](#) and [Automated Snapshots](#).

If you restore a snapshot to a new cluster, GaussDB(DWS) creates a new cluster based on the cluster information recorded in the snapshot, and then restores data from the snapshot. For more information, see [Restoring a Snapshot to a New Cluster](#).

If you restore a snapshot to the original cluster, GaussDB(DWS) clears the existing data in the cluster, and then restores the database information from the snapshot to the cluster. For more information, see [Restoring a Snapshot to the Original Cluster](#).

The snapshot backup and restoration rates are listed below. (The rates are obtained from the test environment with local SSDs as the backup media. The rates are for reference only.) The actual rate depends on your disk, network, and bandwidth resources.)

- Backup rate: 200 MB/s/DN
- Restoration rate: 125 MB/s/DN

 NOTE

- OBS snapshot storage space and billing description
 - The cluster storage is provided by GaussDB(DWS) free of charge. Cluster storage = Storage space per node x Number of nodes
 - GaussDB(DWS) provides some free-of-charge storage space for you to store snapshot data generated in cluster backup. However, if you use more space than the free-of-charge storage space, the exceeded part is charged as per OBS billing rules. For details, see the [OBS pricing details](#).
- The dependency of the snapshot service is as follows:
 - The snapshot management function depends on OBS or NFS.
 - The backup device employs disk mounting mode for NFS backup media, reliant on the cloud-based SFS-Tubor service. For details, see [11.1.3.2 Automatic Snapshot Policy](#).
 - Only the snapshots stored in OBS can be used to restore data to a new cluster.
- The new GaussDB(DWS) cluster created based on the snapshot must have the same configurations as the original cluster. That is, the number and specifications of nodes, memory, and disks in the new cluster must be the same as those in the original cluster.
- If you create a new cluster based on a snapshot without modifying parameters, the parameters of the new cluster will be the same as those of the snapshot.
- The hybrid data warehouse (standalone) does not support snapshots.
- Only DWS 3.0 clusters of 9.0.2 and later versions support the snapshot function.
- During snapshot creation, do not perform the VACUUM FULL operation, or the cluster may become read-only.
- Snapshot creation affects disk I/O performance. You are advised to create snapshots during off-peak hours.
- During the snapshot creation, some intermediate files are retained, which occupy extra disk space. Therefore, create snapshots in off-peak hours and ensure that the disk capacity usage is less than 70%.

7.1.2 Manual Snapshots

7.1.2.1 Creating a Manual Snapshot

Prerequisites

A cluster snapshot is a complete backup that records point-in-time configuration data and service data of a GaussDB(DWS) cluster. This section describes how to create a snapshot on the [Snapshots](#) page to back up cluster data.

A manual snapshot can be created at any time. It will be retained until it is deleted from the GaussDB(DWS) console. Manual snapshots are full backup data, which takes a long time to create.

 NOTE

- Manual cluster snapshots can be backed up to OBS or NFS.
- To create a manual snapshot of a cluster, the cluster state must be **Available**, **To be restarted**, or **Unbalanced**. In cluster versions earlier than 8.1.3.101, you can also create a snapshot of a cluster in the **Read-only** state.

Impact on the System

If a snapshot is being created for a cluster, the cluster cannot be restarted, scaled, its password cannot be reset, and its configurations cannot be modified.

NOTE

To ensure the integrity of snapshot data, do not write data during snapshot creation.

Procedure

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Snapshots**.

Step 3 In the navigation pane, choose **Snapshots**. Click **Create Snapshot** in the upper right corner. Alternatively, choose **More > Create Snapshot** in the **Operation** column.

Step 4 Configure the following snapshot information:

- **Cluster Name:** Select a GaussDB(DWS) cluster from the drop-down list. The drop-down list only displays clusters that are in the **Available** state.
- **Snapshot Name:** Enter a snapshot name. The snapshot name must be 4 to 64 characters in length and start with a letter. It is case-insensitive and contains only letters, digits, hyphens (-), and underscores (_).
- **Snapshot Level:** Select **cluster**.
- **Snapshot Description:** Enter the snapshot information. This parameter is optional. Snapshot information contains 0 to 256 characters and does not support the following special characters: !<>'=&"

Step 5 Click **Create**.

Task status of the cluster for which you are creating a snapshot is **Creating snapshot**. The status of the snapshot that is being created is **Creating**. After the snapshot is created, its status changes to **Available**.

NOTE

If the snapshot size is much greater than that of the data stored in the cluster, the data is possibly labeled with a deletion tag, but is not cleared and reclaimed. In this case, clear the data and recreate a snapshot. For details, see [How Can I Clear and Reclaim the Storage Space?](#)

----End

7.1.2.2 Deleting a Manual Snapshot

On the **Snapshot Management** page of the GaussDB(DWS) management console, you can delete an unwanted snapshot in the **Unavailable** state or delete an available snapshot to release the storage space.

CAUTION

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

Procedure

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Snapshots**. All snapshots are displayed by default.

Step 3 In the **Operation** column of the snapshot that you want to delete, choose **More > Delete**.

 **NOTE**

You can only delete snapshots that were manually created.

Step 4 If the information is correct, enter **DELETE** and click **OK** to delete the snapshot.

----End

7.1.3 Automated Snapshots


7.1.3.1 Automatic Snapshot Overview

Automated snapshots adopt differential incremental backups. The automated snapshot created for the first time is a full backup (base version), and then the system creates full backups at a specified interval. Incremental backups are generated between two full backups. The incremental backup records change based on the previous backup.

During snapshot restoration, GaussDB(DWS) uses all backups between the latest full backup and the current incremental backup to restore the cluster. Therefore, no data loss occurs.

If the retention period of an incremental snapshot exceeds the maximum retention period, GaussDB(DWS) does not delete the snapshot immediately. Instead, GaussDB(DWS) retains it until the next full backup is completed, when the deletion of the snapshot will not hinder incremental data backup and restoration.

Figure 7-1 Snapshot backup process



Automated snapshots are enabled by default when you create a cluster. If automated snapshots are enabled for a cluster, GaussDB(DWS) periodically takes snapshots of that cluster based on the time and interval you set, usually every eight hours. You can configure one or more automated snapshot policies for the cluster as required. For details, see [Configuring an Automated Snapshot Policy](#).

The retention period of an automated snapshot can be set to 1 to 31 days. The default retention period is 3 days. The system deletes the snapshot at the end of

the retention period. If you want to keep an automated snapshot for a longer period, you can create a copy of it as a manual snapshot. The automated snapshot is retained until the end of the retention period, whereas the corresponding manual snapshot is retained until you manually delete it. For details about how to copy an automated snapshot, see [Copying Automated Snapshots](#).

7.1.3.2 Configuring an Automated Snapshot Policy

You can select a snapshot type and set one or more automated snapshot policies for a cluster. After an automated snapshot policy is enabled, the system automatically creates snapshots based on the time, period, and snapshot type you configured.

Procedure



- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane, choose **Clusters > Dedicated Clusters**.
- Step 3** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 4** Click the **Snapshots** tab page and click **Policy List**. All policies of the current cluster are displayed on the **Policy List** page. Toggle on **Snapshot Policy**.
 -  indicates that the policy is enabled (default). The default retention period is three days.
 -  indicates that the automatic snapshot function is disabled.
- Step 5** After this function is enabled, you can set the retention mode and the backup device used by the current cluster for automated snapshots. For more information, see [Table 7-1](#).

Table 7-1 Automated snapshot parameters

Parameter	Description
Retention Days	Retention days of the snapshots that are automatically created. The value ranges from 1 to 31 days. NOTE Snapshots that are automatically created cannot be deleted manually. The system automatically deletes these snapshots when their retention duration exceeds the threshold.
Backup Device	Select OBS and NFS from the drop-down list.
NFS Backup File System Address (NFS)	NFS shared IP address. Enter the IP address of the SFS shared path. After the mounting is successful, a mount directory is created in the <code>/var/chroot/nfsbackup</code> directory of the cluster instance by default.

Step 6 After automated snapshot is enabled, you can configure its parameters. For more information, see [Table 7-2](#).

 **NOTE**

The snapshot creation time is UTC, which may be different from your local time.

- If the snapshot type is set to **Full**, you can choose either **Periodic** or **One-time**, as shown in the following figures.
 - **Periodic**: Specify the days for every week/month and the exact time on the days.

×

Snapshot Policy

DWS provides certain free-of-charge space for you to store automated snapshots. If you use more space than the free-of-charge storage space the excessive part is charged in pay-per-use mode. [Learn more](#)

Snapshot Policy Name ?

Snapshot Type Full Incremental

Snapshot Policy ? Periodic One-off

Periodic Snapshot Settings

Days ? Specified day in week Specified date

1	2	3	4	5	6	7	8	9	10	11	12	13
14	15	16	17	18	19	20	21	22	23	24	25	26
27	28	29	30	31								

⚠ Select at least one of the preceding days.

All Clear

Time ? Specified time


Start time: ? (UTC)

 **WARNING**



Choosing the days in red (29th/30th/31st) may skip some monthly backups. Policy and execution depend on the specific month and date you choose.

- **One-time**: Specify a day and the exact time on the day.



Snapshot Policy  No full snapshot policy is configured for the current cluster. The default policy is used, that is, a full snapshot is taken every 14 incremental snapshots. You can set full snapshot policies as required.

Name 

Type Full Incremental

Policy Periodic  One-time 


One-time Policy Configurations


Time Create a backup at   UTC

Note: The UTC time is used by default. Set the policy based on the time zone and time difference as required.

- Incremental snapshots can be set only to **Periodic**, as shown in the first figure below.

When configuring a periodic incremental snapshot policy, you can specify the days for every week/month and the exact time on the days. You can also specify the start time and interval for the snapshots.



Snapshot Policy  No full snapshot policy is configured for the current cluster. The default policy is used, that is, a full snapshot is taken every 14 incremental snapshots. You can set full snapshot policies as required.

Name 



Type Full Incremental

Policy Periodic

Periodic Policy Configurations

Days Weekly  Monthly 

Sunday Monday Tuesday Wednesday Thursday Friday Saturday

Time Daily  Interval 

Create a backup at UTC

Note: The UTC time is used by default. Set the policy based on the time zone and time difference as required.

Table 7-2 Snapshot policy parameters

Parameter	Description
Name	The policy name must be unique, consist of 4 to 92 characters, and start with a letter. It is case-insensitive and can contain only letters, digits, hyphens (-), and underscores (_).

Parameter	Description
Type	You can choose either full or incremental snapshots. NOTE <ul style="list-style-type: none">• A full snapshot is created after every fifteen incremental snapshots are created.• Incremental snapshot restoration is based on full snapshots. Incremental snapshots are used to restore all data to the time point when they were created.• An incremental snapshot records the changes made after the previous snapshot was created. A full snapshot backs up the data of an entire cluster. It takes a short time to create an incremental snapshot, and a long time to create a full snapshot. When restoring a snapshot to a new cluster, GaussDB(DWS) uses all snapshots between the latest full backup and the current snapshot.
Policy	You can choose either periodic or one-time snapshots. NOTE One-time can be selected only for full snapshots.
One-time	You can create a full snapshot at a specified time in the future. The UTC time is used.
Periodic Policy Configurations	You can create automated snapshots on a daily, weekly, or monthly basis: <ul style="list-style-type: none">• Days: Specify days for every week or every month. Weekly and Monthly cannot be selected at the same time. If the current month does not contain the selected date, the operation will be postponed to the next month.• Time: Specify the exact time on the selected days. For incremental snapshots, you can specify the start time and interval. The interval can be 4 to 24 hours, indicating that a snapshot is created at an interval of 4 to 24 hours. NOTICE Incremental snapshots can be set only to Periodic , as shown in the first figure below.

Step 7 Click **OK**.

 **NOTE**

A maximum of three snapshot policies can be set for a cluster.

Step 8 (Optional) To modify an automated snapshot policy, click **Modify** in the **Operation** column.

Step 9 (Optional) To preview a policy, click **Preview Policy**. The next seven snapshots of the cluster will be displayed. If no full snapshot policy is configured for the cluster, the default policy is used, that is, a full snapshot is taken after every 15 incremental snapshots.

NOTICE

Implementation of the same policy varies according to operations in the cluster. For example:

- The policy preview time is for your reference only. The cluster triggers a snapshot within one hour before and after the preset time.
- The next automated snapshots after cluster scale-out, upgrade, resize, and media modification are full snapshots by default.
- If a periodic policy is used for a cluster, no automatic backup is allowed within 4 hours after the last automated snapshot is complete.
- If the time for triggering snapshots of multiple policies conflicts, the priorities of the policies are as follows: one-time > periodic > full > incremental.
- You can use any backup, full or incremental, to restore the full data of a resource.

----End

7.1.3.3 Copying Automated Snapshots

This section describes how to copy snapshots that are automatically created for long-term retention.

Copying an Automated Snapshot

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Snapshots**.

All snapshots are displayed by default. You can copy the snapshots that were automatically created.

Step 3 In the **Operation** column of the snapshot that you want to copy, choose **More > Copy**.

- **New Snapshot Name:** Enter a new snapshot name.

The snapshot name must be 4 to 64 characters in length and start with a letter. It is case-insensitive and contains only letters, digits, hyphens (-), and underscores (_).

- **Snapshot Description:** Enter the snapshot information.

This parameter is optional. Snapshot information contains 0 to 256 characters and does not support the following special characters: !<>='&"

Step 4 Click **OK**. The system starts to copy the snapshot for the cluster.

The system displays a message indicating that the snapshot is successfully copied and delivered. After the snapshot is copied, the status of the copied snapshot is **Available**.

 NOTE

If the snapshot size is much greater than that of the data stored in the cluster, the data is possibly labeled with a deletion tag, but is not cleared and reclaimed. In this case, clear the data and recreate a snapshot. For details, see [How Can I Clear and Reclaim the Storage Space?](#)

----End

7.1.3.4 Deleting an Automated Snapshot

Only GaussDB(DWS) can delete automated snapshots; you cannot delete them manually.

GaussDB(DWS) deletes an automated snapshot if:

- The retention period of the snapshot ends.
- The cluster is deleted.

 CAUTION

To help users restore a cluster deleted by mistake, GaussDB(DWS) provides the following policies (supported only in 8.2.0 and later) for cluster snapshots:

- If the latest snapshot is an automated snapshot, it will be retained for one day.
 - If the latest snapshot is a manual snapshot, the automated snapshot of the cluster will be deleted.
-

7.1.4 Viewing Snapshot Information

This section describes how to view snapshot information on the **Snapshots** page.


Viewing Snapshot Information

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane on the left, choose **Snapshots**.

In the snapshot list, all snapshots are displayed by default.

Step 3 You can view the **Snapshot Name, Snapshot Status, Cluster Name, Backup Mode, Snapshot Type, Storage Medium**, and creation time of snapshots.

You can also enter a snapshot name or cluster name in the upper right corner of the snapshot list and click  to search for the specified snapshot. GaussDB(DWS) supports fuzzy search.

[Table 7-3](#) describes snapshot status.

Table 7-3 Snapshot status

Status	Description
Available	Indicates that the existing snapshot works properly.
Creating	Indicates that a snapshot is being created.
Unavailable	Indicates that the existing snapshot cannot provide services.

[Table 7-4](#) lists the backup modes.

Table 7-4 Backup modes

Type	Description
Manual	Indicates the snapshot that you manually create through the GaussDB(DWS) management console or using APIs. You can delete the snapshots that are manually created.
Automated	Indicates the snapshot that is automatically created after the automated snapshot backup policy is enabled. You cannot delete the snapshots that are automatically created. The system automatically deletes the snapshots whose retention duration expires.

The following table describes the snapshot types.

Table 7-5 Type

Type	Description
Full	The snapshot is a full backup.
Incremental	The snapshot is an incremental backup.

The following table describes the snapshot media.

Table 7-6 Storage media

Storage Medium	Description
OBS	The created snapshot is an OBS snapshot and the backup data is stored on the OBS server.
NFS	The created snapshot is an NFS snapshot and the backup data is stored on the NFS server.

----End

7.1.5 Restoration Using a Snapshot

7.1.5.1 Constraints on Restoring a Snapshot

Cluster-Level Snapshot Restoration

Cluster-level restoration consists of two steps:

1. Data restoration: Restores data in the backup set to the data directory of each primary DN/CN instance in parallel.
2. Rebuilding the standby DN: After the primary DN is restored, standby DNs are rebuilt with full data in parallel.

NOTE

- The restoration process takes 1.5 to 2 times longer than the backup process.
- The parameters after cluster-level restoration are the same as those before backup. When restoring data to a new cluster, ensure that the flavor of the new cluster is the same as that of the original cluster. If the flavor of the new cluster is smaller, the restoration may fail.

7.1.5.2 Restoring a Snapshot to a New Cluster

Scenario

This section describes how to restore a snapshot to a new cluster when you want to check point-in-time snapshot data of the cluster.

When a snapshot is restored to a new cluster, the restoration time is determined by the amount of data backed up by the snapshot. If a snapshot contains a large amount of data, the restoration will be slow. A small snapshot can be quickly restored.

Automatic snapshots are incremental backups. When restoring a snapshot to a new cluster, GaussDB(DWS) uses all snapshots between the latest full backup and the current snapshot. You can set the backup frequency. If snapshots are backed up only once a week, the backup will be slow if the incremental data volume is large. You are advised to increase the backup frequency.

NOTICE

- Currently, you can only use the snapshots stored in OBS to restore data to a new cluster.
- By default, the new cluster created during restoration has the same specifications and node quantity as the original cluster.
- Restoring data to a new cluster does not affect the services running in the original cluster.
- If cold and hot tables are used, snapshots cannot be used to restore cold data to a new cluster.
- Fine-grained restoration does not support tables in absolute or relative tablespace.
- Logical clusters and resource pools cannot be restored to a new cluster.

Prerequisites

- The resources required for restoring data to a new cluster do not exceed your available resource quota.
- The snapshot is in the **Available** state.

Procedure

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Snapshots**. All snapshots are displayed by default.

Step 3 In the **Operation** column of a snapshot, click **Restore**.

Step 4 On the **Restore Snapshot** page, configure the parameters of the new cluster, as shown in the following figure. **Extreme SSD** and **Extreme SSD V2** can only be selected for clusters that use ECS for computing and EVS for storage.

- Restore to a single-AZ cluster.
- Restore to a multi-AZ cluster.

NOTE

- Only clusters later than 8.2.0.100 can be restored to a multi-AZ cluster.
- Currently, the multi-AZ cluster supports only GaussDB(DWS) 2.0 standard data warehouses.
- The number of AZs in the current region is greater than or equal to 3.
- The number of nodes and CNs must be a multiple of 3.
- DNs in the multi-AZ cluster must be less than or equal to 2.

You can modify cluster parameters. For details, see [Table 7-7](#). By default, other parameters are the same as those in the snapshot. For details, see [Table 7-2](#).

Table 7-7 Parameters for the new cluster

Category	Operation
Basic settings	Region, AZ, node flavor, cluster name, database port, VPC, subnet, security group, public access, and enterprise project
Advanced settings	If Custom is selected, configure the following parameters: <ul style="list-style-type: none">• Tag: If encryption is enabled for the original cluster, you can configure a key name.

Step 5 Click **Restore** to go to the confirmation page.

Step 6 Click **Submit** to restore the snapshot to the new cluster.

When the status of the new cluster changes to **Available**, the snapshot is restored.

After the snapshot is restored, the private network address and EIP (if **EIP** is set to **Buy now**) are automatically assigned.

 **NOTE**

If the number of requested nodes, vCPU (cores), or memory (GB) exceed the user's remaining quota, a warning dialog box is displayed, indicating that the quota is insufficient and displaying the detailed remaining quota and the current quota application. You can click **Increase quota** in the warning dialog box to submit a service ticket and apply for higher node quota. Once approved, we will update your resource quota accordingly and send you a notification.

----End

Viewing Restoration Details

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters > Dedicated Clusters**.

Step 3 In the cluster list, if the cluster status is **Restoring**, click **View Details**.

Step 4 You can view the snapshot restoration progress of the cluster on the task details page.

 **NOTE**

- The estimated duration in the task details is for reference only. The actual duration depends on the current data volume.
- In the restore phase, click **View** to view the kernel restoration process. Note that there may be a time gap between the task time displayed in the task details area and the actual kernel execution time due to task scheduling and restart.

----End

7.1.5.3 Restoring a Snapshot to the Original Cluster

Scenario

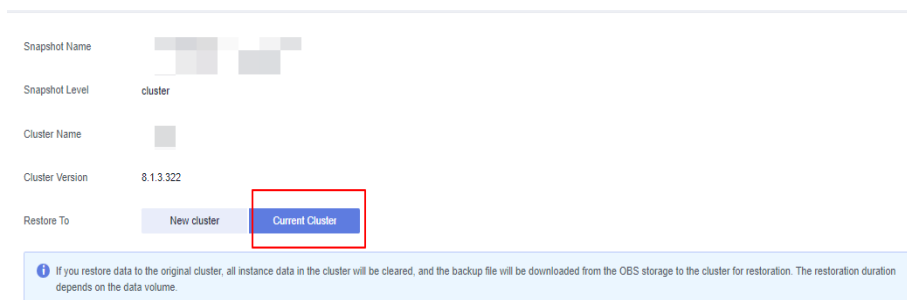
You can use a snapshot to restore data to the original cluster. This function is used when a cluster is faulty or data needs to be rolled back to a specified snapshot version.

NOTICE

- This function is supported only by clusters of version 8.1.3.200 or later.
- Snapshots whose backup device is OBS can be backed up.
- Only a snapshot in the **Available** state can be used for restoration.
- Logical clusters and resource pools cannot be restored to the current cluster.

Procedure

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane, choose **Clusters > Dedicated Clusters**.
- Step 3** Click the name of a cluster and choose **Snapshots**.
- Step 4** Click **Restore**.
- Step 5** Restore the snapshot to the current cluster.



NOTE

If you use a snapshot to restore data to the original cluster, the cluster will be unavailable during the restoration.

----End

7.1.6 Configuring a Snapshot

You can configure the parameters for creating and restoring a snapshot.

 **NOTE**

- This feature applies only to clusters of 8.2.0 or later. (For clusters of versions earlier than 8.2.0, only some parameters can be configured.)
- The parameters take effect on all the snapshot creation and restoration tasks.

Procedure

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane, choose **Clusters** > **Dedicated Clusters**.
- Step 3** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 4** Click the **Snapshots** tab page and click **Configure Parameters**. All the configurable parameters of the current cluster will be displayed.
- Step 5** Configure parameters as required. For details, see [Table 7-8](#).
- Step 6** Click **Save**.

----End

Snapshot parameters

Table 7-8 Snapshot information

Parameter	Type	Description	Default Value
parallel-process	Backup parameter	Number of concurrent processes on each node during Roach backup. NOTE This parameter can be configured for clusters earlier than 8.2.0.	The value is the number of DNs on the current node.
compression-type	Backup parameter	Compression algorithm. <ul style="list-style-type: none">• zlib• LZ4 NOTE This parameter can be configured for clusters earlier than 8.2.0.	LZ4

Parameter	Type	Description	Default Value
compression-level	Backup parameter	Compression level. The value range is 0 to 9. <ul style="list-style-type: none">● 0: fast backup and no compression● 9: slow backup and maximum compression NOTE This parameter can be configured for clusters earlier than 8.2.0.	6
buffer-size	Backup parameter	Buffer size of the Roach upload media. The value range is 256 to 16,384, in MB.	256
buffer-block-size	Backup parameter	Data block size of the data file to be read by Roach. The value range is 5,242,880 to 268,435,456, in bytes.	67108864
cpu-cores	Backup parameter	Number of CPU cores that can be used when Roach starts multiple threads concurrently	1/2 of the total number of logical CPU cores on the node
master-timeout	Backup parameter	Timeout period for the communication between the Roach master and agent nodes. The value range is 600 to 3600, in seconds.	3600
max-backup-io-speed	Backup parameter	I/O flow control during Roach backup. The value range is 0 to 2048, in MB/s. The value must be greater than the value of buffer-block-size . The value 0 indicates no limit.	0
backup-mode	Backup parameter	Full backup mode. <ul style="list-style-type: none">● 0: phase-1 backup● 1: phase-2 backup	0

Parameter	Type	Description	Default Value
cbm-parse-mode	Backup parameter	Incremental backup mode. <ul style="list-style-type: none">● 0: one-time CBM scan (high memory usage and high performance)● 1: multiple CBM scans (stable memory usage and low performance)	0
parallel-process	Restoration parameter	Number of concurrent processes on each node during Roach backup. By default, the value is the number of primary DNS on the current node plus 1.	1
cpu-cores	Restoration parameter	Number of CPU cores that can be used when Roach starts multiple threads concurrently	The default value is 1/2 of the number of CPU cores.
logging-level	Restoration parameter	Log levels: <ul style="list-style-type: none">● FATAL: Unrecoverable faults that cause the system suspension. This is the most severe level.● ERROR: Major errors.● WARNING: Exceptions. In this case, the system may continue to process tasks.● INFO: Notes.● DEBUG: Debugging details.● DEBUG2: Detailed debugging information, which is generally not displayed. This is the least severe level.	INFO

7.1.7 Stopping Snapshot Creation

You can stop snapshot creation on the **Snapshots** page.

NOTE

- This feature is supported only in version 8.1.3.200 and later.
- If the snapshot is ready to complete, the command for stopping the snapshot will not take effect and the snapshot will end normally.

Precautions

Only the snapshots in the **Creating** state can be stopped. A snapshot creation task that just started or is about to complete cannot be stopped.

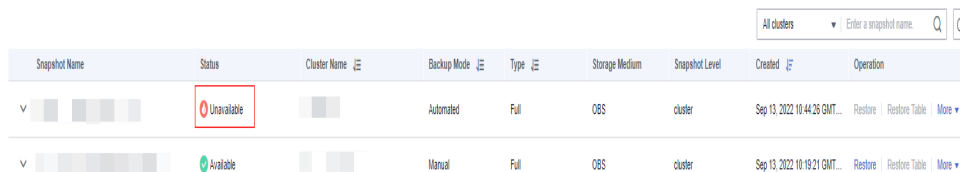
Procedure

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane on the left, choose **Snapshots**.

In the **Operation** column of a snapshot that is being created, and click **Cancel Creation**.

Step 3 In the dialog box that is displayed, click **Yes** to stop the snapshot. The snapshot state will change to **Unavailable**.



Snapshot Name	Status	Cluster Name	Backup Mode	Type	Storage Medium	Snapshot Level	Created	Operation
	Unavailable		Automated	Full	OBS	cluster	Sep 13, 2022 10:44:26 GMT...	Restore Restore Table More
	Available		Manual	Full	OBS	cluster	Sep 13, 2022 10:19:21 GMT...	Restore Restore Table More

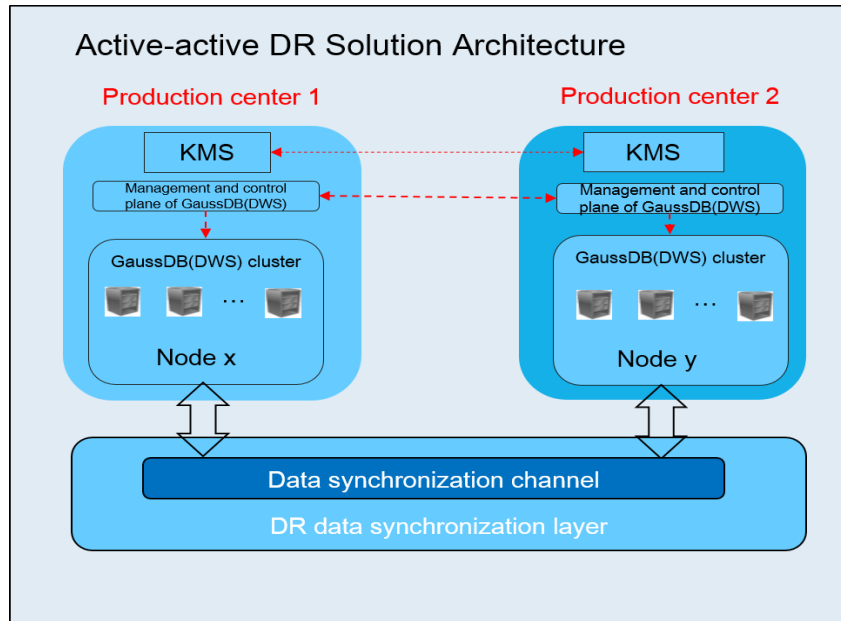
----End

7.2 Cluster DR

7.2.1 DR Overview

Overview

A homogeneous GaussDB(DWS) disaster recovery (DR) cluster is deployed in the same region. If the production cluster fails to provide read and write services due to natural disasters in the specified region or cluster internal faults, the DR cluster becomes the production cluster to ensure service continuity. The following figure shows the architecture.



NOTE

- Intra-region DR is supported only in cluster version 8.1.1 and later.
- The hybrid data warehouse (standalone) does not support disaster recovery.
- Multi-AZ clusters and DWS 3.0 clusters do not support DR.

DR Features

- Multi-form DR
 - Intra-region DR
 - Multiple data synchronization modes: synchronization layer based on mutual trust
- Low TCO
 - Heterogeneous deployment (logical homogeneity)
 - Cluster-level DR
- Visual console
 - Automatic and one-click DR drills

Constraints and Limitations

- During data synchronization, a non-fine-grained DR cluster cannot provide read or write services.
- When the DR task is stopped or abnormal but the DR cluster is normal, the DR cluster can provide the read service. After the DR switchover is successful, the DR cluster can provide the read and write services.
- When the DR task is created, the snapshot function of the production cluster is normal, but that of the DR cluster is disabled. Besides, snapshot restoration of both clusters is disabled.
- Logical clusters are not supported.
- Resource pools are not supported.

- If cold and hot tables are used, cold data is synchronized using OBS.
- DR does not synchronize data from external sources.
- DR management refers to dual-cluster DR under the same tenant.
- The DR cluster and the production cluster must be logically homogeneous and in the same type and version.
- The production cluster and DR cluster used for intra-region DR must be in the same VPC.
- In intra-region DR, after services are switched over from the production cluster to the DR cluster, the bound ELB is automatically switched to the new production cluster. During the switchover, the connection is interrupted for a short period of time. Do not run service statements to write data during the switchover.
- During intra-region DR, the EIP, intranet domain name, and connection IP address of the original production cluster are not automatically switched with the cluster switchover. The EIP, domain name, or IP address used for connection in the service system need to be switched to the new cluster.

7.2.2 Creating a DR Task

Creating an Intra-Region Cluster-Level DR Task

Prerequisites

You can create a DR task only when the cluster is in the **Available** or **Unbalanced** state.

Procedure

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane on the left, choose **DR Tasks**.
- Step 3** On the displayed page, click **Create**.
- Step 4** Select the type and enter the name of the DR task to be created.
 - **Type: Intra-region DR**
 - **Name:** Enter 4 to 64 case-insensitive characters, starting with a letter. Only letters, digits, hyphens (-), and underscores (_) are allowed.
- Step 5** Configure the production cluster.
 - Select a created production cluster from the drop-down list.
 - After a production cluster is selected, the system automatically displays its AZ.
- Step 6** Configure the DR cluster.
 - Select the AZ associated with the region where the DR cluster resides.

NOTE

- The AZ of the DR cluster can be the same as that of the production cluster. In a 3-AZ cluster, any of the three AZs can be selected for DR.
- **Cluster Name:** Upon selecting an AZ for the DR cluster, the system will auto-filter DR clusters that fulfill the logical homogeneity criteria. Should there be

no qualifying DR clusters, you can click **Create DR Cluster** to create a DR cluster with the same configuration as the production cluster.

Step 7 Configure advanced parameters. Select **Default** to keep the default values of the advanced parameters. You can also select **Custom** to modify the values.

- The DR synchronization period indicates the interval for synchronizing incremental data from the production cluster to the DR cluster. Set this parameter based on the actual service data volume.

 **NOTE**

The default DR synchronization period is 30 minutes.

Step 8 Click **OK**.

The DR status will then change to **Creating**. Wait until the creation is complete, and the DR status will change to **Not Started**.

----End

7.2.3 Viewing DR Information

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane on the left, choose **DR Tasks**.

Step 3 In the DR list, click the name of a DR task.

On the page that is displayed, view the following information:

- **DR Information:** You can view the DR ID, DR name, DR creation time, and DR status.
- **Production Cluster Information:** You can view the production cluster ID, cluster name, AZ, used storage capacity, cluster DR status, and the time of the latest successful DR task.
- **DR Cluster Information:** You can view the DR cluster ID, cluster name, AZ, used storage capacity, cluster DR status, and the time of the latest successful DR task.
- **DR Configuration:** Users can view and modify the DR synchronization period.

----End

7.2.4 DR Management

Starting a DR Task

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane on the left, choose **DR Tasks**.

Step 3 Click **Start** in the **Operation** column of the target DR task.

Step 4 In the dialog box that is displayed, click **OK**.

The DR status will change to **Starting**. The process will take some time. After the task is started, the DR status will change to **Running**.

 **NOTE**

- You can start a DR task that is in the **Not started/Startup failed/Stopped** state.
- After you start the DR task, you cannot perform operations such as restoration, scale-out, upgrade, restart, node replacement, and password update, on the production cluster or DR cluster, and backup is also not allowed on the DR cluster. Exercise caution when performing this operation.
- After the DR task is started, if the DR cluster is running properly and DR recovery is in progress, the cluster will be billed.

----End

Stopping the DR Task

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane on the left, choose **DR Tasks**.

Step 3 Click **Stop** in the **Operation** column of the target DR task.

Step 4 In the dialog box that is displayed, click **OK**.

The DR status will change to **Stopping**. The process will take some time. After the DR task is stopped, the status will change to **Stopped**.

 **NOTE**

- Only DR tasks in the **Running** or **Stop failed** state can be stopped.
- Data cannot be synchronized after a DR task is stopped.

----End

Switching to the DR Cluster

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane on the left, choose **DR Tasks**.

Step 3 Click **Switch to DR Cluster** in the **Operation** column of the target DR task.

Step 4 In the dialog box that is displayed, click **OK**.

The DR status will change to **DR switching**.

After the switchover is successful, the DR status will change to the original status.

 **NOTE**

- To perform a switchover when the DR cluster is running properly, click **Switch to DR Cluster**.
- You can perform a DR switchover when the DR task is in the **Running** state.
- During a switchover, the original production cluster is not available.
- The Recovery Point Object (RPO, time point to which the system and data must be recovered after a disaster occurs) in different DR switchover scenarios is described as follows:
 - Production cluster in the **Available** state: RPO = 0
 - Production cluster in the **Unavailable** state: A zero RPO may not be achieved, but data can at least be restored to that of the latest successful DR synchronization (**Last DR Succeeded**). For details, see [Viewing DR Information](#).

----End

Exception Switchover

Scenario

The production cluster is unavailable, the DR cluster is normal, and the DR status is **Abnormal**.

Procedure

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane on the left, choose **DR Tasks**.
- Step 3** Choose **More > Exception Switchover** in the **Operation** column of the target DR task.
- Step 4** In the dialog box that is displayed, click **OK**.

The **Status** will change to **Switchover in progress**.

After the switchover is successful, the DR status will change to the original status. In this procedure, the DR status will change back to **Abnormal**.

 **NOTE**

- To perform a switchover when the DR cluster is abnormal or the production cluster is faulty, click **Exception Switchover**.
- DR exception switchover is supported only by clusters of version 8.1.2 or later.
- Before a switchover, check the latest synchronization time in the DR cluster. The DR cluster will serve as a production cluster after an abnormal switchover, but the data that failed to be synchronized from the original production cluster to the DR cluster will not exist in the DR cluster.

----End

Performing a DR Switchback

Scenario

After abnormal switchover, if you have confirmed that the original production cluster was recovered, you can perform a switchback.

Procedure

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane on the left, choose **DR Tasks**.
- Step 3** Click **DR Recovery** in the **Operation** column of a DR task.
- Step 4** In the displayed dialog box, set **Synchronization Mode** to **Incremental** or **Full**.

NOTE

You are advised to set **Synchronization Mode** to **Incremental** when updating a DR creation task.

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

The **Status** will change to **Recovering**.

After the DR recovery is successful, the **Status** will change to **Running**.

NOTE

- DR is supported only by clusters of 8.1.2 or later.
- During DR recovery, data in the DR cluster will be deleted, and the DR relationship will be re-established with the new production cluster.

----End

Updating DR Configurations

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane on the left, choose **DR Tasks**.
- Step 3** In the DR list, click the DR name to go to the DR information page.
- Step 4** In the **DR Configurations** area, click **Modify**.

NOTE

- Only DR tasks in the **Not started** or **Stopped** state can be modified.
- The new configuration takes effect after DR is restarted.

----End

Deleting DR Tasks

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane on the left, choose **DR Tasks**.
- Step 3** Click **Delete** in the **Operation** column of the target DR task.
- Step 4** In the dialog box that is displayed, click **OK**.

The DR status will change to **Deleting**.

 NOTE

- You can delete a DR task when **DR Status** is **Creation failed**, **Not started**, **Startup failed**, **Stopped**, **Stop failed**, or **Abnormal**.
- Data cannot be synchronized after a DR task is deleted, and the deleted task cannot be restored.

----End

7.2.5 Mutually Exclusive DR Cases

Case 1: How Do I Scale out a Cluster in the DR State?

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Clusters** > **Dedicated Clusters**.

Step 3 In the cluster list, if **Task Information** of the cluster you want to scale out is **DR not started**, perform [Step 5](#) and [Step 7](#).

Step 4 (Optional) If the **Task Information** is other than **DR not started**, delete the DR task. For details, see [Deleting DR Tasks](#).

Step 5 In the **Operation** column of the production and DR clusters, choose **More** > **Scale Out**.

Step 6 Create a DR task. For details, see [Creating a DR Task](#).

Step 7 Start the DR task. For details, see [Starting a DR Task](#).

 NOTE

After scale-out, the number of DNs in the production cluster must be the same as that in the DR cluster.

----End

8 Intelligent O&M

8.1 Overview

GaussDB(DWS) provides the intelligent O&M feature to help users quickly and efficiently execute O&M tasks. Intelligent O&M selects a proper time window and concurrency to complete specified tasks based on the cluster load. During O&M tasks, intelligent O&M monitors user service changes and promptly adapts task execution policies to minimize the impact on user services. Periodic tasks and one-off tasks are supported, and you can configure the time window as required.

Intelligent O&M ensures high availability. When the cluster is abnormal, failed O&M tasks will be retried. If some steps of an O&M task cannot be completed due to an abnormal cluster, the failed steps will be skipped for cost saving.

The intelligent O&M page consists of the following parts:

- Common configuration of O&M tasks: Currently, you can only configure **Maximum number of concurrent O&M tasks in the VacuumFull user table**. This configuration takes effect on all the VACUUM FULL tasks of user tables.
- Information about ongoing O&M tasks. (Currently, only VACUUM tasks are displayed. If disk space is insufficient because of table bloating, you can vacuum tables.)
 - Frequent table creation and deletion can lead to table bloating. To free up space, you can run the **VACUUM** command on system catalogs.
 - Frequently update and delete operations can lead to table bloating. To free up space, you can run the **VACUUM** or **VACUUM FULL** command on system catalogs.
- O&M details: **O&M Plan** and **O&M Status**. **O&M Plan** displays the basic information about all O&M tasks, and **O&M Status** displays the running status.

 NOTE

- This feature is supported only in 8.1.3 or later.
- The intelligent O&M function is not supported in hybrid data warehouses (standalone mode).
- After completing the **VACUUM FULL** O&M task, the system automatically performs the **ANALYZE** operation.
- Only cluster 8.1.3 and later versions support the common configuration module for O&M tasks. For earlier versions, contact technical support to upgrade them.

8.2 Managing O&M Plans

Setting the Common Configurations of O&M Tasks

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Click the name of the target cluster.

Step 3 In the navigation pane, choose **Intelligent O&M**.

Step 4 In the **Common O&M Task Configuration** area, configure **Maximum number of concurrent O&M tasks in the VacuumFull user table**.

 NOTE

- This configuration takes effect for the VACUUM FULL O&M tasks of all user tables.
- The concurrency value range is 1 to 24. Configure it based on the remaining disk space and I/O load. You are advised to set it to 5.

----End

Adding an O&M Plan

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Click the name of the target cluster.

Step 3 In the navigation pane, choose **Intelligent O&M**.

Step 4 In the **O&M Plan** area, click **Add O&M Task**.

Step 5 In the displayed **Add O&M Task** dialog box, configure basic information about the O&M task.

Table 8-1 Basic configuration items of an O&M task

Configurat ion Item	Description	Example
O&M Task	Vacuum (Currently, only Vacuum O&M tasks are supported.)	Vacuum

Configuration Item	Description	Example
Description	Brief description of the intelligent O&M task.	This intelligent O&M task helps users periodically invoke Vacuum commands to reclaim space.
Remarks	Supplementary information.	-
Scheduling Mode	<p>The supports the following scheduling modes:</p> <ul style="list-style-type: none"> • Auto: Intelligent O&M scans the database in a specified time window, and automatically delivers table-level vacuum tasks by service load and reclaimable space of user tables. • Specify: You need to specify a vacuum target. Intelligent O&M will automatically deliver a table-level vacuum task in a specified time window. • Priority: You can specify the preferential vacuum targets. During the remaining time window (if any), Intelligent O&M will automatically scan other tables that can be vacuumed and deliver table-level vacuum tasks. <p>NOTE You are advised to select Specify for VACUUM and VACUUM FULL operations. Do not perform VACUUM FULL on wide column-store tables. Otherwise, memory bloat may occur.</p>	Specify

Configuration Item	Description	Example
Autovacuum	<p>Supported: system catalog Vacuum or user table VacuumFull.</p> <ul style="list-style-type: none">A system catalog VACUUM transaction holds a level-5 lock (share update exclusive lock), which does not affect user services. Only the transactions on the DDL process of the system catalog are blocked.A user table VACUUM FULL transaction holds a level-8 lock (access exclusive lock). All the other transactions on the table are blocked until VACUUM FULL is complete. To avoid affecting services, you are advised to perform VACUUM FULL during off-peak hours. <p>CAUTION During VACUUM FULL, the space usage will first increase and then decrease, because this operation requires the same space as the table to be vacuumed. (Actual table size = Total table size x (1 - dirty page rate). Ensure you have sufficient space before doing VACUUM FULL.</p>	User table (VACUUM FULL)
Vacuum First	You can configure the preferred Vacuum target. Each row corresponds to a table. Each table is represented by the database name, schema name, and table name, which are separated by spaces.	-
Advanced settings	<p>If you select Custom, you can configure the autovacuum triggers, including the table bloat and table reclaimable space. If you select Default, the default configuration is used.</p> <p>NOTE VACUUM bloat rate: After frequent UPDATE and DELETE operations are performed in a database, the deleted or updated rows are logically deleted from the database, but actually still exist in tables. Before VACUUM is complete, such data is still stored in disks, causing table bloat. If the bloat rate reaches the percentage threshold set in an O&M task, VACUUM will be automatically triggered.</p>	Default configuration (table bloat rate 80% or reclaimable space 100 GB.)

Step 6 Click **Next > Schedule** to configure scheduling for O&M tasks.

Select an O&M type.

- **One-off**: Set the start time and end time of the task.
- **Periodic**: Select a time window type, which includes **Daily**, **Weekly**, and **Monthly**, and select a time segment. Intelligent O&M will automatically analyze the time window and deliver O&M tasks accordingly.

 **CAUTION**

- Do not choose peak hours when configuring the time window for autovacuum O&M tasks. Otherwise, automatic Vacuum may cause a deadlock on user services.
- The number of concurrent O&M tasks (vacuum/vacuum full) ranges from 0 to 24 for user tables, and from 0 to 1 for system catalogs. The concurrency value cannot be customized, but can be automatically adjusted based on system **io_util**.
 - Two intervals for 0% to 60%
 - 0% to 30%: The concurrency value increases by 2 each time the value of **io_util** decreases by 15%.
 - 30% to 60%: The concurrency value is incremented by 1 each time the value of **io_util** decreases by 15%.
 - 60% to 70%: The concurrency value remains unchanged.
 - Above 70%: The concurrency value decreases by 1 until it reaches 0.
- The scheduler scans the expansion of column-store compression units (CUs) within the time window. If the average number of CU records in a column-store table is less than 1000, the scheduler scans the table first. The scanning of column-store CUs is not limited by table bloat or table reclaimable space.
- A maximum of 100 tables can be added to the priority list.
- The scheduler autovacuum function depends on the statistics. If the statistics are inaccurate, the execution sequence and results may be affected.
- The scheduler does not support names containing spaces or single quotation marks, including database names, schema names, and table names. Otherwise, the tables will be skipped. Priority tables whose name contains spaces or single quotation marks will also be skipped automatically.

Step 7 Click **Next: Finish**. After you confirm the information, click **Finish** to submit the request.

----End

Modifying an O&M Plan

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Click the name of the target cluster.
- Step 3** In the navigation pane, choose **Intelligent O&M**.
- Step 4** In the **O&M Plan** area, click **Modify** in the **Operation** column of the target task.

Task	Description	Time Window(UTC)	IS PAUSE	Task Type	Operation
<input type="checkbox"/>	Vacuum	12:45:00 - 13:15:00 UTC every day	No	Periodic	Modify Details
<input type="checkbox"/>	Vacuum	2021-09-22 12:47:41 to 2021-09-23 12:47:41	No	One-off	Modify Details
<input type="checkbox"/>	Vacuum	11:00:00 - 13:00:00 UTC every day	Yes	Periodic	Modify Details
<input type="checkbox"/>	Vacuum	07:00:00 - 14:00:00 UTC every day	Yes	Periodic	Modify Details

- Step 5** The **Modify O&M Task** panel is displayed. The configurations are similar to adding an O&M task (see [Adding an O&M Plan](#)).
 - Step 6** Confirm the modification and click **OK**.
- End

Viewing O&M Task Details

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Click the name of the target cluster.
- Step 3** In the navigation pane, choose **Intelligent O&M**.
- Step 4** In the **O&M Plan** area, click **Details** in the **Operation** column of the target task.

Task	Description	Time Window(UTC)	IS PAUSE	Task Type	Operation
<input type="checkbox"/>	Vacuum	12:45:00 - 13:15:00 UTC every day	No	Periodic	Modify Details
<input type="checkbox"/>	Vacuum	2021-09-22 12:47:41 to 2021-09-23 12:47:41	No	One-off	Modify Details
<input type="checkbox"/>	Vacuum	11:00:00 - 13:00:00 UTC every day	Yes	Periodic	Modify Details
<input type="checkbox"/>	Vacuum	07:00:00 - 14:00:00 UTC every day	Yes	Periodic	Modify Details

- Step 5** The **O&M Task Details** panel is displayed for you to check the information.
- End

8.3 Viewing O&M Tasks

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Click the name of the target cluster.
- Step 3** In the navigation pane, choose **Intelligent O&M**.
- Step 4** Switch to the **O&M Status** area.

Task	Status	Progress	Remaining Time Window	Time Window
No data available				

Step 5 Click the name of a specified O&M task to view the status details.

- **O&M Task: Vacuum**
- Status: **Waiting, Running, Completed, or Failed.**
- **Progress**
- **Remaining Time Window**
- **Time Window** (Local Time)
- **Tables Being Vacuumed**
- **Tables to Be Vacuumed**
- **Vacuumed Tables**
- **Failed Tables**

 **NOTE**

- A maximum of 100 tables can be displayed for each category of the tables above.
- If the cluster is read-only, the INSERT statement cannot be executed for intelligent O&M tasks. There may be tasks remaining in the **Running** status. The **Running** status in this case is a historical status, and it indicates that the task is not completed within the specified time. If you manually pause the task and the task is not scheduled, the task may remain in the **Waiting** status. In this case, cancel the cluster read-only state and contact technical support to update the task status.

----End

9 Cluster Management

9.1 Modifying Database Parameters

After a cluster is created, you can modify the cluster's database parameters as required. On the GaussDB(DWS) console, you can configure common database parameters. For details, see [Modifying Parameters](#). You can also view the parameter modification history. For details, see [Viewing Parameter Change History](#).

Prerequisites

You can modify parameters only when no task is running in the cluster.

Modifying Parameters

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 3** In the cluster list, find the target cluster and click the cluster name. The **Cluster Information** page is displayed.
- Step 4** Click the **Parameters** tab and modify the parameter values. Then click **Save**.

Parameter Name	CN Value	DN Value	Unit	Value Range	Restart Cluster ...	Description
ForceUDFMemoryLimit	4411000	44111	KB	0 - 2,147,483,647	No	Controls the virtual memory used by each fenced udf worker process. Default: 0.
UDFWorkMemHardLimit	1048570000	104857	KB	0 - 2,147,483,647	Yes	Specifies the maximum value of fencedUDFMemoryLimit. Unit: KB. Default: 104.
sqz_redistribute_enhancement	off	off	-	-	No	When the aggregate operation is performed, which contains multiple group by col...
alarm_report_interval	123110	100	Second	0 - 2,147,483,647	No	Specifies the interval at which an alarm is reported. Default: 10.
allocate_mem_cost	1.00011e-56	0	-	0 - 1.79709e+208	No	allow_concurrent_hop_update. Default: 0.
allow_concurrent_hop_update	on	on	-	-	No	Specifies whether to allow concurrent update. Default: on.
analysis_options	ALL_ON() ALL_VM_COMPILE OFF NA	ALL_ON() OFF VM_COMPILE HV	-	-	No	Specifies whether to enable function options in the corresponding options to use ...
archive_command		(disabled)	-	-	No	Specifies the command used to archive WALs set by the administrator. You are a...
archive_mode	on	off	-	-	No	Specifies whether to archive WALs. Default: off.
archive_timeout	0	12345611	Second	0 - 1,073,741,823	No	Specifies the archiving period. Default: 0.

Step 5 In the **Modification Preview** dialog box, confirm the modifications and click **Save**.

Step 6 You can determine whether you need to restart the cluster after parameter modification based on the **Restart Cluster** column.

Name	Value	Value Range	Restart Cluster	Description
password_encryption_type	1	0-2	No	Specifies the encryption type of user passwords. 0 indicates that passwords are encrypted in MD5 mode. 1 indicates...
timezone	UTC	--	No	Time zone that will be displayed in the timestamps. Default: UTC.
log_timezone	UTC	--	No	Time zone for timestamps in the server log. Default: UTC.

NOTE

- If cluster restart is not required for a parameter, the parameter modification takes effect immediately.
- If cluster restart is required for parameter modifications to take effect, the new parameter values will be displayed on the page after the modification, but will not take effect until the cluster is restarted. Before a restart, the cluster status is **To be restarted**, and some O&M operations are disabled.

----End

Viewing Parameter Change History

Perform the following steps to view the parameter modification history and check whether the modifications have taken effect:

Procedure

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.

Step 3 In the cluster list, find the target cluster and click the cluster name. The **Cluster Information** page is displayed.

Step 4 Click the **Modify Records** tab.

Status	Result	Operator	Time Range(Start)	Parameter N...	Start Time	End Time
To be restarted	success	E...	2022-09-17	Enter a parameter name	Dec 16, 2022 10:54:55 GMT+08:00	Dec 16, 2022 10:55:16 GMT+08:00

Parameter Name	Pre CN Value	Changed CN Value	Pre DN Value	Changed DN Value	Unit	Effective Or Not
alarm_report_interval	12311	123110	100	100	Second	Yes

Status	Result	Operator	Time Range(Start)	Parameter N...	Start Time	End Time
To be restarted	success	E...	2022-09-17	Enter a parameter name	Dec 16, 2022 09:56:40 GMT+08:00	Dec 16, 2022 09:56:58 GMT+08:00
Synchronized	success	E...	2022-09-17	Enter a parameter name	Dec 15, 2022 11:13:23 GMT+08:00	Dec 15, 2022 11:13:43 GMT+08:00
Applying failed	async job failed. ("result": "failed", "detail": "errorCod...	E...	2022-09-17	Enter a parameter name	Dec 15, 2022 10:32:21 GMT+08:00	Dec 15, 2022 10:32:29 GMT+08:00

NOTE

- If a parameter can take effect immediately after modification, its status will change to **Synchronized** after you modify it.
- If a parameter can take effect only after a cluster restart, its status will change to **To be restarted** after you modify it. You can click the expansion icon on the left to view the parameters that have not taken effect. After the cluster is restarted, the status of the record will change to **Synchronized**.

Step 5 By default, only the change history within a specified period is displayed. To check the entire change history of a parameter, search for it in the search box in the upper right corner.

----End

Parameter Description

There are a large number of database parameters. You can search for and view the parameters on the **Parameter Modification** page. For details, see [Modifying Parameters](#). The default values of the parameters are for reference only. For more information, see [Setting GUC Parameters](#).


9.2 Checking the Cluster Status

On the **Clusters > Dedicated Clusters** page of the GaussDB(DWS) console, you can view the general information about a cluster in the cluster list, such as the cluster status, task information, recent events, and node flavor.

Querying General Information of a Cluster

Log in to the GaussDB(DWS) console. In the navigation pane, click **Clusters > Dedicated Clusters**. The cluster list displays all clusters. If there are a large number of clusters, you can turn pages to view the clusters in any status.

In the upper part of the cluster list, click the search box and search for the required cluster based on the filter criteria (cluster name, cluster status, task information, node specifications, billing mode, recent events, and enterprise

project). Click  to refresh the cluster list and billing mode. You can also click **Search by Tag** to search for clusters based on cluster tags. For details, see [Searching for Clusters Based on Tags](#).

Clusters are listed in chronological order by default, with the most recent clusters displayed at the top. [Table 9-1](#) describes the cluster list parameters.

Table 9-1 Cluster list parameters

Parameter	Description
Cluster Name	Cluster name specified when a cluster is created. NOTE If the cluster name cannot be changed on the console, contact technical support.
Cluster Status	Cluster running status. For details, see Cluster Status .
Task Information	Cluster task status. For details, see Cluster Task Information .
Node Flavor	Node flavors of clusters.

Parameter	Description
Billing Mode	Cluster billing mode. <ul style="list-style-type: none">• In pay-per-use mode, the cluster creation time is displayed.
Recent Events	Number of recent events in a cluster. You can click the number to view event details.
Enterprise Project	Enterprise project to which a cluster belongs.
Operation	<ul style="list-style-type: none">• Monitoring Panel: For details, see Databases Monitoring (DMS).• More<ul style="list-style-type: none">– View Metric: For details, see Monitoring Clusters Using Cloud Eye.– Restart: Click Restart to restart a cluster. For details, see Starting and Stopping a Cluster.– Scale Out: For details, see Scaling Out a Cluster.– Change all specifications: For details, see Changing All Specifications.– Scale In: For details, see Scaling In a Cluster.– Redistribute: For details, see Redistributing Data.– View Scaling Details: For details, see Viewing Redistribution Details.– Expand Disk Capacity: For details, see Disk Capacity Expansion of an EVS Cluster.– Reset Password: For details, see Resetting a Password.– Create Snapshot: For details, see Manual Snapshots.– Delete: Click Delete to delete a cluster. For details, see Deleting a Cluster.– Change node flavor: For details, see Using the Elastic Specification Change.– Manage CN: For details, see Managing CNs.

Cluster Status

Table 9-2 Cluster status description

Status	Description
Available	Indicates that the cluster runs properly.

Status	Description
Read-only	<p>A cluster goes into this state when the disk usage of the cluster or a single node in the cluster is greater than 90%. The cluster can still work in this state but supports only query operations. Write operations are not supported. When the cluster status becomes read-only, remove the status by referring to Removing the Read-only Status. If the status cannot be removed, contact technical support engineers.</p> <p>After the read-only status is canceled for the cluster, you are advised to perform the following operations:</p> <ul style="list-style-type: none">• Use the SQL client tool to connect to the database as the administrator and run the following command to periodically clear and reclaim the storage space: <pre>VACUUM FULL;</pre>After you delete data stored in GaussDB(DWS) data warehouses, dirty data may be generated possibly because the disk space is not released. This results in disk space waste. It is recommended that the storage space be cleared periodically.• You are advised to check the disk capacity and analyze whether the existing cluster specifications meet service requirements. If not, expand the cluster capacity. For details, see Scaling Out a Cluster.
Unbalanced	<p>If the role of a GTM or DN in the cluster is different from the initial role, the cluster is in the Unbalanced state. In the Unbalanced state, the number of primary instances on some nodes increases. As a result, the load pressure is high. In this case, the cluster is normal, but the overall performance is not as good as that in a balanced state. You are advised to switch a cluster to the Available state during off-peak hours. For details, see Performing a Primary/Standby Switchback.</p>
Redistributing	<p>A cluster goes into this state when it detects that the service data on the original nodes is significantly larger than that on the new node after a new node is added to the cluster. In this case, the system automatically redistributes data on all nodes. The cluster can still work in this state.</p>
Redistribution failed	<p>A cluster goes into this state when data redistribution fails, but no data loss occurs. The cluster can still work in this state. You are advised to contact technical support.</p>
Degraded	<p>A cluster goes into this state when some nodes in the cluster are faulty, but the whole cluster runs properly. You are advised to contact technical support.</p>
Unavailable	<p>A cluster goes into this state when it cannot provide database services. You are advised to contact technical support.</p>
Creating	<p>A cluster goes into this state when it is being created.</p>
Creation failed	<p>A cluster goes into this state when it fails to be created.</p>

Status	Description
Creating, restoring	A cluster goes into this state when it is being restored from a snapshot.
Deleting	A cluster goes into this state when it is being deleted.
Frozen for legal reasons	Indicates that the cluster is frozen for legal reasons. In this case, the cluster cannot be deleted or unsubscribed, and its name cannot be changed.
Frozen	Indicates that the cluster is frozen (excluding legal reasons). In this case, the cluster name cannot be changed. If your account balance is insufficient and fee deduction fails, the retention period starts. During the retention period, the service resources will be frozen and cannot be used, but resources and data are reserved. To unfreeze the clusters, you need to top up your account to ensure that the account balance is not 0. For details, see How Do I Renew My Service?
To be restarted	This status indicates that GUC parameters have been modified in the cluster and the modification can take effect only after the cluster is restarted. Before the cluster is restarted, some O&M operations cannot be performed. After you manually restart the cluster, the GUC parameter takes effect and the cluster status changes to Available .
Stopped	Indicates that the cluster is stopped.

Cluster Task Information

Table 9-3 Task information description

Status	Description
Creating snapshot	Indicates that a snapshot is being created in the cluster.
Snapshot creation failed	Indicates that a snapshot fails to be created.
Observing	Indicates that the cluster is to be submitted after the automatic upgrade.
Configuring	Indicates that the system is storing modifications of cluster parameters.
Restarting	Indicates that a cluster is being restarted.
Restart failed	Indicates that a cluster fails to be restarted.

Status	Description
Converting to encryption cluster	Indicates that the cluster is being converted to an encrypted cluster.
Encryption cluster conversion failed	Indicates that the cluster fails to be encrypted.
Scaling out	Indicates that a cluster is being scaled out.
Scale-out failed	Indicates that a cluster fails to be scaled out.
Expanding disk capacity	Indicates that disk capacity is being expanded.
Disk expansion failed	Indicates that disk capacity fails to be expanded.
Associating ELB	Indicates that ELB is being associated.
Failed to associate ELB	Indicates that ELB fails to be associated.
Disassociating ELB	Indicates that ELB is being disassociated.
Failed to disassociate ELB	Indicates that ELB fails to be disassociated.
Checking scale-in	The service is checking whether a cluster can be scaled in.
Scale-in check failed	A cluster does not meet the scale-in requirements. For example: <ul style="list-style-type: none">• The value of default_storage_nodegroup is not installation.• In the cluster database, data_redis is a reserved redistribution schema, but the schema contains user tables.• The cluster disk space does not meet the scale-in requirements. For details, see Scaling In a Cluster.
Scaling in	A cluster is being scaled in.
Scale-in failed	The cluster scale-in fails. You need to manually scale in the cluster again as soon as possible, or your services will be affected.
Switching back	The primary/standby relationship of a cluster is being restored.

Status	Description
Switchback failed	The primary/standby relationship of a cluster fails to be restored. Possible causes are as follows. <ul style="list-style-type: none">• Redo operations are being performed on DNs. Wait until the operations are completed and try again.• Failed to query DN redo information. Check tenant logs to identify the failure cause.• Primary/standby catchup is in progress. Wait until it is completed and try again.• Failed to query primary/standby catchup information. Check tenant logs to identify the failure cause.• Primary/standby catchup failed. Contact technical support or try again later. Check tenant logs to identify the failure cause.• The cluster is abnormal.
Changing node flavor	The cluster is being scaled.
Node flavor change failed	All specifications change failed
Waiting for payment	The order for changing a pay-per-use cluster to a yearly/monthly cluster has not been paid. After the order is paid or canceled, the status will change.
Changing all specifications	All the specifications of the cluster being changed.
All specifications change failed	Specifications change failed because of insufficient quotas or permissions, or abnormal cluster status.
Maintaining	A maintenance change operation, such as cluster upgrade or plug-in upgrade, is being performed on the cluster.
Maintain_failure	A cluster fails to be restarted.
Stopping	Indicates that the cluster is being stopped.
Starting	Indicates that the cluster is being started.
Inspecting	Indicates that the cluster is being inspected before the change.
Inspection failed	Indicates that the cluster inspection fails.

9.3 Viewing Cluster Details

Log in to the GaussDB(DWS) console. In the navigation tree on the left, click **Clusters > Dedicated Clusters**. In the cluster list, locate the required cluster and click its name. The **Cluster Information** page is displayed.

On the **Cluster Information** page, you can view the following information:

- **Basic Information:** [Table 9-4](#) lists the related parameters.
- **Connection:** [Table 9-5](#) describes the parameters.
- **Network:** [Table 9-6](#) lists the related parameters.
- **Billing Information:** [Table 9-7](#) describes the parameters.
- **Storage/Backup Capacity:** [Table 9-8](#) describes the parameters.
- **O&M Account:** [Table 9-9](#) describes the related parameters.
- **Data Encryption Information:** [Table 9-10](#) lists the related parameters.

NOTE

You can view this module if you enable the data encryption function when creating a cluster.

Table 9-4 Basic information

Parameter	Description
Cluster Name	Cluster name specified when a cluster is created.
Cluster Status	Cluster running status. For details, see Cluster Status .
Parameter Configuration Status	Parameter configuration status of a cluster.
Task Information	Cluster task status. For details, see Cluster Task Information .
Current Specifications	Current node specifications.
Nodes	Number of nodes in the cluster.
Logical Clusters	You can enable it as required. The Logical Clusters menu item will be displayed after you enable it.
Cluster ID	ID of the cluster.
Cluster Version	Cluster version information.

Parameter	Description
Node Flavor	Node flavor of the cluster.
Enterprise Project	Enterprise project to which a cluster belongs. You can click the enterprise project name to view and edit it on the console of the Enterprise Project service.
Time Zone	The cluster time zone affects the node OS, log files, and data warehouse. You can change the time zone for the node OS and log files, but not for the data warehouse databases. To change the time zone of the data warehouse databases, use the GUC parameter timezone . For details, see Modifying Database Parameters .

Table 9-5 Connection

Parameter	Description
Private Network Domain Name	<p>Domain name for accessing the cluster database through the internal network. The domain name corresponds to all CN IP addresses. The private network domain address is automatically generated when a cluster is created.</p> <p>NOTE</p> <ul style="list-style-type: none">• If the cluster name does not comply with the domain name standards, the prefix of the default access domain name will be adjusted accordingly.• Load balancing is not supported. <p>You can click Modify to change the private network domain name. The access domain name contains 4 to 63 characters, which consists of letters, digits, and hyphens (-), and must start with a letter.</p> <p>For details, see Managing Access Domain Names.</p>
Private Network IP Address	<p>IP address for accessing the database in the cluster over the private network.</p> <p>NOTE</p> <ul style="list-style-type: none">• A private IP address is automatically generated when you create a cluster. The IP address is fixed.• The number of private IP addresses equals the number of CNs. You can log in to any node to connect to the cluster.• If you access a fixed IP address over the internal network, all the resource pools will run on a single CN.• If IPv6 is enabled for a cluster, both IPv4 and IPv6 private addresses will be displayed. You can use either of them as needed.

Parameter	Description
Public Network Domain Name	Name of the domain for accessing the database in the cluster over the public network. For details, see Managing Access Domain Names . NOTE Load balancing is not supported.
Public Network IP Address	IP address for accessing the database in the cluster over the public network. NOTE <ul style="list-style-type: none">If no EIP is assigned during cluster creation and Public Network IP Address is empty, click Edit to bind an EIP to the cluster.If an EIP is bound during cluster creation, click Edit to unbind the EIP.
Initial Administrator	Database administrator specified during cluster creation. When you connect to the cluster for the first time, you need to use the initial database administrator and password to connect to the default database.
Port	Port number for accessing the cluster database through the public network or private network. The port number is specified when the cluster is created.
Default Database	Database name specified when the cluster is created. When you connect to the cluster for the first time, connect to the default database.
ELB Address	To achieve high availability and avoid single-CN failures, a new cluster needs to be bound to ELB. You are advised to use the ELB address to connect to the cluster.

Table 9-6 Network

Parameter	Description
Region	Current working zone of the cluster.
AZ	AZ selected during cluster creation
VPC	VPC selected during cluster creation. A VPC is a secure, isolated, and logical network environment. After a data warehouse cluster is created, its VPC cannot be changed. However, you can edit and modify the current VPC. You can click the VPC name to go to the VPC details page to configure it.

Parameter	Description
Subnet	<p>Subnet selected during cluster creation.</p> <p>A subnet provides dedicated network resources that are isolated from other networks, improving network security.</p> <p>After a data warehouse cluster is created, its subnet cannot be changed. However, you can edit and modify the current subnet. You can click the subnet name to go to the subnet details page to configure it.</p>
Security Group	<p>Security group selected during cluster creation.</p> <p>After a GaussDB(DWS) cluster is created, you can change the security group. You can also add, delete, or modify security group rules in the current security group. Changing the security group of a cluster may cause brief service disruption. Exercise caution when performing this operation. For better network performance, do not select more than five security groups.</p> <ul style="list-style-type: none">• To change the security group, click Modify on the right of the security group name, select the security group name to be changed, and click OK.• Modifying an existing security group rule: Click the security group name to go to the security group details page.

Table 9-7 Billing information

Parameter	Description
Billing Mode	<p>Billing mode.</p> <ul style="list-style-type: none">• Pay-per-use
Created	Time when a pay-per-use cluster is created.

Table 9-8 Storage/Backup capacity

Parameter	Description
Storage	<p>The storage class Ultra-high I/O and the storage space usage are displayed.</p> <p>NOTE</p> <ul style="list-style-type: none">• The used storage capacity does not include data on OBS foreign tables. It includes only GaussDB(DWS) data, including files, logs, snapshots, and indexes.• The available storage space is half of the actual disk capacity.
Backup	The space in use, free space, and charged space of the cluster are displayed.

Parameter	Description
OBS Hot Data Used Capacity (DWS 3.0 parameter)	OBS hot data capacity used by DWS 3.0.
Cold Data Used Capacity	OBS capacity used by cold data. NOTE OBS capacity usage. It is synchronized every hour.
Used Capacity of OBS Foreign Tables	OBS capacity used by the foreign tables of the default OBS server of the cluster: default_obs_foreign_table_server . NOTE OBS capacity usage. It is synchronized every hour.

Table 9-9 O&M account

Parameter	Description
O&M Account	Specifies whether to enable the cluster O&M account. Check the created O&M account. Its name format is om_user_first_eight_numbers_in_cluster_ID . The gs_role_analyze_any , gs_role_vacuum_any , gs_role_read_all_stats , and gs_role_signal_backend roles will be assigned to the account. For details, see Preset Roles .
Account Status	Displays the status of the current cluster O&M account, which can be Normal or Expired .
Expires	Indicates the expiration time of the O&M account of the current cluster.
Extend by 8h	<ul style="list-style-type: none">For a normal account, its validity period is extended to 8 hours later than its expiration time.For an expired account, its validity period is extended to 8 hours later than the current time.

Table 9-10 Data encryption information

Parameter	Description
Key Name	Indicates the database encryption key of the cluster when Encrypt DataStore is enabled.
Last Key Rotation Time	Indicates the time when the last encryption key is rotated when Encrypt DataStore is enabled.

9.4 Changing a Cluster Name

Scenario

After a GaussDB(DWS) cluster is created, you can change the cluster name as required.

After the cluster name is changed, the names of all nodes in the current cluster are changed accordingly.

NOTE

- If the cluster name cannot be changed on the console, contact technical support to upgrade the console.
- If the cluster name fails to be changed, the cluster functions are not affect. You can contact technical support rectify the fault.

Constraints and Limitations

If the cluster is in the **Unavailable** status or is performing other tasks, the cluster name cannot be changed. You can change the cluster name only after the cluster status changes to **Available** or the running tasks are complete.


Procedure

Method 1:

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the cluster list, click the modification icon next to a cluster name to modify the cluster.

Figure 9-1 Changing the name of a cluster in the cluster list



Step 3 In the displayed dialog box, enter a new cluster name.

Step 4 Confirm the information and click **OK**.

----End


Method 2:

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the cluster list, click the name of a cluster.

Step 3 On the displayed **Cluster Details** page, click the modification icon next to the cluster name in the **Basic Information** area.

Figure 9-2 Changing the cluster name on the cluster details page



Step 4 After confirming that the information is correct, click **OK** to deliver the cluster modification task. After the task is complete, the cluster name is changed successfully.

----End

9.5 O&M Account

Context

If you need technical support when using a cluster, you can authorize them to use an O&M account on the GaussDB(DWS) console to access the cluster for fault locating.

NOTE


Only cluster 8.1.3.110 and later versions support O&M accounts. For earlier versions, contact technical support.

Overview

You can perform the following operations:

- Enable or disable the O&M account.
- Check the O&M account status.
- View the validity period of an O&M account and extend it as needed. (If the validity period is not extended, the account will expire automatically.)

Figure 9-3 O&M Account



Enabling the O&M Account

Step 1 Log in to the GaussDB(DWS) management console.

- Step 2** In the cluster list, click the name of a cluster.
- Step 3** On the cluster details page, and enable **O&M Account** in the **O&M Account** area.
- Step 4** In the displayed dialog box, click **OK**.
- Step 5** Check the created O&M account. Its name format is **om_user_First_eight_numbers_in_cluster_ID**.

Assign the **gs_role_analyze_any**, **gs_role_vacuum_any**, **gs_role_read_all_stats**, and **gs_role_signal_backend** roles to the account. For details, see [Preset Roles](#).

 **NOTE**

You can toggle off the switch and delete the O&M account if it is no longer needed.

----End

Extending the Validity Period

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** In the cluster list, click the name of a cluster.
- Step 3** On the cluster details page, click **Extend by 8h** in the **O&M Account** area.
- Step 4** In the displayed dialog box, click **OK**.
- For a normal account, its validity period is extended to 8 hours later than its expiration time.
 - For an expired account, its validity period is extended to 8 hours later than the current time.

----End

9.6 Managing Access Domain Names

Overview

A domain name is a string of characters separated by dots to identify the location of a computer or a computer group on the Internet, for example, www.example.com. You can enter a domain name in the address box of the web browser to access a website or web application.

On GaussDB(DWS), you can access clusters using the private network domain name or the public network domain name.

Private network domain name: Name of the domain for accessing the database in the cluster through the private network. The private network domain name is automatically generated when you create a cluster.

Public network domain name: Name of the domain for accessing the database in the cluster through the public network. If a cluster is not bound to an EIP, it cannot be accessed using the public network domain name. If you bind an EIP during cluster creation, the public network domain name is automatically generated.

 NOTE

Neither public nor private domain names support load balancing. To use load balancing, see [Configuring JDBC to Connect to a Cluster \(Load Balancing Mode\)](#).

After a cluster is created, you can set private and public domain names for accessing the cluster as required. The operations are as follows:

- [Modifying a Private Network Domain Name](#)
- [Creating a Public Network Domain Name](#)
- [Modifying a Public Network Domain Name](#)
- [Releasing a Public Network Domain Name](#)

Modifying a Private Network Domain Name

The private network domain name is automatically generated during cluster creation. After the cluster is created, you can modify the default domain name based on site requirements.

To modify the private network domain name, perform the following steps:

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 3** In the cluster list, find the target cluster and click the cluster name. The **Cluster Information** page is displayed.
- Step 4** In the **Connection** area, click **Modify** next to the automatically generated **Private Network Domain Name**.
- Step 5** In the **Modify Private Network Domain Name** dialog box, enter the target domain name and click **OK**.

The private network domain name contains 4 to 63 characters, which consists of letters, digits, and hyphens (-) and must start with a letter.

After the domain name is modified, click copy button next to the private network domain name to copy it.

----End

Creating a Public Network Domain Name

A cluster is not bound to an EIP by default during cluster creation. That is, cluster access using the public network is disabled. After a cluster is created, if you want to access it over the public network, bind an EIP to the cluster and create a public network domain name.

 NOTE

By default, only Huawei Cloud accounts or users with Security Administrator permissions can query and create agencies. By default, the IAM users in those accounts cannot query or create agencies. When the users use the EIP, the system makes the binding function unavailable. Contact a user with the **DWS Administrator** permissions to authorize the agency on the current page.

To create a public network domain name, perform the following steps:


- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 3** In the cluster list, find the target cluster and click the cluster name. The **Cluster Information** page is displayed.
- Step 4** In the **Connection** area, **Public Network Domain Name** and **Public Network IP Address** are empty. Click **Edit** to bind the cluster with an EIP.
- Step 5** In the **Edit Elastic IP** dialog box, select an EIP from the drop-down list to bind it to a specified CN.

If no available EIPs are displayed, click **View EIP** to go to the **Elastic IP** page and create an EIP that satisfies your needs. After the new EIP is created, click the refresh icon next to the drop-down list. The newly created EIP will be displayed in the **EIP** drop-down list.

After the EIP is bound successfully, the specific public network IP address is displayed in the **Connection** area.

- Step 6** In the **Connection** area, click **Create** next to **Public Network Domain Name** to create a public network domain name for the cluster.
- Step 7** In the **Apply for Public Network Domain Name** dialog box, enter the target domain name and click **OK**.

The public network domain name contains 4 to 63 characters, which consists of letters, digits, and hyphens (-) and must start with a letter.

The specific public network domain name is displayed in the **Connection** area after being created. Click copy button  to copy the public network domain name.

----End

Modifying a Public Network Domain Name

If you bind an EIP during cluster creation, the public network domain name is automatically generated. After a cluster is created, you can modify the public network domain name as required.

To modify the public network domain name, perform the following steps:

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 3** In the cluster list, find the target cluster and click the cluster name. The **Cluster Information** page is displayed.
- Step 4** Click **Modify** next to the **Public Network Domain Name** in the **Connection** area.
- Step 5** In the **Modify Public Network Domain Name** dialog box, enter the target domain name and click **OK**.

----End

Releasing a Public Network Domain Name

After a cluster is created, you can release unnecessary public network domain names.

To do so, perform the following steps:

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 3** In the cluster list, find the target cluster and click the cluster name. The **Cluster Information** page is displayed.
- Step 4** Click **Release** next to the **Public Network Domain Name** in the **Connection** area.
- Step 5** In the **Release Domain Name** dialog box, click **Yes**.

----End

9.7 Cluster Topology

Overview

A topology shows all the nodes in a cluster. You can check the node statuses, processes, and IP addresses.

NOTE

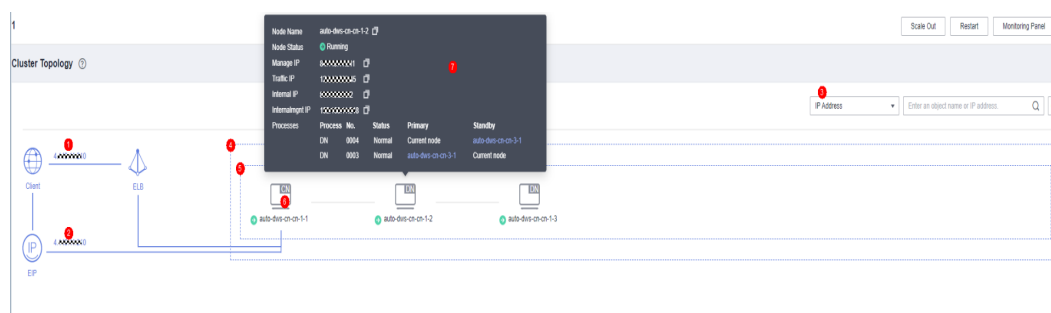
- You can check the topology structure and node processes.
- Only cluster versions 8.0.0 and later can display the topology structure. Only cluster versions 8.2.0 and later can display node processes.

Viewing the Cluster Topology

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the cluster list, click the name of a cluster.
- Step 3** On the **Cluster Details** page, click the **Cluster Topology** tab.
- Step 4** In the upper part of the page, you can select **IP Address** or **Node Name**. After entering the IP address or node name in the search box, you can view the location of the IP address or node name in the cluster topology.

----End

Topology Overview



This figure shows a topology. The elements marked in the figure are as follows:

1. Public IP address of the ELB bound to the cluster. If no public IP addresses are bound to the ELB, the service address is displayed.
2. EIP bound to the cluster.
3. Search category. You can perform exact search by IP address or node name.
4. Rings in the cluster.
5. A ring. Each ring occupies a line. An icon in a ring indicates a node.
6. A node. The type of the node is displayed in the upper right corner of the icon. Currently, the type can only be CN or DN. If there is a CN process on the node, **CN** is displayed. If there are no CN processes on the node, **DN** is displayed.
7. Node details, including the node name, status, IP addresses, and task process. Node details are displayed when you hover your cursor over a node icon.

Terms in the Topology View

Table 9-11 Cluster structure description

Name	Description	Usage
ELB	Elastic Load Balance (ELB) automatically distributes incoming traffic across multiple backend servers based on listening rules you configure.	If the private IP address or EIP of a CN is used to connect to a GaussDB(DWS) cluster, the failure of this CN will lead to a cluster connection failure. If a private or public domain name is used for connection, the DNS service randomly selects a private IP address or EIP for each client. This cannot balance loads or avoid single-CN failures. ELB is used to solve these problems. For details, see Associating and Disassociating ELB .
EIP	The Elastic IP (EIP) service provides static public IP addresses and scalable bandwidths that enable your cloud resources to communicate with the Internet.	EIPs can be bound to or unbound from ECSs, BMSs, virtual IP addresses, load balancers, and NAT gateways.

Name	Description	Usage
Ring	A security ring is used for isolating faulty servers. A fault in a ring does not affect servers outside the ring.	<p>Data on a DN has multiple copies in a ring, and will not be lost even if the DN server is faulty.</p> <p>For example, if Server1 in a ring is faulty, the standby DN1 on Server2, the standby DN2 on Server3, and the standby DN3 on Server3 are still running. The loads of servers in a ring are still balanced.</p> <p>A cluster can run properly as long as the number of faulty servers does not exceed the number of rings.</p> <p>NOTE The ring is the minimum unit for a scale-out. When you scale out a cluster, the added nodes must be a multiple of the ring quantity.</p>

Table 9-12 Node IP addresses

Name	Description	Usage
Manage IP	IP address used by a data warehouse node to communicate with the management plane	It is used by the management plane to deliver commands, and used by the node to report node status and monitoring information.
Traffic IP	IP address of a data warehouse node for external access.	This IP address can be bound to an EIP or ELB, or directly connect to a VPC.
Internal IP	IP address used for communication inside a data warehouse cluster.	-
Internalmgnt IP	IP address used by nodes to send internal management commands in a data warehouse cluster.	-

Table 9-13 Node processes

Name	Description	Usage
CMS	<p>A Cluster Manager (CM) manages and monitors the running status of functional units and physical resources in the distributed system, ensuring system stability. CM Server (CMS) is a module of CM.</p>	<p>A CM consists of CM Agent, OM Monitor, and CM Server.</p> <ul style="list-style-type: none"> • CM Agent monitors the running status of primary and standby GTMs, CNs, and primary and standby DNs on the host, and reports the status to CM Server. In addition, it executes the arbitration instruction delivered by CM Server. A CM Agent process runs on each server. • OM Monitor monitors scheduled tasks of CM Agent and restarts CM Agent when CM Agent stops. If CM Agent cannot be restarted, the server will be unavailable. In this case, you need to manually rectify this fault. <p>NOTE A CM Agent restart fails probably because of lack of system resources, which rarely happens.</p> <ul style="list-style-type: none"> • CM Server checks whether the current system is normal according to the instance status reported by CM Agent. In the case of exceptions, CM Server delivers recovery commands to CM Agent. <p>GaussDB(DWS) deploys CM Server in primary/standby mode to ensure system HA. CM Agent</p>

Name	Description	Usage
		connects to the primary CM Server. If the primary CM Server is faulty, the standby CM Server is promoted to primary to prevent single-CM faults.
GTM	A Global Transaction Manager (GTM) generates and maintains the globally unique information, such as the transaction ID, transaction snapshot, and timestamp.	A cluster includes only one pair of GTMs: one primary and one standby GTM.
CN	A Coordinator (CN) receives access requests from applications, and returns execution results to the client; splits tasks and allocates task fragments to different DNs for parallel processing.	<p>CNs in a cluster have equivalent roles and return the same result for the same DML statement. Load balancers can be added between CNs and applications to ensure that CNs are transparent to applications. If a CN is faulty, the load balancer connects its applications to another CN.</p> <p>CNs need to connect to each other in the distributed transaction architecture. To reduce heavy load caused by excessive threads on GTMs, no more than 10 CNs should be configured in a cluster.</p>

Name	Description	Usage
CCN	Central Coordinator (CCN)	GaussDB(DWS) handles the global resource load in a cluster using the Central Coordinator (CCN) for adaptive dynamic load management. When the cluster is started for the first time, the CM selects the CN with the smallest ID as the CCN. If the CCN is faulty, CM replaces it with a new one.
DN	A Data Node (DN) stores data in row-store, column-store, or hybrid mode, executes data query tasks, and returns execution results to CNs.	There are multiple DNs in the cluster. Each DN stores part of data. If DNs are not deployed in primary/standby mode and a DN is faulty, data on the DN will be inaccessible.

9.8 Managing Tags

9.8.1 Overview

A tag is a key-value pair customized by users and used to identify cloud resources. It helps users to classify and search for cloud resources.

Tags are composed of key-value pairs.

- A key in a tag can have multiple values.
- A cloud resource must have a unique key.

On GaussDB(DWS), after creating a cluster, you can add identifiers to items such as the project name, service type, and background information using tags. If you use tags in other cloud services, you are advised to create the same tag key-value pairs for cloud resources used by the same business to keep consistency.

GaussDB(DWS) supports the following two types of tags:

- Resource tags
Non-global tags created on GaussDB(DWS)
- Predefined tags
Predefined tags created on Tag Management Service (TMS). Predefined tags are global tags.

For details about predefined tags, see the *Tag Management Service User Guide*.

On GaussDB(DWS), tags can be added to the following resources:

- Cluster

Tags can be added to a cluster when the cluster is being created or after it is successfully created. You can search for the cluster in the cluster list using tags.

Each cluster can have a maximum of 20 tags.

After you add tags to a cluster and then create a snapshot for the cluster, the tags cannot be restored if you use the snapshot to restore the cluster. Instead, you need to add tags again.

When a cluster is deleted, non-predefined tags associated with the cluster are also deleted. Predefined tags need to be deleted on TMS.

9.8.2 Managing Tags

This section describes how to search for clusters based on tags and how to add, modify, and delete tags.

Adding a Tag to a Cluster

- Step 1** On the **Clusters > Dedicated Clusters** page, click the name of the cluster to which a tag is to be added, and choose **Tag**.
- Step 2** Click **Add Tag**.
- Step 3** Configure tag information in the **Add Tag** dialog box. The value of a key cannot be left blank.

Table 9-14 Tag parameters

Parameter	Description	Example Value
Tag key	<p>You can:</p> <ul style="list-style-type: none">• Select a predefined tag key or an existing resource tag key from the drop-down list of the text box. <p>NOTE To add a predefined tag, you need to create one on TMS and select it from the drop-down list of Tag key. You can click View predefined tags to enter the Predefined Tags page of TMS. Then, click Create Tag to create a predefined tag.</p> <ul style="list-style-type: none">• Enter a tag key in the text box. The tag key can contain a maximum of 128 characters and cannot be an empty string. It cannot start with _sys_. Only letters, digits, spaces, and the following characters are allowed: _ . : = + - @ <p>NOTE A key must be unique in a given cluster.</p>	key01

Parameter	Description	Example Value
Tag value	You can: <ul style="list-style-type: none">• Select a predefined tag value or resource tag value from the drop-down list of the text box.• Enter a tag value in the text box. The tag key can contain a maximum of 255 characters and cannot be an empty string. Only letters, digits, spaces, and the following characters are allowed: _ . : = + - @	value01

Step 4 Click **OK**.

----End

Searching for Clusters Based on Tags

You can quickly locate a tagged cluster using tags.

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters > Dedicated Clusters**.

Step 3 Click the search box above the cluster list and select the **Resource Tag** filter.

Step 4 Click the tag key to be searched for and select the corresponding tag value. Click the search box again to add more tag filters.

Search by tag supports only the keys and values that exist in the drop-down list. If no tag key or value is available, create a tag for the cluster. For details, see [Adding a Tag to a Cluster](#).

Step 5 Click **Search**. The target cluster will be displayed in the cluster list.

----End

Modifying a Tag

Step 1 On the **Clusters > Dedicated Clusters** page, click the name of the cluster to which a tag is to be added, and choose **Tag**.

Step 2 Locate the row that contains the tag to be modified, and click **Edit** in the **Operation** column. The **Edit Tag** dialog box is displayed.

Step 3 Enter the new key value in the **Value** text box.

Step 4 Click **OK**.

----End

Deleting a Tag

Step 1 On the **Clusters > Dedicated Clusters** page, click the name of the cluster from which a tag is to be deleted, and click the **Tags** tab.

Step 2 Locate the row that contains the tag to be deleted, click **Delete** in the **Operation** column. The **Delete Tag** dialog box is displayed.

Step 3 Click **Yes** to delete the tag.

----End

9.9 Managing Enterprise Projects

An enterprise project is a cloud resource management mode. Enterprise Management provides users with comprehensive management in cloud-based finance. The Enterprise Management console differs from typical management consoles as it focuses on resource management rather than independent control and configuration of cloud products. It assists enterprises in managing finances within the hierarchy of companies, departments, and projects.

Users who have enabled the Enterprise Project Management service can use it to manage cloud service resources.

Binding an Enterprise Project

You can select an enterprise project during cluster creation to associate it with the cluster. For details, see [Creating a DWS 2.0 Cluster](#). The **Enterprise Project** drop-down list displays the projects you created. In addition, the system has a built-in enterprise project (**default**). If you do not select an enterprise project for the cluster, the default project is used.

Note that the Enterprise Project Management service is still in the OBT. Only users with the OBT permission can set enterprise projects. Common users cannot view the enterprise project information.

During cluster creation, if the cluster is successfully bound to an enterprise project, the cluster will be successfully created. If the binding fails, the system sends an alarm and the cluster fails to be created.

Snapshots of a cluster retain the association between the cluster and its enterprise project. When the cluster is restored, the association is also restored.

When you delete a cluster, the association between the cluster and its enterprise project is automatically deleted.

Viewing Enterprise Projects

After a cluster is created, you can view the associated enterprise project in the cluster list and **Cluster Information** page. You can query only the cluster resources of the project on which you have the access permission.

- In the cluster list on the **Clusters** page, view the enterprise project to which the cluster belongs.
- In the cluster list, find the target cluster and click the cluster name. The **Cluster Information** page is displayed, on which you can view the enterprise project associated with the cluster. Click the enterprise project name to view and edit it on the Enterprise Management console.
- When querying the resource list of a specified project on the Enterprise Management console, you can also query the GaussDB(DWS) resources.

Searching for Clusters by Enterprise Project

Log in to the GaussDB(DWS) console and click **Clusters > Dedicated Clusters**. Click the search box above the cluster list and select **Enterprise Project**. Enter the project name and click the search button to view all clusters associated with the project.

Migrating a Cluster to or Out of an Enterprise Project

A GaussDB(DWS) cluster can be associated with only one enterprise project. After a cluster is created, you can migrate it from its current enterprise project to another one on the Enterprise Management console, or migrate the cluster from another enterprise project to a specified enterprise project. After the migration, the cluster is associated with the new enterprise project. The association between the cluster and the original enterprise project is automatically released. For details, see "Resource Management" > "Managing Enterprise Project Resources" in the *Enterprise Management User Guide*.

Enterprise Project-Level Authorization

If permissions preset in the system cannot meet requirements, you can customize policies and grant the policies to user groups for refined access control. As an independent managed object, the enterprise project can be bound to a user group, and the customized policy can be granted to the user group. This implements refined authorization at the enterprise project level.

Step 1 Log in to the IAM console and create a custom policy.

Refer to the following to create the policy:

- Use the IAM administrator account, that is, the user in the admin user group, because only the IAM administrator has the permissions to create users and user groups and modify user group permissions.
- GaussDB(DWS) is a project-level service, so its **Scope** must be set to **Project-level services**. If this policy is required to take effect for multiple projects, authorization is required to each project.
- Some GaussDB(DWS) policy templates are preconfigured on IAM. When creating a custom policy, you can select one of the following templates and modify the policy authorization statement based on the template:
 - **DWS FullAccess**: all execution permissions for GaussDB(DWS)
 - **DWS ReadOnlyAccess**: read-only permission for GaussDB(DWS)
 - **DWS Administrator**: all execution permissions for GaussDB(DWS)
 - **DWS Database Access**: Users granted this permission can generate temporary database user credentials based on IAM users to connect to databases in the data warehouse clusters.
- You can add permissions corresponding to GaussDB(DWS) operations or RESTful APIs listed in [List of Supported Actions](#) to the action list in the policy authorization statement, so that the policy can obtain the permissions.
For example, if **dws:cluster:create** is added to the action list of a policy statement, the policy has the permission to create or restore clusters.
- If you want to use other services, grant related operation permissions on these services. For details, see the help documents of related services.

For example, when creating a GaussDB(DWS) cluster, configure the VPC to which the cluster belongs. To obtain the VPC list, add action **vpc:*:get*** to the policy statement.

Policy example:

- Example in which multiple operation permissions are supported

The following policy has the permissions to create/restore/restart/delete a cluster, set security parameters, and reset passwords.

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "dws:cluster:create",
        "dws:cluster:restart",
        "dws:cluster:delete",
        "dws:cluster:setParameter",
        "dws:cluster:resetPassword",
        "ecs:*:get*",
        "ecs:*:list*",
        "vpc:*:get*",
        "vpc:*:list*"
      ]
    }
  ]
}
```

- Example of wildcard (*) usage

The following policy has all operation permissions on GaussDB(DWS) snapshots.

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "dws:snapshot:*",
        "ecs:*:get*",
        "ecs:*:list*",
        "vpc:*:get*",
        "vpc:*:list*"
      ]
    }
  ]
}
```

Step 2 Click **Enterprise** in the upper right corner of the management console to enter the Enterprise Management console.

Step 3 Choose **Personnel Management > User Group Management** in the left navigation tree. Then, create a user group and add users to it, add the user group to a project, and grant the newly created custom policy to the group so that users in the group can obtain the permissions defined by the policy.

----End

9.10 Managing Clusters That Fail to Be Created

If a cluster fails to be created, you can go to the **Clusters > Dedicated Clusters** page of the GaussDB(DWS) console to view the cluster status and the cause of failure.

Checking the Cause of a Creation Failure

Step 1 Log in to the GaussDB(DWS) console. In the navigation pane, choose **Clusters > Dedicated Clusters**.

Step 2 In the cluster list, locate the cluster whose **Cluster Status** is **Creation failed**.

Step 3 Click  in the **Cluster Status** column to view the cause of the creation failure.

You can rectify the fault based on the error code displayed. For details, see [Error Code Reference](#). If the fault persists, contact technical support.

----End

Deleting a Cluster That Fails to Be Created

You can delete a cluster that fails to be created if you do not need it. Before deletion, check the cause of creation failure.

Step 1 Log in to the GaussDB(DWS) console. In the navigation pane, choose **Clusters > Dedicated Clusters**.

Step 2 In the cluster list, locate the row containing the failed cluster to be deleted, and choose **More > Delete**.

Step 3 In the displayed dialog box, confirm the deletion. You can determine whether to perform the following operations:

- Create a snapshot for the cluster.

If the cluster status is normal, you can click **Create Snapshot**. In the dialog box that is displayed, enter the snapshot name and click **OK** to create a snapshot for the cluster to be deleted. In the row of a cluster, choose **More > Delete**.

- Delete associated resources:

- Release the EIP bound to the cluster.

If an EIP is bound to the cluster, you are advised to select **EIP** to release the EIP. If you do not release the EIP, you can bind it to another cluster or cloud resource and it will be billed based on the EIP pricing rule of VPC.

- Delete automated snapshots.
- Delete manual snapshots.

If you have created a manual snapshot, you can select **Manual Snapshot** to delete it.

Step 4 After confirming that the information is correct, enter **DELETE** or click **Auto Enter** and click **OK** to delete the cluster. The cluster status in the cluster list will change to **Deleting**, and the cluster deletion progress will be displayed.

If the cluster to be deleted uses an automatically created security group that is not used by other clusters, the security group is automatically deleted when the cluster is deleted.

----End

9.11 Removing the Read-only Status

A cluster in read-only status does not allow write operations. You can remove this status on the management console. A cluster becomes read-only probably because of high disk usage. For how to solve this problem, see [High Disk Usage and Read Only Status](#).

NOTE

- The read-only status can be canceled for version 1.7.2 or later.
- In 8.2.0 and later versions, you can free up disk space by using **DROP/TRUNCATE TABLE** in a read-only cluster.

Impact on the System

- You can cancel the read-only status only when a cluster is read-only.
- When a cluster is in read-only status, stop the write tasks to prevent data loss caused by used up disk space.
- After the read-only status is canceled, clear the data as soon as possible to prevent the cluster from entering the read-only status again after a period of time.

Removing Read-only Status

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters > Dedicated Clusters**.

All clusters are displayed by default.

Step 3 In the row containing the cluster whose cluster status is **Read-only**, click **Cancel Read-only**.

Step 4 In the dialog box that is displayed, click **OK** to confirm and remove the read-only status for the cluster.

----End

9.12 Performing a Primary/Standby Switchback

Context

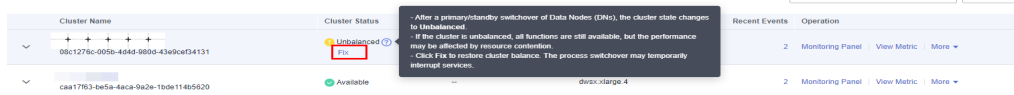
In the **Unbalanced** state, the number of primary instances on some nodes increases. As a result, the load pressure is high. In this case, the cluster is normal, but the overall performance is not as good as that in a balanced state. Restore the primary-standby relationship to recover the cluster to the available state.

NOTE

- Only 8.1.1.202 and later versions support primary/standby cluster restoration.
- Cluster restoration interrupts services for a short period of time. The interruption duration depends on the service volume. You are advised to perform this operation during off-peak hours.

Procedure

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster in unbalanced state.
- Step 3** In the **Cluster Status** column of the cluster, click **Fix** under **Unbalanced**.



- Step 4** In the dialog box that is displayed, confirm that the service is in off-peak hours, and click **Yes**. A message will be displayed in the upper right corner, indicating that the switchback request is being processed.
- Step 5** Check the cluster status. During the switchback, the cluster status is **Switching back**. After the switchback, the cluster status will change to **Available**.

----End

9.13 Starting and Stopping a Cluster

Restarting a cluster

If a cluster is in the **Unbalanced** state or cannot work properly, you may need to restart it for restoration. After modifying a cluster's configurations, such as security settings and parameters, manually restart the cluster to make the configurations take effect.

NOTE

If your cluster is in arrears, this function may be unavailable. Please top up your account in time.

Impact on the System

- A cluster cannot provide services during the restart. Therefore, before the restart, ensure that no task is running and all data is saved.

If the cluster is processing service data, such as importing data, querying data, creating snapshots, or restoring snapshots, cluster restarting will cause file damage or restart failure. You are advised to stop all cluster tasks before restarting the cluster.

View the **Session Count** and **Active SQL Count** metrics to check whether the cluster has active events. For details, see [Monitoring Clusters Using Cloud Eye](#).

- The time required for restarting a cluster depends on the cluster scale and services. Generally, it takes about 3 minutes to restart a cluster. The duration does not exceed 20 minutes.
- If the restart fails, the cluster may be unavailable. Try again later or contact technical support.

Procedure

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters > Dedicated Clusters**.

Step 3 In the **Operation** column of the cluster to be restarted, click **Restart**.

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

Task Information changes to **Restarting**. When **Cluster Status** changes to **Available** again, the cluster is successfully restarted.

----End

Stopping a Cluster

If a cluster is no longer used, you can stop the cluster to bring services offline.

NOTE

- If the current console does not support this feature, contact technical support. Billing resumes after the cluster is started.
- After the cluster is stopped, ECS basic resources (vCPUs and memory) are no longer reserved. When you start the service again, it may fail to be started due to insufficient resources. In this case, wait for a while and try again later.

Procedure

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Clusters > Dedicated Clusters**.

Step 3 On the **Dedicated Clusters** page, locate the row that contains the target dedicated cluster, click **More > Stop** in the **Operation** column.

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

The **Task Information** of the cluster changes to **Stopping**. If the **Cluster Status** changes to **Stopped**, the cluster is stopped successfully.

----End

Starting a Cluster

You can start a stopped cluster to restore cluster services.

NOTE

If the current console does not support this feature, contact technical support. Billing resumes after the cluster is started.

Procedure

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Clusters > Dedicated Clusters**.

Step 3 On the **Dedicated Clusters** page, locate the row that contains the target dedicated cluster, click **More > Start** in the **Operation** column.

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

The **Task Information** of the cluster changes to **Starting**. If the **Cluster Status** changes to **Available**, the cluster is started successfully.

----End

9.14 Resetting a Password

GaussDB(DWS) allows you to reset the password of the database administrator. If a database administrator forgets their password or the account is locked because the number of consecutive incorrect password attempts reaches the upper limit, the database administrator can reset the password on the **Clusters > Dedicated Clusters** page. After the password is reset, the account can be automatically unlocked. You can set the maximum number of incorrect password attempts (10 by default) by configuring the **failed_login_attempts** parameter on the **Parameter** page of the cluster. For details, see [Modifying Database Parameters](#).

Resetting a Password

- Step 1** Log in to the GaussDB(DWS) management console.
- Step 2** Choose **Cluster > Dedicated Cluster**.
- Step 3** In the **Operation** column of the target cluster, choose **More > Reset Password**.
- Step 4** On the displayed **Reset Password** page, set a new password, confirm the password, and then click **OK**.

The password complexity requirements are as follows:

- Contains 12 to 32 characters.
- Cannot be the username or the username spelled backwards.
- Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters (~!?,.,;_){}[]/<>@#%^&*+|\=-)
- Passes the weak password check.
- Cannot be the same as the old password and cannot be the reverse of the old password.
- Cannot use a historical password.

NOTE

If the default database administrator account of the cluster is deleted or renamed, password resetting fails.

----End

9.15 Cluster Upgrade

GaussDB(DWS) allows users to upgrade clusters on the console. For details, see [Upgrading a Cluster](#).

During cluster O&M operations, GaussDB(DWS) will send SMS notifications to keep you informed. It is important to be careful when performing any operations on the cluster during this time.

If a node needs to be replaced due to a hardware fault, the `repairCluster` event will be triggered. You can check the event progress by [Subscribing to Event Notifications](#).

During the upgrade, the cluster will be restarted and cannot provide services for a short period of time.


 NOTE

- After a cluster is upgraded to 8.1.3 or later, it enters the observation period. During this period, you can check service status and roll back to the earlier version if necessary.
- Upgrading the cluster does not affect the original cluster data or specifications.

Upgrade Version Description

The following figure shows the cluster version.

Figure 9-4 Version description



- **Service patch upgrade:** The last digit of cluster version $X.X.X$ is changed. For example, the cluster is upgraded from 1.1.0 to 1.1.1.
 - Duration: The whole process will take less than 10 minutes.
 - Impact on services: During this period, if the source version is upgraded to 8.1.3 or later, online patching is supported. During the patch upgrade, you do not have to stop services, but the services will be intermittently interrupted for seconds. If the destination version is earlier than 8.1.3, services will be interrupted for 1 to 3 minutes. Therefore, you are advised to perform this operation during off-peak hours.
- **Service upgrade:** The first two digits of cluster version $X.X.X$ are changed. For example, the cluster is upgraded from 1.1.0 to 1.2.0.
 - Duration: The whole process will take less than 30 minutes.
 - Impact on services: Online upgrade is supported for update to 8.1.1 or later. During the upgrade, you are not required to stop services, but services are intermittently interrupted for seconds. You are advised to perform the upgrade during off-peak hours.
- **Hot patch upgrade:** A hot patch upgrade involves adding a one-digit version number (in the format of **0001-9999**) to the current cluster version.
 - Duration: The upgrade of a single hot patch takes less than 10 minutes.

- Impact on services: The hot patch upgrade will not affect services, but there is a chance that the issues resolved by the current hot patch may come back after it is uninstalled.

Upgrading a Cluster

Prerequisites

Cluster 8.1.1 and later versions allow users to deliver cluster upgrade operations on the console.

Procedure

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the cluster list, click the name of a cluster.

Step 3 In the navigation pane, choose **Upgrade Management**.

Step 4 Choose either **Upgrade** or **Hot patch** from the **Type** drop-down list depending on the type of upgrade you want to perform.

Step 5 On the **Upgrade Management** page, select a version from the **Target Version** drop-down list.

NOTE

- Before the upgrade, it is crucial to verify if it meets the inspection conditions. Click **Immediate Inspection** to complete the inspection and proceed to the next step only if it passes. For more information, see [Viewing Inspection Results](#).
- DR cannot be established after a hot patch is installed in a cluster.

Step 6 Click **Upgrade**. Click **OK** in the displayed dialog box.

Step 7 Check whether the cluster is successfully upgraded.

- If the cluster version is 8.1.3 or later, the cluster enters the service observation period after the upgrade is complete. If you have verified your services, click **Submit** on the **Upgrade Management** page to complete the cluster upgrade. If you find your cluster performance affected by the upgrade, you can click **Rollback** to roll back the upgrade.

NOTE

- If you are using a version of 8.1.3 or earlier, you will not be able to roll back or submit upgrade tasks until the cluster upgrade is finished.
- After an upgrade task is successfully delivered, if the upgrade task is not submitted, the **wlm** thread occupies the system storage space and affects the system performance.
- If the cluster upgrade fails, click **Rollback** to roll back to the original cluster version, or click **Retry** to deliver the upgrade again.

----End

9.16 Associating and Disassociating ELB

Overview

If the private IP address or EIP of a CN is used to connect to a cluster, the failure of this CN will lead to cluster connection failure. If a private or public domain name is used for connection, the DNS service randomly selects a private IP address or EIP for each client. This cannot balance loads or avoid single-CN failures. ELB is used to solve these problems.

An ELB distributes access traffic to multiple ECSs for traffic control based on forwarding policies. It improves the fault tolerance capability of application programs. For details, see [Elastic Load Balance User Guide](#).

With ELB health checks, CN requests of a cluster can be quickly forwarded to normal CNs. If a CN is faulty, the workload can be immediately shifted to a healthy node, minimizing cluster access faults. Currently, ELBs can be bound in the same VPC or across VPCs.

NOTE

- This feature is supported only in cluster version 8.1.1.200 or later.
- For load balancing and high availability purposes, and to prevent single CN failures, a cluster must be bound to ELB.
- When you bind a cluster to ELB across VPCs, you can bind it to a dedicated load balancer.
- ELB does not support cross-database access.

Constraints and Limitations

- To bind an ELB to a GaussDB(DWS) cluster, the ELB must be in the same region, VPC, and enterprise project as the cluster.
- Only dedicated load balancers can be bound to GaussDB(DWS).

NOTICE

Load balancing is not supported in regions where the dedicated load balancer is not available. You can check whether dedicated load balancers are supported on the ELB console.

-
- The ELB to be associated must use TCP and has a private IP address.
 - When creating an ELB instance, determine its specifications based on your service access traffic. You are advised to select the maximum specifications. On the GaussDB(DWS) console, you can bind to an ELB instance but cannot change its specifications.
 - You only need to create a load balancer if you want to use ELB. GaussDB(DWS) automatically creates the required ELB listeners and backend server groups.
 - When creating a load balancer, ensure that the listeners do not use the same port as the database. Otherwise, ELB cannot be associated.

- When you associate ELB, the **ROUND_ROBIN** policy is set by default. In addition, the health check interval is set to 10 seconds, the timeout duration is set to 50 seconds, and the number of maximum retries is set to 3. Exercise caution when you modify these ELB parameters.
- When you bind a cluster to ELB across VPCs, you can only bind it to a dedicated load balancer.
- Before you bind a cluster to ELB across VPCs, ensure that the subnet segment of the cluster VPC is different from that of the ELB VPC.
- When you disassociate ELB from a cluster, related cluster information is cleared on GaussDB(DWS) but the load balancer is not deleted. Delete the load balancer in time to prevent unnecessary costs.
- If you need to access the ELB cluster using a public IP address or domain name, bind an EIP or domain name on the ELB management console.
- If the cluster is an IPv4 cluster, only IPv4 ELBs can be bound. If the cluster is an IPv6 dual-stack cluster, only IPv6 dual-stack ELBs can be bound.

Associating ELB

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters > Dedicated Clusters**. All clusters are displayed by default.

Step 3 In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.

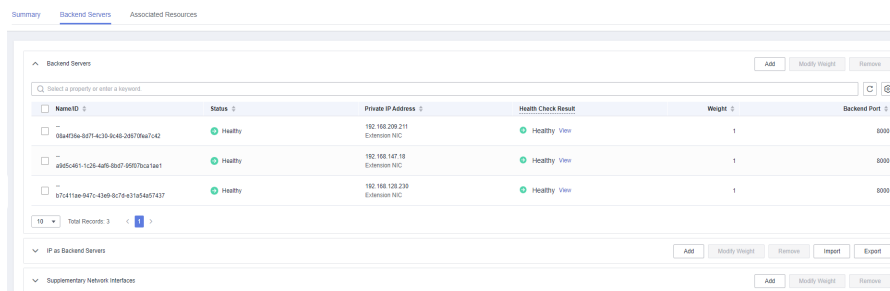
Step 4 On the **Basic Information** page that is displayed, click **Associate ELB** and select the ELB name. If no load balancer exists, create one on the ELB management console. Then refresh the GaussDB(DWS) page and associate ELB with the cluster.

NOTE

By default, the ELB in the VPC of the cluster is selected for GaussDB(DWS). If you select **Bind to ELB in another VPC**, the list of ELBs in other VPCs will be displayed for you to choose from. Before binding your cluster to an ELB across VPCs, ensure the cluster VPC has been connected to the ELB VPC. For details, see [Prerequisites for Binding an ELB to a Cluster Across VPCs](#).

Step 5 After the request is delivered, go back to the **Clusters** page. Task information **Associating ELB** of the cluster is displayed. The process takes some time.

Step 6 Log in to the ELB management console, choose **Elastic Load Balance > Load Balancers**, click the name of the bound load balancer, switch to the **Backend Server Groups** tab, and check whether the cluster CNs are associated with the load balancer.



NameID	Status	Private IP Address	Health Check Result	Weight	Backend Port
08a4f86-8d71-4c38-9c48-26879aa7c42	Healthy	192.168.209.211 Elasticson N.C.	Healthy View	1	8000
29950461-1c29-4a95-8b87-99579ca1aa1	Healthy	192.168.147.18 Elasticson N.C.	Healthy View	1	8000
07d411ee-9472-47a9-8c79-671a5a857437	Healthy	192.168.120.230 Elasticson N.C.	Healthy View	1	8000

NOTE

If the health check result indicates that the ELB backend nodes are deleted, ignore the problem.

Step 7 In the **Basic Information** area of the **Cluster Information** page, check the **ELB Address**, which is used for connecting to the cluster.

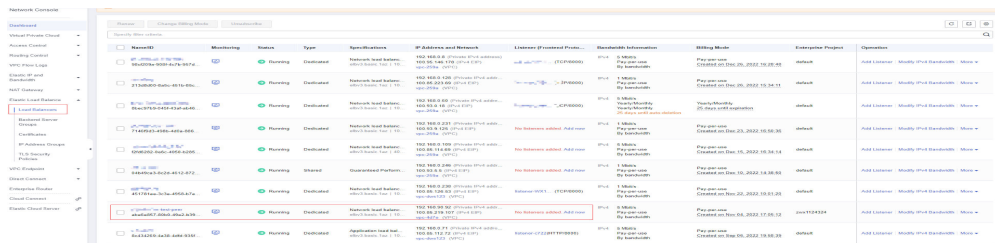
----End

Prerequisites for Binding an ELB to a Cluster Across VPCs

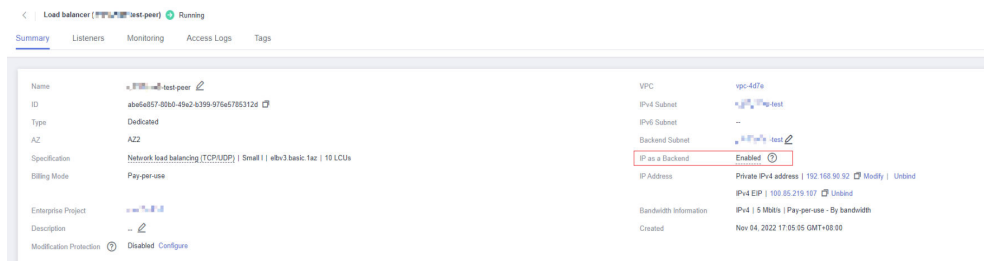
Enabling ELB for a cross-VPC backend server

Step 1 Log in to the ELB console.

Step 2 In the ELB list, click the name of a dedicated ELB to go to its details page.



Step 3 On the **Summary** page, enable **IP as a Backend**, confirm the information, and click **OK**.



Step 4 Check the VPC and subnet segment.



----End

Connecting the cluster VPC and the ELB VPC (by using VPC peering as an example)

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters > Dedicated Clusters**. All clusters are displayed by default.

Step 3 In the cluster list, click the name of a cluster to go to the cluster details page. Check the VPC and subnet segment of the cluster.

Network			
Region	North-Ulanqab203	AZ	AZ3
VPC	vpc-259a	Subnet	subnet-2600 (192.168.0.0/24)
Security Group	sg-9-8000		

- Step 4** Log in to the VPC management console. choose **My VPCs** in the navigation pane, and locate the VPC for which you want to create a VPC peering connection.
- Step 5** Choose **VPC Peering Connections**. In the upper right corner of the page, click **Create VPC Peering Connection**.
- Step 6** On the displayed page, set **Local VPC** to the cluster VPC, and set **Peer VPC** to the VPC of the ELB. Confirm the settings and click **OK**.

Create VPC Peering Connection

A VPC peering connection allows two VPCs to communicate with each other if they are in the same region. If you need two VPCs in different regions to communicate with each other, use [Cloud Connect](#).

Local VPC Settings

- Name: peering-68a3
- Local VPC: vpc-259a (Cluster VPC)
- Local VPC CIDR Block: 192.168.0.0/16

Peer VPC Settings

- Account: My account
- Peer Project: cn-north-7
- Peer VPC: vpc-4d7e (ELB VPC)
- Peer VPC CIDR Block: 192.168.0.0/16

OK Cancel

- Step 7** Click **Add Route** to add the route information.
- Step 8** Click the name of the created VPC peering connection. On the displayed page, click the **Local Routes** tab, click **Route Tables**, and add the route table of the cluster VPC.

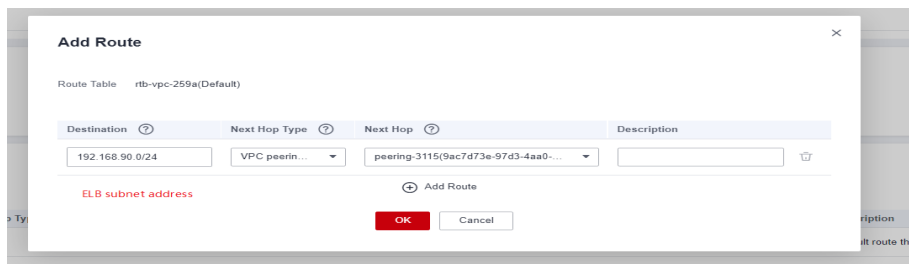
peering-3115

Basic Information Local Routes Peer Routes


Switch to the [Route Tables](#) page to add routes for the VPC peering connection.

Destination	Next Hop Type	Next Hop	Route Table
-------------	---------------	----------	-------------

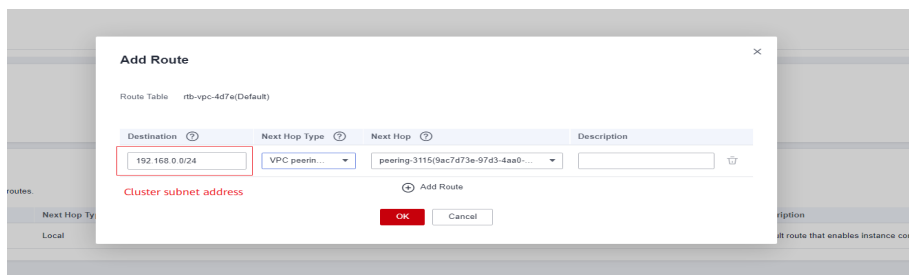
- Step 9** In the local route table, set **Destination** to the subnet CIDR block of the ELB VPC, set **Next Hop Type** to **VPC peering connection**, and set **Next Hop** to the created VPC peering connection. Click **OK**.



- Step 10** Go to the basic information page of the created VPC peering connection, click the **Peer Routes** tab, click **Route Tables**, and add the route table of the ELB VPC.



- Step 11** In the peer route table, set **Destination** to the subnet CIDR block of the cluster VPC, set **Next Hop Type** to **VPC peering connection**, and set **Next Hop** to the created VPC peering connection. Click **OK**.



- Step 12** After the cluster is created, the network between the VPC where the cluster resides and the VPC where the load balancer resides is connected. For details, see section [Binding an ELB](#).

----End

Disassociating ELB

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Clusters > Dedicated Clusters**. All clusters are displayed by default.
- Step 3** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 4** On the **Basic Information** page that is displayed, click **Disassociate ELB**.
- Step 5** After the request is delivered, go back to the **Clusters** page. Task information **Dissociating ELB** of the cluster is displayed. The process takes some time.
- Step 6** Log in to the ELB management console, click the name of the dissociated ELB, switch to the **Backend Server Groups** tab, and check whether the cluster CNs are deleted.

----End

9.17 Managing CNs

Purpose

After a cluster is created, the number of required CNs varies with service requirements. The CN management function enables you to adjust the number of CNs in the cluster. The operations are as follows:

- [Adding CNs](#)
- [Deleting CNs](#)

NOTE

- This feature is supported only in cluster version 8.1.1 or later.
- Only cluster versions 8.1.3.300 and later (excluding 8.2.0) support online CN addition, deletion, and concurrent addition of multiple CNs.

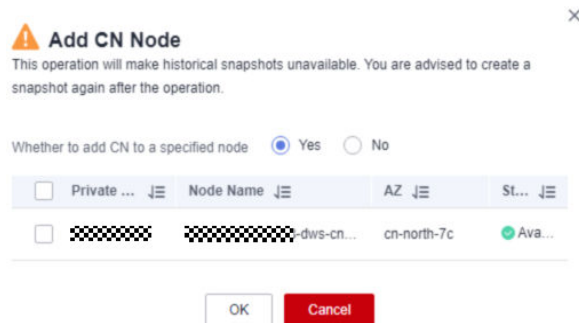
Constraints and Limitations

- During resource provisioning, the default number of CNs is 3. You can adjust the number of CNs based on the number of provisioned nodes. The number of CNs ranges from 2 to 20.
- Do not perform other O&M operations when adding or deleting a CN.
- Adding CNs consumes lots of CPU and I/O resources, which will greatly impact job performance. You are advised to perform this operation during off-peak hours or after services are stopped.
- If a fault occurs when you add a CN node and the rollback fails, try adding the CN again. The deletion of a CN node cannot be rolled back.
- For a CN that fails to be added, you can only retry the addition. For a CN that fails to be deleted, you can only retry the deletion. Other O&M operations are not allowed for such CNs.
- If DDL operations, such as schema and function creation, are performed during CN deletion, an error may be reported because the deleted CN cannot be found. In this case, try again.
- If one of your CNs is abnormal, you can only delete this abnormal CN. If two or more CNs are abnormal, you can delete CNs only after the CNs are recovered from faults.

Adding CNs

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** On the **Clusters > Dedicated Clusters** page, locate the cluster to which you want to add CNs.
- Step 3** In the **Operation** column of the specified cluster, choose **More > Manage CN > Add CN Node**.
- Step 4** On the displayed page, determine whether to add a CN to a specified node.
 - If you select **No**, you can set the **CN quantity after adjustment**.

- If you select **Yes**, specify the node.



NOTICE

- Before adding a CN, ensure that the cluster is in the **Available** or **Unbalanced** state.
- The number of CNs cannot exceed the total number of nodes after adjustment.
- You cannot add more CNs than the number of CNs that have already been deployed.

Step 5 Click **OK**.

----End

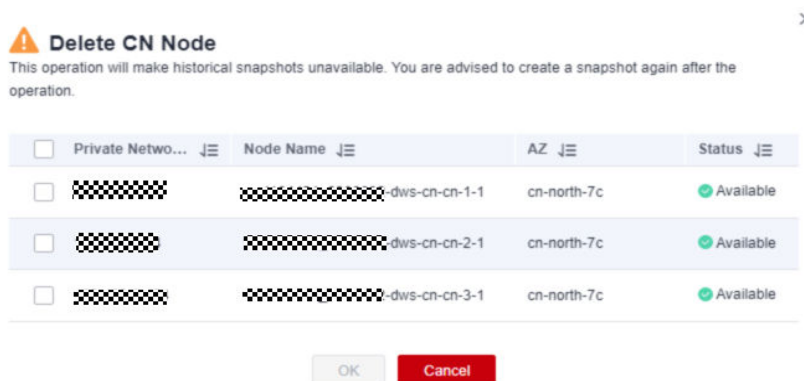
Deleting CNs

Step 1 Log in to the GaussDB(DWS) console.

Step 2 On the **Clusters > Dedicated Clusters** page, locate the cluster from which you want to delete CNs.

Step 3 In the **Operation** column of the specified cluster, choose **More > Manage CN > Delete CN Node**.

Step 4 On the displayed page, select the CN node to be deleted. After confirming that the information is correct, enter **DELETE** or click **One-Click Input** and click **OK** to delete the CN node.



NOTICE

- At least two CN must be retained.
 - When deleting a CN from a multi-AZ cluster, reserve a normal CN node in each AZ. Faulty CN nodes (if any) can be deleted.
 - When you delete a CN, the cluster must be in the **Available**, **Degraded**, or **Unbalanced** state.
 - If an elastic IP address has been bound to a CN, the CN cannot be deleted.
 - If abnormal nodes exist, only the abnormal CNs can be deleted.
 - If one CN is faulty, only this CN can be deleted.
 - If two or more CNs are faulty, no CN can be deleted.
-

----End

10 Data Integration

10.1 Overview

GaussDB(DWS) helps you migrate data from multiple sources and integrate diverse data sources, quick and easy. Currently, data can be migrated from Kafka and MRS to GaussDB(DWS).

 **NOTE**

This feature is supported only in 8.2.0 or later.

Supported Data Sources

Source Data Source	Destination Data Source	Description
Kafka	GaussDB(DWS)	-
MySQL	GaussDB(DWS)	-
Oracle	GaussDB(DWS)	-
IOT	GaussDB(DWS)	-
DWS	GaussDB(DWS)	-

10.2 Managing Instances

Overview

Data migration provides independent clusters for secure and reliable data migration. Clusters are isolated from each other and cannot access each other. With instance management, you can easily create and manage clusters by purchasing GDS-Kafka instances. GDS-Kafka consumes and caches data from Kafka. If the data cache time or size reaches a preconfigured threshold, GDS-Kafka

will copy the data to a GaussDB(DWS) temporary table, and then insert or update data in the temporary table.

1. The format of data generated by the Kafka message producer is specified by the **kafka.source.event.type** parameter. For details, see [Message Formats Supported by GDS-Kafka](#).
2. In GDS-Kafka, you can directly insert data for tables without primary keys, or update data by merging. Direct insert can achieve better performance, because it does not involve update operations. Determine your update mode based on the target table type in GaussDB(DWS). The data import mode is determined by the **app.insert.directly** parameter and whether a primary key exists. For details, see [GDS-Kafka Data Import Modes](#).

NOTE

- GDS-kafka only allows lowercase target table and column names.
- GDS-Kafka deletes historical data based on **pos** in the extended field. If imported data involves the delete operation, the extended field must be used.

Purchasing a GDS-Kafka Instance

To use the data migration feature, you need to purchase a GDS-kafka instance (cluster). Cluster instances provide secure and reliable data migration services. Clusters are isolated from each other.

Constraints

- Currently, only standalone clusters are supported.
- Only the pay-per-use billing mode is supported.

Procedure

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane, choose **Data > Data Integration > Instances**.
- Step 3** In the upper right corner of the page, click **Buy GDS-Kafka Instance**. Configure cluster parameters.

Table 10-1 Parameter description

Parameter	Description	Example Value
CPU Architecture	The following CPU architectures can be selected: <ul style="list-style-type: none">• x86• Kunpeng NOTE The only difference between the x86 and Kunpeng architectures lies in the underlying architecture, of which the application layer is unaware. The same SQL syntax is used. If x86 servers are sold out when you create a cluster, select the Kunpeng architecture.	x86

Parameter	Description	Example Value
Flavor	Select a node flavor.	-
Capacity	Storage capacity of a node.	-
Current Flavor	Current flavor of the cluster.	-
Name	Set the name of the data warehouse cluster. Enter 4 to 64 characters. Only case-insensitive letters, digits, hyphens (-), and underscores (_) are allowed. The value must start with a letter. Letters are not case-sensitive.	-
Version	Version of the database instance installed in the cluster. The version in the screenshot is for reference only.	8.2.1.300
VPC	Specify a VPC to isolate the cluster's network. If you create a data warehouse cluster for the first time and have not configured the VPC, click View VPC . On the VPC management console that is displayed, create a VPC as needed.	-
Subnet	Specify a VPC subnet. A subnet provides dedicated network resources that are isolated from other networks, improving network security.	-
Security Group	Specify a VPC security group. A security group restricts access rules to enhance security when GaussDB(DWS) and other services access each other.	-

Parameter	Description	Example Value
EIP	<p>Specify whether users can use a client to connect to a cluster's database over the Internet. The following methods are supported:</p> <ul style="list-style-type: none">• Do not use: Do not specify any EIPs here. If GaussDB(DWS) is used in the production environment, first bind it to ELB, and then bind it to an EIP on the ELB page.• Automatically assign: Specify bandwidth for EIPs, and the system will automatically assign EIPs with dedicated bandwidth to clusters. You can use the EIPs to access the clusters over the Internet. The bandwidth name of an automatically assigned EIP starts with the cluster name.• Specify: Specify an EIP to be bound to the cluster. If no available EIPs are displayed in the drop-down list, click Create EIP to go to the Elastic IP page and create an EIP as needed. The bandwidth can be customized.	-
Enterprise Project	Select the enterprise project of the cluster. You can configure this parameter only when the Enterprise Project Management service is enabled. The default value is default .	default

Step 4 Confirm the information and click **Submit**.

----End

Viewing Instance Details

On the instance details page, you can view the basic information and network information about the cluster.

Procedure

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Data > Data Integration > Instances**.

Step 3 Click the name of an instance to go to the instance details page.

----End

Message Formats Supported by GDS-Kafka

Table 10-2 Message formats supported by GDS-Kafka

Kafka Source Event Type	Format	Description
cdc.drs.avro	Internal format of Huawei Cloud DRS. DRS generates data in the avro format used by Kafka. GDS-Kafka can directly interconnect with DRS to parse and import the data.	None

Kafka Source Event Type	Format	Description
<p>drs.cdc</p>	<p>To use the avro format for drs.cdc, specify the Maven dependency of GDS-Kafka-common and GDS-Kafka-source in the upstream programs of Kafka, and then create and fill in the Record object. A Record object represents a table record. It will be serialized into a byte[] array, produced and sent to Kafka, and used by the downstream GDS-Kafka.</p> <p>In the following example, the target table is the person table in the public schema. The person table consists of the id, name, and age fields. The op_type is U, which indicates an update operation. This example changes the name field from a to b in the record with the ID 0, and changes the value of the age field from 18 to 20.</p> <pre>Record record = new Record(); // Set the schema and table name of the target table. record.setTableName("public.person"); // Set the field list. List<Field> fields = new ArrayList<>(); fields.add(new Field("id", 0)); fields.add(new Field("name", 1)); fields.add(new Field("age", 2)); record.setFields(fields); // Set the field value list before the table record is updated. List<Object> before = new ArrayList<>(); before.add(new Integer(0, "0")); before.add(new Character("utf-8", ByteBuffer.wrap("a".getBytes(StandardCharsets.UTF_8)))); before.add(new Integer(0, "18")); record.setBeforeImages(before); // Set the field value list after the table record is updated. List<Object> after = new ArrayList<>(); after.add(new Integer(0, "0")); after.add(new Character("utf-8", ByteBuffer.wrap("b".getBytes(StandardCharsets.UTF_8)))); after.add(new Integer(0, "20")); record.setAfterImages(after); // Set the operation type. record.setOperation("U"); // Set the operation time. record.setUpdateTimestamp(325943905); // Serialize the record object into a byte[] array. byte[] msg = Record.getEncoder().encode(record).array();</pre>	<p>Standard avro format:</p> <ul style="list-style-type: none"> • The tableName field is used to describe the target table and schema names that the current record belongs to. [Mandatory] • The operation field is used to describe the operation type of the current record. I indicates insert, U indicates update, and D indicates deletion. [Mandatory] • updateTimestamp indicates the time when an operation is performed on the source end. [Optional] • The beforeImages list describes the information before the current record is updated or deleted. The fields in the before body correspond to those in the target table. [Mandatory for U/D] • The afterImages list describes the updated or

Kafka Source Event Type	Format	Description
		<p>newly inserted information of the current record. [Mandatory for U/D]</p> <ul style="list-style-type: none"> • The fields list describes the field list of the current table record. The index values of the fields must be in the same sequence as those in beforeImage and afterImage. [Mandatory]

Kafka Source Event Type	Format	Description
cdc.json	<p>In the following example, the target table is the person table in the public schema. The person table consists of the id, name, and age fields. The op_type is U, which indicates an update operation. This example changes the name field from a to b in the record with the ID 1, and changes the value of the age field from 18 to 20.</p> <pre data-bbox="644 667 1152 1077"> { "table": "public.person", "op_type": "U", "op_ts": "1668426344", "current_ts": "1668426344", "before": { "id": "1", "name": "a", "age": 18 }, "after": { "id": "1", "name": "b", "age": 20 } } </pre>	<p>Standard JSON format:</p> <ul style="list-style-type: none"> • The table field describes the target table and schema names that the current record belongs to. [Mandatory] • The op_type field is used to describe the operation type of the current record. I indicates insert, U indicates update, and D indicates deletion. [Mandatory] • op_ts indicates the time when an operation is performed on the source end. [Optional] • current_ts indicates the time when a message is imported to Kafka. [Optional] • The before object describes the information before the current record is updated or deleted. The fields in the before body correspond to those in the target table. [Mandatory for U/D]

Kafka Source Event Type	Format	Description
		<ul style="list-style-type: none"> The after object list describes the update or newly inserted information of the current record. [Mandatory for U/D]
industrial.iot.json	<pre>{ "header": { "thing_id": "a0001", "instance_id": "1", "thing_model_name": "computer", "timestamp": "1668426344" }, "body": { "status": "Normal", "temperature": "10", "working_time": "10000" }, }</pre>	<p>IoT data format:</p> <ul style="list-style-type: none"> thing_model_name in header indicates the table name. [Mandatory] The values of thing_id, instance_id, and timestamp in header and the content in the body comprise the fields of the current record. IoT data is time series data and does not involve update or deletion. Only insert operations are involved.

Kafka Source Event Type	Format	Description
industrial.iot.recursion.json	<pre> { "header": { "thing_id": "a0001", "instance_id": "1", "thing_model_name": "computer", "timestamp": "1668426344" }, "body": { "status": "Normal", "temperature": "10", "property": { "key1": "1", "key2": "2" }, "working_time": "10000" }, }, } </pre>	IoT data format: <ul style="list-style-type: none"> • thing_model_name in header indicates the table name. [Mandatory] • The values of thing_id, instance_id, and timestamp in header and the content in the body comprise the fields of the current record. • IoT data is time series data and does not involve update or deletion. Only insert operations are involved. • In this data format, the key and value of body are added to the property and value fields in the new format to generate multiple pieces of new data. In this way, rows are converted to columns.

Kafka Source Event Type	Format	Description
industrial.iot.event.json.independent.table	<pre>{ "event_id": "1", "event_name": "test", "start_time": "1970-1-1T00:00:00.000Z", "end_time": "1970-1-1T00:00:00.000Z", "fields": { "field1": "value1", "field2": 2 } }</pre>	<p>IoT event stream data format:</p> <ul style="list-style-type: none"> • event_name indicates a table name. [Mandatory] • event_id, start_time, end_time, and fields comprise the field content of a record. [Mandatory] • IoT event stream data is time series data and does not involve update or deletion. Only insert operations are involved.

Kafka Source Event Type	Format	Description
industrial.iot.json. multi.events	<pre> { "event_id": "1", "event_name": "test", "start_time": "1970-1-1T00:00:00.000Z", "end_time": "1970-1-1T00:00:00.000Z", "fields": { "field1": "value1", "field2": 2, "field3": { "key1": "1", "key2": 2 } } } </pre>	IoT event stream data format: <ul style="list-style-type: none"> event_name indicates a table name. [Mandatory] event_id, start_time, end_time, and fields comprise the field content of a record. [Mandatory] IoT event stream data is time series data and does not involve update or deletion. Only insert operations are involved. In this data format, the key and value of fields are added to the field_name and field_value fields in the new format to generate multiple pieces of new data. In this way, rows are converted to columns.

GDS-Kafka Import Modes

To import GDS-Kafka data to the database, copy the data to a temporary table, and then merge or insert the data. The following table describes their usage and scenarios.

Table 10-3 GDS-Kafka import modes

Operation	Direct Insertion	Primary Key Table	Import Mode
insert	true (only for tables without primary keys)	No	Use INSERT SELECT to write data from the temporary table to the target table.
	false	Yes	Merge data from the temporary table to the target table based on the primary key.
		No	Use INSERT SELECT to write data from the temporary table to the target table.
delete	true (only for tables without primary keys)	No	Use INSERT SELECT to write data from the temporary table to the target table.
	false NOTE You can mark deletion by configuring the app.del.flag parameter. The flag of a deleted record will be set to 1 .	Yes	<ul style="list-style-type: none"> If the delflag field is set, merge will be performed based on the primary key. If a matched primary key is found, and the value of pos in the target table is smaller than that in the temporary table, the delflag field will be set to 1. Otherwise, a new record will be inserted. If the delflag field is not set, a matched primary key is found, and the value of pos in the target table is smaller than that in the temporary table, the record will be deleted from the target table.

Operation	Direct Insertion	Primary Key Table	Import Mode
		No	<ul style="list-style-type: none"> • If the delflag field is set, all the fields in the temporary table will be used to match and merge with the target table. If a matched record is found, and the value of pos in the target table is smaller than that in the temporary table, the delflag field will be set to 1. Otherwise, a new record will be inserted. • If the delflag field is not set, all the fields in the temporary table will be used to match the target table. If a matched record is found, and the value of pos in the target table is smaller than that in the temporary table, the matched record will be deleted from the target table.
update	true (only for tables without primary keys)	No	Use INSERT SELECT to write data from the temporary table to the target table.
	false NOTE The update operation is split. The message in before or beforeImage is processed as a delete operation, and the message in after or afterImage is processed as an insert operation. Then, the message is saved to the database based on the insert and delete operations.	Yes	Equivalent to the insert+delete operation on a table with a primary key.
		No	Equivalent to the insert+delete operation on a table without a primary key.

10.3 Managing Connections

Description

Before creating a data migration task, you need to create a connection, so that the cluster can read and write the data source. A migration job requires a source connection and a destination connection. Data sources that support exporting are used as source connections and data sources that support importing are used as destination connections.

The connection parameters you can configure vary according to the data source. This section describes how to create these connections.

Prerequisites

- A GDS-kafka cluster has been created.
- The GDS-kafka cluster can communicate with the destination data source.
 - If the destination data source is an on-premises database, you need the Internet or Direct Connect. If the Internet is used for communication, ensure that an EIP has been bound to the GDS-kafka cluster, the security group of GDS-kafka allows outbound traffic from the host where the off-cloud data source is located, the host where the data source is located can access the Internet, and the connection port has been enabled in the firewall rules.
 - If the destination data source is a cloud service, the following requirements must be met for network interconnection:
 - If the GDS-kafka cluster and the cloud service are in different regions, the Internet or a Direct Connect is required for enabling communication between the CDM cluster and the cloud service. If the Internet is used for communication, ensure that an EIP has been bound to the GDS-kafka cluster, the host where the data source is located can access the Internet, and the port has been enabled in the firewall rules.
 - If the GDS-kafka cluster and the cloud service are in the same region, VPC, subnet, and security group, they can communicate with each other by default. If they are in the same VPC but in different subnets or security groups, you must configure routing rules and security group rules.
 - The cloud service instance and the cluster belong to the same enterprise project. If they do not, you can modify the enterprise project of the workspace.
- You have obtained the URL, account, and password for accessing the destination data source. The account is granted with the read and write permissions on the data source.

Creating a Connection

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane, choose **Data > Data Integration > Connections**.
- Step 3** Click **Create Connection**.
- Step 4** Configure connection parameters. For more information, see [Connection parameters](#).

Table 10-4 Connection parameters

Protocol	Parameter	Mandatory	Description
Kafka	Connection Name	Yes	Connection name, which can be customized. Only letters, numbers, underscores (_), and hyphens (-) are allowed.
	Type	Yes	Currently, MRS-Kafka, IoT-Kafka, DMS-Kafka, and Default-Kafka are supported. Default-Kafka is an open-source Kafka.
	Service Address	Yes	Kafka connection address. Format: Domain name + Port number or IP address + Port number
	Topics	Yes	List of Kafka topics, which are separated by commas (,).
	Ciphertext Access	No	This function must be enabled during Kafka authentication. The SASL_SSL and SASL_PLAINTEXT protocols are supported.
	User	No	Username for connecting to Kafka
	Password	No	Password for connecting to Kafka.
	SSL Authentication	No	Whether the SSL protocol is supported.
	Certificate	No	SSL certificate in binary JKS format.
	Certificate Password	No	Certificate encryption password.

Protocol	Parameter	Mandatory	Description
	Host Configuration	No	<p>MRS-Kafka configuration parameter. When you connect to MRS-Kafka in security mode, you need to configure the host file of the VM where Gds-Kafka resides. Therefore, you need to upload the host file to be modified. The file format can only be TXT. The file content is as follows:</p> <pre>192.168.4.111 node-master1JuQr.mrs-yd8z.com 192.168.4.204 node-master3mgqy.mrs-yd8z.com 192.168.4.221 node-master2Ktgg.mrs-yd8z.com</pre> <p>The information on the left is the IP address of the Kafka broker. If MRS-Kafka and GDS-Kafka are not in the same VPC, replace the IP address with a public IP address. The information on the right is the host name of the broker. You can log in to FusionInsightManage and access the Kafka cluster to obtain the host name corresponding to the broker instance.</p>
	Security mode	No	MRS-Kafka configuration parameter. When the security mode is enabled, Kerberos authentication is required.
	Krb5 File	No	<p>MRS-Kafka configuration parameter. When the security mode is enabled, you need to upload the krb5 file. This file is the authentication credential of the machine-machine account applied for on FusionInsight Manager of MRS.</p> <p>NOTE If MRS-Kafka and GDS-Kafka are not in the same VPC, replace the internal IP address of the broker in the file with the public IP address.</p>

Protocol	Parameter	Mandatory	Description
	Keytab File	No	MRS-Kafka configuration parameter. When the security mode is enabled, you need to upload the Keytab file. This file is the authentication credential of the machine-machine account applied for on FusionInsight Manager of MRS.
	Account	No	MRS-Kafka configuration parameter. It is a machine-machine account applied for on FusionInsight Manager of MRS.
	SSL	No	MRS-Kafka configuration parameter. When SSL is enabled, you need to upload the SSL certificate and key.
	Authentication Mechanism	No	DMS-Kafka configuration parameter. It indicates the security authentication protocol.
MySQL	Connection Name	Yes	Connection name, which can be customized. Only letters, numbers, underscores (_), and hyphens (-) are allowed.
	Service Address	Yes	MySQL connection address. Format: Domain name + Port number or IP address + Port number
	User	Yes	Username for logging in to the database.
	Password	Yes	Password used to log in to the database.
	Database	Yes	MySQL database name.
Oracle	Connection Name	Yes	Connection name, which can be customized. Only letters, numbers, underscores (_), and hyphens (-) are allowed.

Protocol	Parameter	Mandatory	Description
	Service Address	Yes	Oracle connection address. Format: Domain name + Port number or IP address + Port number
	User	Yes	Username for logging in to the database.
	Password	Yes	Password used to log in to the database.
	Database	Yes	Oracle database name.
	Schema	Yes	Schema name. You can configure one or more schema names and use commas (,) to separate them.
IoT	Service Address	Yes	Address of the <code>iot-edge-node</code> page. Format: Domain name or IP address
	User	Yes	Account for logging in to the IoT platform.
	Password	Yes	Password for logging in to the IoT platform.
DWS	Connection Name	Yes	Connection name, which can be customized. Only letters, numbers, underscores (_), and hyphens (-) are allowed.
	Service Address	Yes	GaussDB(DWS) connection address.
	User	Yes	Username for logging in to the database.
	Password	Yes	Password used to log in to the database.
	Database	Yes	GaussDB(DWS) database name.
	Schema	Yes	Name of a schema in the GaussDB(DWS) database.

Step 5 Confirm the information and click **OK**.

----End

Modifying a Connection

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Data > Data Integration > Connections**.

Step 3 In the **Operation** column of a connection, click **Modify**.

Step 4 In the **Modify Connection** dialog box, configure **Connection Name**, **Connection Address**, **Topics**, **User**, and **Password**.

Step 5 Confirm the information and click **OK**.

----End

Deleting a Connection

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Data > Data Integration > Connections**.

Step 3 In the **Operation** column of a connection, click **Delete**.

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

----End

10.4 Managing Table Mappings

Mapping Overview

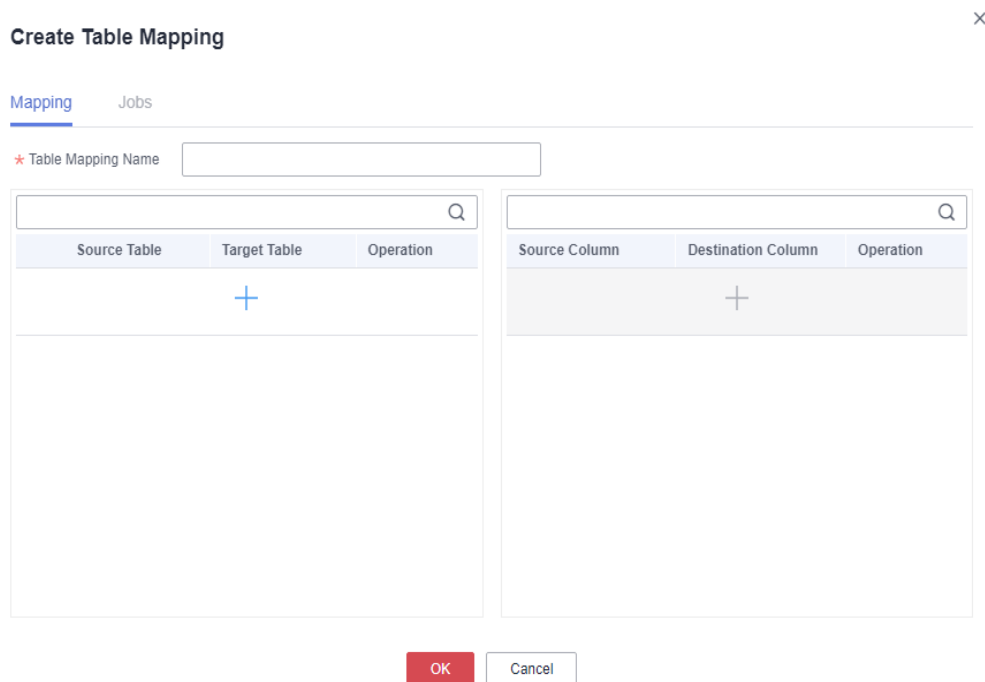
Before creating a job, you need to create a mapping to map the table structures of the source and destination databases, facilitating data migration between databases.

Creating a Table Mapping



Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Data > Data Integration > Table Mappings**.

Step 3 Click **Create Table Mapping**.



Step 4 Configure parameters.

1. Click  in the list on the left. Configure **Table Mapping Name**, **Source Table**, and **Target Table**.
2. Click  in the list on the right and configure the parameters.

 **NOTE**

If no column mappings are specified in the list on the right, all the columns with the same name will be mapped with by default.

Step 5 Confirm the information and click **OK**.

----End

Modifying a Table Mapping

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Data > Data Integration > Table Mappings**.

Step 3 In the **Operation** column of a table mapping, click **Modify**.

Step 4 In the **Modify Table Mapping** dialog box, configure **Table Mapping Name**, **Source Table**, and **Target Table**.

Step 5 Confirm the information and click **OK**.

----End

Checking a Table Mapping

- Step 1** Log in to the GaussDB(DWS) console.
 - Step 2** In the navigation pane, choose **Data > Data Integration > Table Mappings**.
 - Step 3** In the **Operation** column of a table mapping, click **Modify**.
 - Step 4** In the **Modify Table Mapping** dialog box, click the **Jobs** tab to view the bound jobs.
- End

Deleting a Table Mapping

- Step 1** Log in to the GaussDB(DWS) console.
 - Step 2** In the navigation pane, choose **Data > Data Integration > Table Mappings**.
 - Step 3** In the **Operation** column of a table mapping, click **Delete**.
 - Step 4** In the displayed dialog box, click **OK**.
- End

10.5 Managing Jobs

After creating a cluster instance, you can customize a job, enable a job, and migrate data.

You can create jobs to migrate data or automatically create tables.

- Data migration: Data is migrated from Kafka to GaussDB(DWS).
- Automatic table creation: Tables and fields in the source database are synchronized to GaussDB(DWS), but data is not migrated.

Creating a Job

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane, choose **Data > Data Integration > Instances**.
- Step 3** Click the name of an instance to go to the details page.
- Step 4** In the navigation pane, click **Manage Job**.
- Step 5** Click **Data Migration** or **Create Table**. (By default, the **Kafka Connection** parameter cannot be configured if you click **Create Table**.)
- Step 6** Enter the job name, configure **Kafka Connection**, **DWS Cluster Connection**, and **Customized Table/Field Mapping**, and click **Test Connection**.
- Step 7** Check to ensure the connection passes the test, and click **Next**.
- Step 8** Click **Next** and confirm the settings.
- Step 9** Click **OK**.

Step 10 Return to the job list. In the **Operation** column of the job, click **Start**. For details, see [Starting a Job](#).

----End

Viewing Job Details


Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Data > Data Integration > Instances**.

Step 3 Click the name of an instance to go to the details page.

Step 4 In the navigation pane, click **Manage Job**.

Step 5 Click a job name to go to the details page. Check the job information, including the connections, service parameters, and table/column mappings.



The diagram illustrates the data integration process. It shows three components: DWS (Data Warehouse Service), GaussDB(DWS), and KafkaServer. Arrows indicate the flow of data from DWS to GaussDB(DWS), and from GaussDB(DWS) to KafkaServer. Both connections are marked with a green checkmark, indicating they are successful.

DWSConnection		KafkaConnection	
Connection Name	Test_DWS	Connection Name	Test_Kafka
Connection Address		Connection Address	
Database	gaussdb	Topics	test
Username	dbadmin		
Schema Name	public		

Parameter Group	Parameter	Value	Value Range	Mandatory	Description
Kafka source	kafka source event type	cdc-ds-avro	cdc-ds-avro;cdc-periodic-join;am-join;industrial-ud-join	Yes	Data format in Kafka. Options: cdc-ds-avro; cdc-join; cdc-periodic-join; am-join; industrial-ud-join
	kafka.partition	--	--	No	Partitions consumed by each Kafka topic. If this parameter is not specified, the default value is 1.
Kafka consumer	kafka.consumer.group	--	--	No	User-defined consumer group name. Multiple names are allowed.
	kerberos.domain.name	--	--	No	Kerberos service domain name. This parameter is mandatory for PLAINTEXT and SASL_PLAINTEXT security protocols.
Kafka connection group	security.protocol	PLAINTEXT	PLAINTEXT;SASL_PLAINTEXT	No	Authentication protocol.
	sasl.kerberos.service.name	--	--	No	Service name used by Gds-Kafka for Kerberos authentication.
	java.security.krb5.conf	--	--	No	Address of the krb5 configuration file.
	principal	--	--	No	Kerberos username used by Gds-Kafka for Kerberos authentication.

----End

Starting a Job

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Data > Data Integration > Instances**.

Step 3 Click the name of an instance to go to the details page.

Step 4 In the navigation pane, click **Manage Job**.

Step 5 In the **Operation** column of a job, click **Start**.

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

----End

Stopping a Job

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Data > Data Integration > Instances**.

Step 3 Click the name of an instance to go to the details page.

Step 4 In the navigation pane, click **Manage Job**.

Step 5 In the **Operation** column of a job, click **Stop**.

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

----End

Deleting a Job

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane, choose **Data > Data Integration > Instances**.

Step 3 Click the name of an instance to go to the details page.

Step 4 In the navigation pane, click **Manage Job**.

Step 5 In the **Operation** column of a job, click **Delete**.

Step 6 Click **OK**.

----End

11 Cluster Log Management

Overview

Cluster logs are collected and sent to Log Tank Service (LTS). You can check or dump the collected cluster logs on LTS.

The following log types are supported: CN logs, DN logs, OS messages logs, audit logs, cms logs, gtm logs, Roach client logs, Roach server logs, upgrade logs, and scale-out logs.

NOTE

- Only 8.1.1.300 and later versions support cluster log management.
- Only 8.3.0 and later versions support CMS logs, GTM logs, Roach client logs, Roach server logs, scaling logs, and upgrade logs.

Enabling LTS

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters** > **Dedicated Clusters**. All clusters are displayed by default.

Step 3 Click the name of the target cluster. Choose **Logs**.



Log Type	Description	Operation
messages	operating system messages log	View Log
expand	dms-expand log	View Log
roach_controller	dms-roach-controller log	View Log
audit	audit log	View Log
gtm	dms-gtm log	View Log
roach_agent	dms-roach-agent log	View Log
cms	dms-cms log	View Log
CN	dms-CN-node log	View Log
upgrade	dms-upgrade log	View Log
DN	dms-DN-node log	View Log

Step 4 On the **Logs** tab, enable LTS. If LTS is enabled for the first time, the following dialog box will be displayed. Confirm the information and click **Yes**.

 **NOTE**

- If LTS has been enabled and authorized to create an agency, no authorization is required when LTS is enabled again.
- By default, only Huawei Cloud accounts or users with Security Administrator permissions can query and create agencies. IAM users under an account do not have the permission to query or create agencies by default. Contact a user with that permission and complete the authorization on the current page.
- When interconnecting with LTS, you need to grant LTS-related permission policies (**LTS Admin**, **LTS Administrator**, **LTS FullAccess**, and **LTS ReadOnlyAccess**) to users.

----End

Checking Cluster Logs

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters > Dedicated Clusters**. All clusters are displayed by default.

Step 3 Click the name of the target cluster. Choose **Logs**.

Step 4 On the **Logs** tab, click **View Log** in the **Operation** column of a log type to go to the Log Tank Service (LTS) page and view logs.



Log Type	Description	Operation
messages	operating system messages log	View Log
expand	dmis-expand log	View Log
reach-controller	dmis-reach-controller log	View Log
audit	audit Log	View Log
gfm	dmis-gfm log	View Log
reach-agent	dmis-reach-agent log	View Log
cms	dmis-cms log	View Log
CH	dmis-CH node log	View Log
upgrade	dmis-upgrade log	View Log
DH	dmis-DH node log	View Log

----End

Disabling LTS

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters > Dedicated Clusters**. All clusters are displayed by default.

Step 3 Click the name of the target cluster. Choose **Logs**.

Step 4 Toggle off the LTS switch.

Step 5 Click **OK** in the dialog box.

----End

12 Database User Management

12.1 Managing Users

GaussDB(DWS) allows you to manage database users on the console. You can create, delete, and update database users and manage their permissions on the console.

Constraints and Limitations

- If the current console does not support this feature, contact technical support.
- After a cluster is created, the users or roles created with it cannot be modified.
- Before using this function, ensure that the cluster is available.

Creating a User

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 2** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 3** In the navigation pane, choose **User Management**.
- Step 4** On the **Users** tab, click **Create User**. The user creation page is displayed.
- Step 5** The Basic Settings page is displayed. The parameters are described as follows:

Table 12-1 Basic parameters of user information

Parameter	Description	Example Value
User Name	The value must start with a letter and can contain a maximum of 63 characters, including letters, digits, and underscores (_).	Dws-demo

Parameter	Description	Example Value
Password	The value must start with a letter and can contain a maximum of 63 characters, including letters, digits, and underscores (_). NOTE Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters (~!?,.,:;_(){}[]/<>@#%^&*+ \ \=-)	-
Maximum Connections	Maximum number of connections between the user and the database. The value -1 indicates that the number of connections is not limited.	-1
Expires	Expiration time of the user's permissions.	-
System Administrator	Indicates whether the user is a system administrator.	-
Create Database	Specifies whether the user has the permission to create databases.	-
Create Role	Specifies whether the user has the permission to create users and roles.	-
Inherit Permissions	Specifies whether the user inherits permissions from its user group. This function is enabled by default. You are advised to retain this setting.	-

Step 6 If all the parameters are correctly set, click **Next**.

Step 7 On the **Configure Roles** page, select the role to be assigned to the user and click **Next**.

Step 8 Configure permissions not included in the roles of the user.

Click **Add** to add a permission configuration. Select the database object type and corresponding database object, and select the permission to complete assignment. For details about permission definitions, see "DCL Syntax" > "GRANT" in the [GaussDB\(DWS\) SQL Overview](#).

Step 9 After the authorization is complete, click **Create**.

----End

Modifying a User

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 2** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 3** In the navigation pane, choose **User Management**.
- Step 4** In the user list, select a user and click **Modify**. The page for modifying user details is displayed.
- Step 5** Modify the user information. For details, see [Table 12-1](#). After confirming that the information is correct, click **Next**.
- Step 6** Select the role to be granted to the user and click **Next**.
- Step 7** After selecting a permission type, you can click **Modify** in the row where the permission is located to add or remove the permission.
- Step 8** Confirm the permissions. Click **Save**.

----End

Exporting a User

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 2** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 3** In the navigation pane, choose **User Management**.
- Step 4** Click **Export** in the upper part of the user list and select the number of records to be exported to export the user list.

Figure 12-1 Exporting a user



- Step 5** Confirm the configurations and click **Export**.

----End

Exporting Permissions

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.

- Step 2** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 3** In the navigation pane, choose **User Management**.
- Step 4** Select a user from the user list and click **Export Permissions** to export the user permission list.

Figure 12-2 Exporting permissions



----End

Deleting a User

Prerequisites

To prevent any problems with deleting a user, check for dependencies between database objects (such as tables) beforehand. If there are any dependencies, delete them first before proceeding with the user deletion.

Procedure

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 2** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 3** In the navigation pane, choose **User Management**.
- Step 4** Select a user from the user list and click **Delete**. A confirmation dialog box is displayed.
- Step 5** Click **OK**.

----End

12.2 Managing Roles

Creating a Role

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 2** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 3** In the navigation pane, choose **User Management**.
- Step 4** Click the **Roles** tab and click **Create Role**. The role creation page is displayed.

Step 5 Configure role information. The parameters are described as follows:

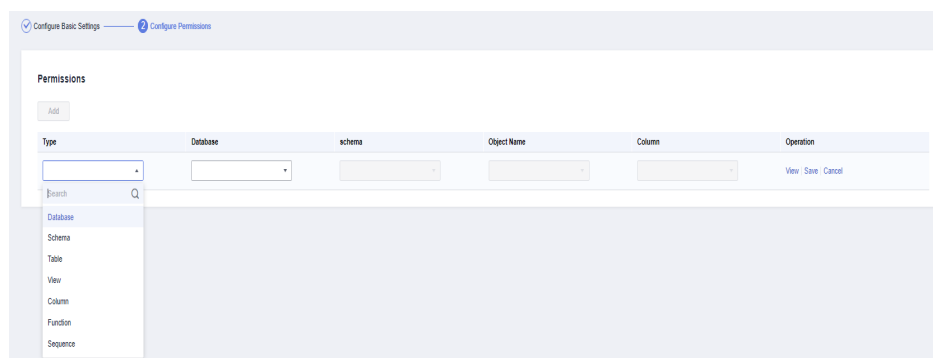
Table 12-2 Parameters for configuring role information

Parameter	Description	Example Value
Role Name	The value must start with a letter and can contain a maximum of 63 characters, including letters, digits, and underscores (_).	dws-demo
Expires	Expiration time of the role permissions.	-
System Administrator	Indicates whether the role has the system administrator rights.	-
Create Database	Specifies whether the role has the permission to create databases.	-
Create Role	Specifies whether the role has the permission to create users and roles.	-
Inherit Permissions	Indicates whether the role inherits the permissions from its role group. This function is enabled by default. You are advised to retain this setting.	-

Step 6 If all the parameters are correctly set, click **Next**.

Step 7 Configure the permissions of the role.

Click **Add** to add a permission configuration. Select the database object type and the corresponding objects. Then, select permissions. For details about permission definitions, see "DCL Syntax" > "GRANT" in the [GaussDB\(DWS\) SQL Overview](#).



Step 8 After the authorization is complete, click **Create**. The role is created.

----End

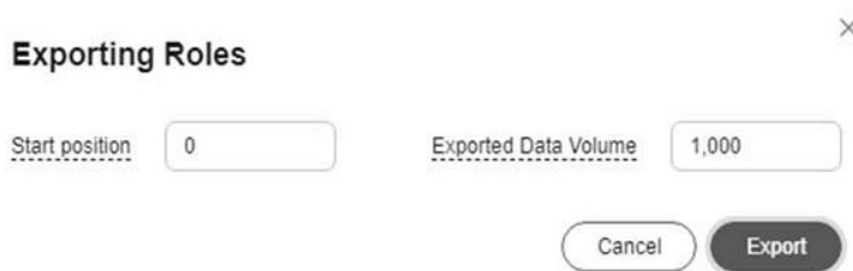
Modifying a Role

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
 - Step 2** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
 - Step 3** In the navigation pane, choose **User Management**.
 - Step 4** In the role list, select a user and click **Modify**. The page for modifying role details is displayed.
 - Step 5** Modify the role information. For the parameter description, see [Table 12-2](#).
 - Step 6** If all the parameters are correctly set, click **Next**.
 - Step 7** Configure permissions. Select a permission as required and click **Modify** in the row where the permission is located to add or remove the permission.
 - Step 8** Confirm the permissions. Click **Save**.
- End

Exporting Roles

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 2** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 3** In the navigation pane, choose **User Management** and click **Roles** to switch to the role list page.
- Step 4** Click **Export** in the upper part of the role list and select the number of roles to be exported.

Figure 12-3 Exporting roles



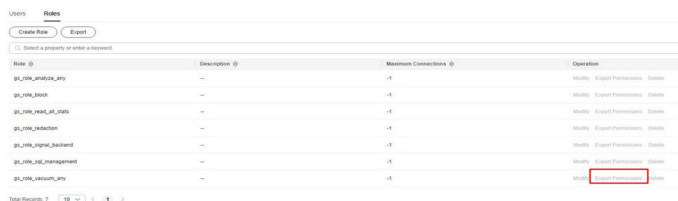
- Step 5** Confirm the information and click **Export**.
- End

Exporting role permissions

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.

- Step 2** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 3** In the navigation pane, choose **User Management** and click **Roles** to switch to the role list page.
- Step 4** Select a user from the role list, click **Export Permissions**, and select the number of records to be exported.

Figure 12-4 Exporting role permissions



Role ID	Description	Maximum Connections	Operation
gaussdb_role_analyze		-1	View Export Permissions Delete
gaussdb_role_backup		-1	View Export Permissions Delete
gaussdb_role_etl		-1	View Export Permissions Delete
gaussdb_role_etl_job		-1	View Export Permissions Delete
gaussdb_role_etl_job_admin		-1	View Export Permissions Delete
gaussdb_role_etl_job_executor		-1	View Export Permissions Delete
gaussdb_role_etl_job_monitor		-1	View Export Permissions Delete
gaussdb_role_etl_job_scheduler		-1	View Export Permissions Delete
gaussdb_role_etl_job_trigger		-1	View Export Permissions Delete
gaussdb_role_etl_job_trigger_admin		-1	View Export Permissions Delete
gaussdb_role_etl_job_trigger_executor		-1	View Export Permissions Delete
gaussdb_role_etl_job_trigger_monitor		-1	View Export Permissions Delete
gaussdb_role_etl_job_trigger_scheduler		-1	View Export Permissions Delete
gaussdb_role_etl_job_trigger_trigger		-1	View Export Permissions Delete

----End

Deleting a Role

Prerequisites

To prevent any problems with deleting a role, check for dependencies such as database objects beforehand. If there are any dependencies, delete them first before proceeding with the role deletion.

Procedure

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 2** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 3** In the navigation pane, choose **User Management**.
- Step 4** Select a role from the role list and click **Delete**. A confirmation dialog box is displayed.
- Step 5** Click **OK** to delete the role.

----End

13 Audit Logs

13.1 Audit Log Overview

GaussDB(DWS) provides management console audit logs and database audit logs for users to query service logs, analyze problems, and learn product security and performance status.

Management Console Audit Logs

GaussDB(DWS) uses Cloud Trace Service (CTS) to record mission-critical operations performed on the GaussDB (DWS) management console, such as cluster creation, snapshot creation, cluster scale-out, and cluster restart. The logs can be used in purposes such as security analysis, compliance audit, resource tracing, and fault locating.

For details about how to enable and view management console audit logs, see [Management Console Audit Logs](#).

Database Audit Logs

If the **Security** function is enabled, GaussDB(DWS) records any DML and DDL operations performed by the database. You can locate and analyze faults based on the database audit logs, and perform behavior analysis and security auditing on historical database operations to improve GaussDB (DWS) security.

For details about how to enable and view database audit logs, see [Configuring the Database Audit Logs](#) and [Viewing Database Audit Logs](#).

13.2 Management Console Audit Logs

Enabling CTS

A tracker will be automatically created after CTS is enabled. All traces recorded by CTS are associated with a tracker. Currently, only one tracker can be created for each account.

- Step 1** Log in to the management console, choose **Service List > Management & Governance > Cloud Trace Service**. The CTS management console is displayed.
- Step 2** In the navigation tree on the left, click **Trackers**.
- Step 3** Enable CTS.

If you have enabled CTS, the system has automatically created a management tracker. Only one management tracker can be created and it cannot be deleted.

----End

Disabling the Audit Log Function

If you want to disable the audit log function, disable the tracker in CTS.

- Step 1** Log in to the management console, choose **Service List > Management & Governance > Cloud Trace Service**. The CTS management console is displayed.
- Step 2** Disable the audit log function by disabling the tracker. To enable the audit log function again, you only need to enable the tracker.

----End

Key Operations

With CTS, you can record operations associated with GaussDB(DWS) for later query, audit, and backtrack operations.

NOTE

The creation and deletion of automatic snapshots are not performed by users, therefore not recorded in audit logs.

Table 13-1 GaussDB(DWS) operations that can be recorded by CTS

Operation	Resource	Event Name
Creating/Restoring a cluster	cluster	createCluster
Deleting a cluster	cluster	deleteCluster
Scaling out a cluster	cluster	resizeCluster
Restarting a cluster	cluster	restartCluster
Creating a snapshot	backup	createBackup
Deleting a snapshot	backup	deleteBackup
Setting security parameters	configurations	updateConfigurations

Operation	Resource	Event Name
Creating an MRS data source	dataSource	createExtDataSource
Deleting an MRS data source	dataSource	deleteExtDataSource
Updating an MRS data source	dataSource	updateExtDataSource

Viewing Traces

Step 1 Log in to the management console, choose **Service List > Management & Governance > Cloud Trace Service**. The CTS management console is displayed.

Step 2 In the navigation pane on the left, choose **Trace List**.

Step 3 Click the search box above the trace list and set the search criteria.

The following filters are available:

- **Trace Name:** If you select this option, you also need to select a specific trace name.
- **Cloud Service:** Select **GaussDB(DWS)**.
- **Resource Type:** Select **All resource types** or specify a resource type.
- **Resource Name:** If you select this option, you also need to select or enter a specific resource name.
- **Resource ID:** If you select this option, you also need to select or enter a specific resource ID.
- **Operator:** Select a specific operator (at user level rather than tenant level).
- **Event ID:** If you select this option, you also need to select or enter an event ID.
- **Trace Status:** Available options include **All trace statuses**, **normal**, **warning**, and **incident**. You can only select one of them.
- **Enterprise Project ID:** If you select this option, you also need to select or enter a specific enterprise project ID.
- **Access Key ID:** If you select this option, you also need to select or enter a specific access key ID.

Step 4 Click **Query**.

Step 5 Click the name of the trace to be viewed. A window is displayed, showing the trace details.

For details about key fields of a CTS trace, see "Trace References" > "Trace Structure" and "Trace References" > "Example Traces" in [Cloud Trace Service User Guide](#).

----End

13.3 Database Audit Logs

13.3.1 Configuring the Database Audit Logs

Prerequisites

Database audit logs are configured on the **Security Settings** page. You can change security settings only when the cluster status is **Available** and **Unbalanced**, and **Task Information** cannot be **Creating snapshot**, **Scaling out**, **Configuring**, or **Restarting**.

Procedure

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters** > **Dedicated Clusters**.

Step 3 In the cluster list, click the name of a cluster. Choose **Security**.

By default, **Configuration Status** is **Synchronized**, which indicates that the latest database results are displayed.

Step 4 In the **Audit Settings** area, set the audit items:

 **NOTE**

The default audit log retention policy is space-first, which means audit logs will be automatically deleted when the size of audit logs on a single node exceeds 1 GB. This function prevents node faults or low performance caused by high disk space occupied by audit logs.

[Table 13-2](#) describes the detailed information about the audit items.

Table 13-2 Audit items

Audit Item	Description
Unauthorized access	Specifies whether to record unauthorized operations. This parameter is disabled by default.
DQL operations	SELECT operations can be selected. NOTE 8.1.1.100 and later versions support the DQL operations audit item.
DML operations	Specifies whether to record INSERT , UPDATE , and DELETE operations on tables. This parameter is disabled by default. NOTE 8.1.1.100 and later versions support fine-grained splitting of audit items, and the COPY and MERGE options are added.

Audit Item	Description
DDL operations	<p>Specifies whether to record the CREATE, DROP, and ALTER operations of specified database objects. DATABASE, SCHEMA, and USER are selected by default.</p> <p>NOTE 8.1.1.100 and later versions support TABLE, DATA SOURCE, and NODE GROUP operations. These operations are enabled by default.</p>
Other operations	<p>Specifies whether to record other operations. Only the TRANSACTION and CURSOR operations are selected by default.</p> <p>NOTE</p> <ul style="list-style-type: none"> • 8.1.1.100 and later versions support the Other operations audit item. • You are advised to select TRANSACTION. Otherwise, statements in a transaction will not be audited. • You are advised to select CURSOR. Otherwise, SELECT statements in a cursor will not be audited. The Data Studio client automatically encapsulates SELECT statements using CURSOR.

Except the audit items listed in [Table 13-2](#), key audit items in [Table 13-3](#) are enabled by default on GaussDB(DWS).

Table 13-3 Key audit items

Parameter	Description
Key audit items	Records successful and failed logins and logout.
	Records database startup, stop, recovery, and switchover.
	Records user locking and unlocking.
	Records the grants and reclaims of user permissions.
	Records the audit function of the SET operation.

Step 5 Enable or disable audit log dumps.

For more information, see [Enabling Audit Log Dumps](#).

Step 6 Click **Apply**.

Click . The configuration status **Applying** indicates that the configurations are being saved.

When the status changes to **Synchronized**, the configurations are saved and take effect.

----End

13.3.2 Dumping the Database Audit Logs

GaussDB(DWS) records information (audit logs) about connections and user activities in your database. The audit logs help you monitor the database to ensure security, rectify faults, and locate historical operation records.

GaussDB(DWS) audit logs are stored in the database by default. You can dump the audit logs to OBS so that users who monitor database activities can view the logs.

NOTE

- This function cannot be used if OBS is not available.
- Data may during cluster specifications change, CN addition, or CN deletion. You are advised to disable audit log dump during these operations.
- If a CN node is faulty, data on the CN node may be lost.
- After audit log dumping is enabled, audit logs will be dumped if the size of saved audit logs exceeds 1 GB. This may cause abnormal query results. Exercise caution when performing this operation.
- Version support for the audit log dump directory partition is as follows:
 - For 8.1.3.x clusters, only 8.1.3.322 and later versions support this feature. For 8.2.0.x clusters, only 8.2.0.106 and later versions support this feature. By default, the audit log dump directory partition is enabled and cannot be disabled.
 - To use this feature in earlier versions, contact technical support to upgrade your cluster first. Manually enable this feature after the upgrade.

Prerequisites

After a GaussDB(DWS) cluster is created, you can enable log dump for it to dump audit logs to OBS. **Before enabling audit log dump, ensure the following conditions are met:**

You have created an OBS bucket for storing the audit logs. For details, see "Managing Buckets > Creating a Bucket" in the *Object Storage Service Console Operation Guide*.

Enabling Audit Log Dumps

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.

Step 3 In the cluster list, click the name of the cluster for which you want to enable audit log dump. In the navigation pane, choose **Security Settings**.

Step 4 In the **Audit Settings** area, enable **Audit Log Dump**.

When you enable audit log dump for a project in a region for the first time, the system prompts you to create an agency named **DWSAccessOBS**. After the agency is created, GaussDB(DWS) can dump audit logs to OBS.

By default, only Huawei Cloud accounts or users with Security Administrator permissions can query and create agencies. IAM users under an account do not

have the permission to query or create agencies by default. Contact a user with that permission and complete the authorization on the current page.

- **OBS Foreign Table:** Audit logs can be read using OBS foreign tables during dumping. Audit logs are stored in CSV format and compressed in GZ format.
- **OBS Bucket:** Name of the OBS bucket used to store the audit data. If no OBS bucket is available, click **View OBS Bucket** to access the OBS console and create one. For details, "Managing Buckets" > "Creating a Bucket" in *Object Storage Service Console Operation Guide*.
- **OBS Path:** User-defined directory on OBS for storing audit files. Different directory levels are separated by forward slashes (/). The value is a string containing 1 to 50 characters, which cannot start with a forward slash (/). If the entered OBS path does not exist, the system creates one and dumps data to it.
- **Dump Interval (Minute):** Interval based on which GaussDB(DWS) periodically dumps data to OBS. The value range is 5 to 43200. The unit is minute.

Step 5 Click **Apply**.

If **Configuration Status** is **Applying**, the system is saving the settings.

Wait for a moment and then refresh **Configuration Status**. When **Configuration Status** is **Synchronized**, the configuration is saved and takes effect.

----End

Modifying Audit Log Dump Configurations

After audit log dump is enabled, you can modify the dump configuration. For example, you can modify the OBS bucket and path for storing logs and the dump period.

The procedure is as follows:

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.

Step 3 In the cluster list, click the name of the cluster for which you want to modify the audit log dump configurations. In the navigation pane, choose **Security**.

Step 4 In the **Audit Settings** area, modify the **Audit Log Dump** configurations.

Step 5 Click **Apply**.

If **Configuration Status** is **Applying**, the system is saving the settings.

Wait for a moment and then refresh **Configuration Status**. When **Configuration Status** is **Synchronized**, the configuration is saved and takes effect.

----End

Viewing Audit Log Dumps

After audit log dump is enabled, you can view the dumped audit logs on OBS.

The procedure is as follows:

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 3** In the cluster list, click the name of the target cluster for which you want to view the log dump history. In the navigation pane, choose **Security**.
- Step 4** In the **Audit Settings** area, click **View Dump Record**.
- Step 5** In the **Audit Log Dump Records** dialog box, click **View OBS Bucket**. The OBS console page is displayed.
- Step 6** Select the OBS bucket and folder where the logs are stored to view the log files.

You can download and decompress the files to view. The fields of audit log files are described as follows:

Table 13-4 Log file fields

Field	Type	Description
begintime	timestamp with time zone	Operation start time
endtime	timestamp with time zone	Operation end time
operation_type	text	Operation type. For details, see Table 13-5 .
audit_type	text	Audit type. For details, see Table 13-6 .
result	text	Operation result
username	text	Name of the user who performs the operation
database	text	Database name
client_conninfo	text	Client connection information, that is, gsql, JDBC, or ODBC.
object_name	text	Object name
command_text	text	Command used to perform the operation
detail_info	text	Operation details
transaction_xid	text	Transaction ID
query_id	text	Query ID
node_name	text	Node name
thread_id	text	Thread ID
local_port	text	Local port
remote_port	text	Remote port

Table 13-5 Operation types

Operation Type	Description
audit_switch	Indicates that the operations of enabling and disabling the audit log function are audited.
login_logout	Indicates that user login and log-out operations are audited.
system	Indicates that the system startup, shutdown, and instance switchover operations are audited.
sql_parse	Indicates that SQL statement parsing operations are audited.
user_lock	Indicates that user locking and unlocking operations are audited.
grant_revoke	Indicates that user permission granting and revoking operations are audited.
violation	Indicates that user's access violation operations are audited.
ddl	Indicates that DDL operations are audited. DDL operations are controlled at a fine granularity based on operation objects. Therefore, audit_system_object is used to control the objects whose DDL operations are to be audited. (The audit function takes effect as long as audit_system_object is configured, no matter whether ddl is set.)
dml	Indicates that the DML operations are audited.
select	Indicates that the SELECT operations are audited.
internal_event	Indicates that internal incident operations are audited.
user_func	Indicates that operations related to user-defined functions, stored procedures, and anonymous blocks are audited.
special_func	Indicates that special function invoking operations are audited. Special functions include pg_terminate_backend and pg_cancel_backend .
copy	Indicates that the COPY operations are audited.
set	Indicates that the SET operations are audited.
transaction	Indicates that transaction operations are audited.
vacuum	Indicates that the VACUUM operations are audited.
analyze	Indicates that the ANALYZE operations are audited.

Operation Type	Description
cursor	Indicates that cursor operations are audited.
anonymous_block	Indicates that the anonymous block operations are audited.
explain	Indicates that the EXPLAIN operations are audited.
show	Indicates that the SHOW operations are audited.
lock_table	Indicates that table lock operations are audited.
comment	Indicates that the COMMENT operations are audited.
preparestmt	Indicates that the PREPARE , EXECUTE , and DEALLOCATE operations are audited.
cluster	Indicates that the CLUSTER operations are audited.
constraints	Indicates that the CONSTRAINTS operations are audited.
checkpoint	Indicates that the CHECKPOINT operations are audited.
barrier	Indicates that the BARRIER operations are audited.
cleanconn	Indicates that the CLEAN CONNECTION operations are audited.
seclabel	Indicates that security label operations are audited.
notify	Indicates that the notification operations are audited.
load	Indicates that the loading operations are audited.

Table 13-6 audit_type parameters

Parameter	Description
audit_open/audit_close	Indicates that the audit type is operations enabling or disabling audit logs.
user_login/user_logout	Indicates that the audit type is operations and users with successful login/logout.
system_start/system_stop/ system_recover/system_switch	Indicates that the audit type is system startup, shutdown, and instance switchover.
sql_wait/sql_parse	Indicates that the audit type is SQL statement parsing.
lock_user/unlock_user	Indicates that the audit type is successful user locking and unlocking.

Parameter	Description
grant_role/revoke__role	Indicates that the audit type is user permission granting and revoking.
user_violation	Indicates that the audit type is unauthorized user access operations.
ddl_database_object	Indicates that successful DDL operations are audited. DDL operations are controlled at a fine granularity based on operation objects. So, audit_system_object is used to control the objects whose DDL operations are to be audited. (The audit function takes effect as long as audit_system_object is configured, no matter whether ddl is set.) For example, ddl_sequence indicates that the audit type is sequence-related operations.
dml_action_insert/ dml_action_delete/ dml_action_update/ dml_action_merge/ dml_action_select	Indicates that the audit type is DML operations such as INSERT , DELETE , UPDATE , and MERGE .
internal_event	Indicates that the audit type is internal events.
user_func	Indicates that the audit type is user-defined functions, stored procedures, or anonymous block operations.
special_func	Indicates that the audit type is special function invocation. Special functions include pg_terminate_backend and pg_cancel_backend .
copy_to/copy_from	Indicates that the audit type is COPY operations.
set_parameter	Indicates that the audit type is SET operations.
trans_begin/trans_commit/ trans_prepare/ trans_rollback_to/ trans_release/trans_savepoint/ trans_commit_prepare/ trans_rollback_prepare/ trans_rollback	Indicates that the audit type is transaction-related operations.
vacuum/vacuum_full/ vacuum_merge	Indicates that the audit type is VACUUM operations.
analyze/analyze_verify	Indicates that the audit type is ANALYZE operations.

Parameter	Description
cursor_declare/cursor_move/ cursor_fetch/cursor_close	Indicates that the audit type is cursor-related operations.
codeblock_execute	Indicates that the audit type is anonymous blocks.
explain	Indicates that the audit type is EXPLAIN operations.
show	Indicates that the audit type is SHOW operations.
lock_table	Indicates that the audit type is table locking operations.
comment	Indicates that the audit type is COMMENT operations.
prepare/execute/deallocate	Indicates that the audit type is PREPARE , EXECUTE , or DEALLOCATE operations.
cluster	Indicates that the audit type is CLUSTER operations.
constraints	Indicates that the audit type is CONSTRAINTS operations.
checkpoint	Indicates that the audit type is CHECKPOINT operations.
barrier	Indicates that the audit type is BARRIER operations.
cleanconn	Indicates that the audit type is CLEAN CONNECTION operations.
seclabel	Indicates that the audit type is security label operations.
notify	Indicates that the audit type is notification operations.
load	Indicates that the audit type is loading operations.

----End

Disabling Audit Log Dumps

You can disable the audit log dump function if you do not want to dump audit logs to OBS.

The procedure is as follows:

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 3** In the cluster list, click the name of the cluster for which you want to disable audit log dump. In the navigation pane, choose **Security Settings**.
- Step 4** In the **Audit Settings** area, disable audit log dump.
- Step 5** Click **Apply**.

If **Configuration Status** is **Applying**, the system is saving the settings.

Wait for a moment and then refresh **Configuration Status**. When **Configuration Status** is **Synchronized**, the configuration is saved and takes effect.

----End

13.3.3 Viewing Database Audit Logs

Prerequisites

- The audit function has been enabled by setting **audit_enabled**. The default value of **audit_enabled** is **ON**. To disable audit, set it to **OFF** by referring to [Modifying Database Parameters](#).
- The audit items have been configured. For details about how to enable audit items, see [Configuring the Database Audit Logs](#).
- The database is running properly and a series of addition, modification, deletion, and query operations have been executed in the database. Otherwise, no audit result is generated.
- The audit logs of each database node are recorded separately.
- Only users with the **AUDITADMIN** permission can view audit records.

Viewing Database Audit Logs

Method 1: Audit logs will occupy disk space. To prevent excessive disk usage, GaussDB(DWS) supports audit log dumping. You can enable the **Log Dump** function to dump audit logs to OBS (you need to create an OBS bucket for storing audit logs first). For details about how to view the dumped logs, see [Viewing Audit Log Dumps](#).

Method 2: Use the **Log** function of LTS to view or download the collected database audit logs. For details, see [Checking Cluster Logs](#).

Method 3: Database audit logs are stored in the database by default. After connecting to the cluster, you can use the **pg_query_audit** function to view the logs. For details, see [Using Functions to View Database Audit Logs](#).

Using Functions to View Database Audit Logs

- Step 1** Use the SQL client tool to connect to the database cluster. For details, see [Cluster Connection](#).
- Step 2** Use the **pg_query_audit** function to query the audit logs of the current CN. The syntax is as follows:

```
pg_query_audit(timestampz starttime,timestampz endtime,audit_log)
```

starttime and **endtime** indicate the start time and end time of the audit record, respectively. **audit_log** indicates the physical file path of the queried audit logs. If **audit_log** is not specified, the audit log information of the current instance is queried.

For example, view the audit records of the current CN node in a specified period.
SELECT * FROM pg_query_audit('2021-02-23 21:49:00','2021-02-23 21:50:00');

The query result is as follows:

begintime	endtime	operation_type	audit_type	result	username	database
client_conninfo	object_name	command_text	detail_info			
transaction_xid	query_id	node_name	session_id	local_port	remote_port	
2021-02-23 21:49:57.76+08	2021-02-23 21:49:57.82+08	login_logout	user_login	ok	dbadmin	gaussdb
gsq@[local]	gaussdb	login db	login db(gaussdb) successfully, the current user is: dbadmin 0			
	0	coordinator1	140324035360512.667403397820909.coordinator1	27777		

This record indicates that user **dbadmin** logged in to the **gaussdb** database at 2021-02-23 21:49:57.82 (GMT+08:00). After the host specified by **log_hostname** is started and a client is connected to its IP address, the host name found by reverse DNS resolution is displayed following the at sign (@) in the value of **client_conninfo**.

Step 3 Use the **pgxc_query_audit** function to query audit logs of all CNs. The syntax is as follows:

```
pgxc_query_audit(timestampz starttime,timestampz endtime)
```

For example, view the audit records of all CN nodes in a specified period.
SELECT * FROM pgxc_query_audit('2021-02-23 22:05:00','2021-02-23 22:07:00') where audit_type = 'user_login' and username = 'user1';

The query result is as follows:

begintime	endtime	operation_type	audit_type	result	username	database
client_conninfo	object_name	command_text	detail_info			
query_id	node_name	session_id	local_port	remote_port	transaction_xid	
2021-02-23 22:06:22.219+08	2021-02-23 22:06:22.271+08	login_logout	user_login	ok	user1	gaussdb
gsq@[local]	gaussdb	login db	login db(gaussdb) successfully, the current user is: user1 0			
	0	coordinator2	140689577342720.667404382271356.coordinator	27782		
2021-02-23 22:05:51.697+08	2021-02-23 22:05:51.749+08	login_logout	user_login	ok	user1	gaussdb
gsq@[local]	gaussdb	login db	login db(gaussdb) successfully, the current user is: user1 0			
	0	coordinator1	140525048424192.667404351749143.coordinator1	27777		

The query result shows the successful login records of **user1** in to CN1 and CN2.

Step 4 Query the audit records of multiple objects.

```
SET audit_object_name_format TO 'all';
SELECT object_name,result,operation_type,command_text FROM pgxc_query_audit('2022-08-26 8:00:00','2022-08-26 22:55:00') where command_text like '%student%';
```

The query result is as follows:

object_name	result	operation_type
command_text		

```
-----  
student | ok | ddl | CREATE TABLE student(stuNo int,  
stuName TEXT);  
studentscore | ok | ddl | CREATE TABLE studentscore(stuNo  
int, stuscore int);  
["public.student_view01","public.studentscore","public.student"] | ok | ddl | CREATE OR REPLACE  
VIEW student_view01 AS SELECT * FROM student t1 where t1.stuNo in (select stuNo from studentscore t2  
where t1.stuNo = t2.stuNo);  
["public.student_view01","public.student","public.studentscore"] | ok | dml | SELECT * FROM  
student_view01;
```

In the **object_name** column, the table, view, and base table associated with the view are displayed.

----End

14 Cluster Security Management

14.1 Configuring Separation of Duties for the GaussDB(DWS) Cluster

Scenario

By default, the administrator specified when you create a GaussDB(DWS) cluster is the database system administrator. The administrator can create other users and view the audit logs of the database. That is, separation of permissions is disabled.

GaussDB(DWS) supports role-based separation of permissions. In this way, different roles have different permissions and cluster data can be better protected.

For details about the default permissions mode and the separation of permissions mode, see [Separation of Permissions](#) in the *Data Warehouse Service (DWS) Developer Guide*.

Impact on the System

- After you modified the security parameters and the modifications take effect, the cluster may be restarted, which makes the cluster unavailable temporarily.
- When a GaussDB(DWS) 3.0 cluster is created, a logical cluster is created by default. After the separation of duties is enabled, only the system administrator has the permission to create, modify, delete, and allocate logical clusters. Accessing a logical cluster requires permissions.

Prerequisites

To modify the cluster's security configuration, ensure that the following conditions are met:

- The cluster status is **Available**, **To be restarted**, or **Unbalanced**.
- The **Task Information** cannot be **Creating snapshot**, **Scaling out**, **Configuring**, or **Restarting**.

Procedure

Step 1 Log in to the GaussDB(DWS) console.

Step 2 In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.

Step 3 In the cluster list, click the name of a cluster. On the page that is displayed, click **Security Settings**.

By default, **Configuration Status** is **Synchronized**, which indicates that the latest database result is displayed.

Step 4 On the **Security Settings** page, configure separation of permissions.

When separation of permissions is enabled, configure the username and password for **Security Administrator** and **Audit Administrator**. Then the system automatically creates these two users. You can use these two users to connect to the database and perform database-related operations. **By default, this function is disabled.**

Table 14-1 Security parameters

Parameter	Description	Example Value
Security Administrator	The username must meet the following requirements: <ul style="list-style-type: none">• Consists of lowercase letters, digits, or underscores.• Starts with a lowercase letter or an underscore.• Contains 6 to 64 characters.	security_admin
Password	The password complexity requirements are as follows: <ul style="list-style-type: none">• Contain 12 to 32 characters.• Cannot be the username or the username spelled backwards.• Contains at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters (~!?,.,:;_){}[]/<>@#%^&*+ \=-)• Perform the weak password check on the passwords that you have created.	-
Confirm Password	Enter the password of the security administrator again.	-

Parameter	Description	Example Value
Audit Administrator	The username must meet the following requirements: <ul style="list-style-type: none">• Consists of lowercase letters, digits, or underscores.• Starts with a lowercase letter or an underscore.• Contains 6 to 64 characters.	audit_admin
Password	The password complexity requirements are as follows: <ul style="list-style-type: none">• Contain 12 to 32 characters.• Cannot be the username or the username spelled backwards.• Must contain at least 3 of the following character types: uppercase letters, lowercase letters, digits, and special characters ~!@#%^&*()-_+= []{};:;<.>/?• Passes the weak password check.	-
Confirm Password	Enter the password of the audit administrator again.	-

Step 5 Click **Apply**.

Step 6 In the displayed **Save Configuration** dialog box, select or deselect **Restart the cluster** and click **Yes**.

- If you select **Restart the cluster**, the system saves the settings on the **Security Settings** page and restarts the cluster immediately. After the cluster is restarted, the security settings take effect immediately.
- If you do not select **Restart the cluster**, the system only saves the settings on the **Security Settings** page. Later, you need to manually restart the cluster for the security settings to take effect.

After the security settings are complete, **Configuration Status** can be one of the following on the **Security Settings** page:

- **Applying**: The system is saving the settings.
- **Synchronized**: The settings have been saved and taken effect.
- **Take effect after restart**: The settings have been saved but have not taken effect. Restart the cluster for the settings to take effect.

----End

14.2 Encrypting Databases

14.2.1 Overview

Encrypting GaussDB(DWS) Databases

In GaussDB(DWS), you can enable database encryption for a cluster to protect static data. After you enable encryption, data of the cluster and its snapshots is encrypted. Encryption is an optional and immutable setting that can be configured during cluster creation. To encrypt an unencrypted cluster, or vice versa, you must export all data from the unencrypted cluster and import it into a new cluster that has database encryption enabled. Database encryption is performed when data is written to GaussDB(DWS). That is, GaussDB(DWS) encrypts data when the data is written to GaussDB(DWS). If you want to query the data, GaussDB(DWS) automatically decrypts it and returns the result to you.

If encryption is required, enable it during cluster creation. Although encryption is an optional setting of GaussDB(DWS), you are advised to enable this setting for clusters to protect data.

NOTICE

- The DWS 3.0 cluster does not support database encryption.
- The database encryption function can be enabled or disabled only when a cluster is created. It cannot be enabled after a cluster is created. Once enabled, it cannot be disabled. For details, see [Encrypting the Database](#).
- After **Encrypt DataStore** is enabled, the key cannot be disabled, deleted, or frozen when being used. Otherwise, the cluster becomes abnormal and the database becomes unavailable.
- Snapshots created after the database encryption function is enabled cannot be restored using open APIs.

Viewing Database Encryption Information

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation pane on the left, choose **Clusters > Dedicated Clusters**.
- Step 3** Click the name of a cluster. The **Cluster Information** page is displayed.
- Step 4** In the **Data Encryption Information** area on the cluster information page, view the database encryption information, as shown in [Table 14-2](#).

Table 14-2 Data encryption information

Parameter	Description
Key Name	Indicates the database encryption key of the cluster when Encrypt DataStore is enabled.
Last Key Rotation Time	Indicates the time when the last encryption key is rotated when Encrypt DataStore is enabled.

 **NOTE**

If database encryption is disabled by default during cluster creation, the encryption module is not displayed on the cluster details page.

----End

Encrypting GaussDB(DWS) Databases Using KMS

When you choose KMS to manage GaussDB(DWS) keys, a three-layer key management structure is adopted, including the cluster master key (CMK), cluster encryption key (CEK), and database encryption key (DEK).

- The CMK is used to encrypt the CEK and is stored in KMS.
- The CEK is used to encrypt the DEK. The CEK plaintext is stored in the data warehouse cluster's memory, and the ciphertext is stored in GaussDB(DWS).
- The DEK is used to encrypt database data. The DEK plaintext is stored in the data warehouse cluster's memory, and the ciphertext is stored in GaussDB(DWS).

The procedure of using the keys is as follows:

1. You choose a CMK.
2. GaussDB(DWS) randomly generates the CEK and DEK plaintext.
3. KMS uses the CMK you choose to encrypt the CEK plaintext and imports the encrypted CEK ciphertext to GaussDB(DWS).
4. GaussDB(DWS) uses the CEK plaintext to encrypt the DEK plaintext and saves the encrypted DEK ciphertext.
5. GaussDB(DWS) transfers the DEK plaintext to the cluster and loads it to the cluster's memory.

When the cluster is restarted, it automatically requests the DEK plaintext from GaussDB(DWS) through an API. GaussDB(DWS) loads the CEK and DEK ciphertext to the cluster's memory, invokes KMS to decrypt the CEK using the CMK, loads the CEK to the memory, decrypts the DEK using the CEK plaintext, loads the DEK to the memory, and returns it to the cluster.

Rotating Encryption Keys

Encryption key rotation is used to update the ciphertext stored on GaussDB(DWS). On GaussDB(DWS), you can rotate the encrypted CEK of an encrypted cluster.

The procedure of rotating the keys is as follows:

1. The GaussDB(DWS) cluster starts key rotation.
2. GaussDB(DWS) decrypts the CEK ciphertext stored on GaussDB(DWS) based on the CMK to obtain the CEK plaintext.
3. Use the obtained CEK plaintext to decrypt the DEK ciphertext in GaussDB(DWS) to obtain the DEK plaintext.
4. GaussDB(DWS) randomly generates new CEK plaintext.

5. GaussDB(DWS) uses the new CEK plaintext to encrypt the DEK and saves the encrypted DEK ciphertext.
6. Use the CMK to encrypt the new CEK plaintext and import the encrypted CEK ciphertext to GaussDB(DWS).

You can plan the key rotation interval based on service requirements and data types. To improve data security, you are advised to periodically rotate the keys to prevent the keys from being cracked. Once you find that your keys may have been disclosed, rotate the keys in time.

 **NOTE**

- When GaussDB(DWS) rotates the cluster's CEK, snapshots of the cluster do not need CEK rotation, because the CEK is not stored in snapshots. The CEK plaintext is stored in the GaussDB(DWS) cluster memory, and the ciphertext is stored in GaussDB(DWS).
- The DEK is not updated during key rotation, so data encryption and decryption are not affected.

14.2.2 Rotating Encryption Keys

If you have enabled the **Encrypt DataStore** function in **Advanced Settings** during cluster creation, you can rotate the encryption keys for the cluster after the cluster is created successfully. Each key rotation will update the CEK once. During the key rotation, the cluster is still in **Available** status.

Rotating Encryption Keys for Data Warehouse Clusters

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** In the navigation tree on the left, choose **Clusters > Dedicated Clusters**.
- Step 3** In the cluster list, find the target cluster and click the cluster name. The **Cluster Information** page is displayed.
- Step 4** In the **Data Encryption Information** area, click **Key Rotation**.
- Step 5** In the dialog box that is displayed, click **Yes**.

----End

14.2.3 Converting an Ordinary Cluster to an Encrypted Cluster

GaussDB(DWS) allows you to convert an unencrypted cluster to an encrypted cluster when the cluster status is **Available** on the console. To ensure data security, converting a cluster to an encrypted cluster is an **irreversible high-risk operation** and will restart the cluster. As a result, services may be unavailable for a short period of time. Exercise caution when performing this operation.

 **NOTE**

If the current console does not support this feature, contact technical support.

Procedure

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane, choose **Clusters > Dedicated Clusters**.

Step 2 In the cluster list, locate the row that contains the target cluster and choose **More > Convert to Encrypted Cluster** in the **Operation** column.

Step 3 In the dialog box that is displayed, select the key source and encryption algorithm to convert the cluster into an encrypted cluster.

- Method 1: Select a key name.
- Method 2: Enter the key ID. Enter the key ID used for authorizing the current tenant..

When you grant permissions on the [Creating a Grant](#) page, the authorized object must be an account instead of a user. The authorized operations must at least contain **Querying key details**, **Encrypting data**, and **Decrypting data**.

Table 14-3 Parameter description

Parameter	Description
Key Source	You can select a key name from the key list or directly enter a key name.
Cryptographic Algorithm	The encryption algorithms are as follows: <ul style="list-style-type: none">• AES256 (general encryption algorithm, SM algorithms not supported)• SM4 (compatible with international algorithms)

 **NOTE**

- The database encryption function cannot be disabled once it is enabled.
- After **Encrypt DataStore** is enabled, the key cannot be disabled, deleted, or frozen when being used. Otherwise, the cluster becomes abnormal and the database becomes unavailable.
- Snapshots created after the database encryption function is enabled cannot be restored using open APIs.
- By default, only Huawei Cloud accounts or users with Security Administrator permissions can query and create agencies. IAM users under an account do not have the permission to query or create agencies by default. Contact a user with that permission and complete the authorization on the current page.

Step 4 After the conversion, you can click the cluster name to go to the **Cluster Details** page to view the cluster details. For details, see [Viewing Database Encryption Information](#).

----End

14.3 Permissions

14.3.1 Creating a User and Granting GaussDB(DWS) Permissions

This topic describes how to use [IAM](#) to implement fine-grained permissions control for your GaussDB(DWS) resources. With IAM, you can:

- Create IAM users for employees based on the organizational structure of your enterprise. Each IAM user has their own security credentials, providing access to GaussDB(DWS) resources.
- Grant only the permissions required for users to perform specific tasks.
- Entrust a Huawei Cloud account or service to perform professional and efficient O&M on your GaussDB(DWS) resources.

If your Huawei Cloud account does not need individual IAM users, you may skip this section.

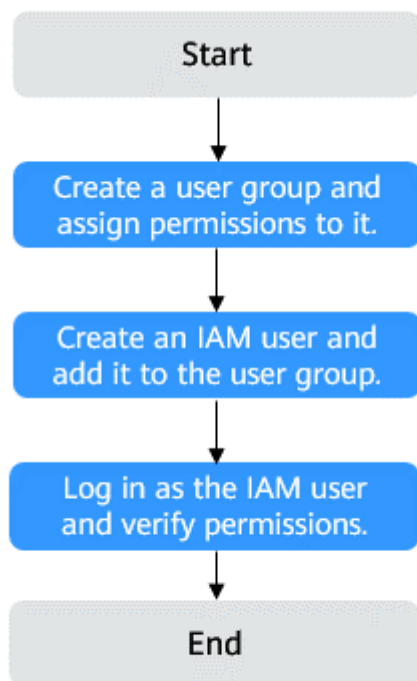
This section describes the procedure for granting permissions (see [Procedure](#)).

Prerequisites

Before assigning permission policies to a user group, you need to understand the GaussDB(DWS) permission policies. For details about the system policies supported by GaussDB(DWS), see [Supported System Policies](#). For the system policies of other services, see [System Permissions](#).

Procedure

Figure 14-1 Procedure



1. **Create a user group and assign permissions.**
Use the Huawei Cloud account to log in to the **IAM console**, create a user group, and attach the **DWS ReadOnlyAccess** policy to the group.
2. **Create a user and add it to a user group .**
Create a user on the IAM console and add the user to the group created in Step 1.
3. **log in** and verify the permissions.
Log in to the management console by using the user created and verify the user permissions.
 - Choose **Service List > Data Warehouse Service** to access the GaussDB(DWS) console, and click **Create GaussDB(DWS) Cluster** to create a GaussDB(DWS) cluster. If you cannot create one, **DWS ReadOnlyAccess** has taken effect.
 - Choose any other service in **Service List**. If only the **DWS ReadOnlyAccess** policy is added and a message is displayed indicating that you have insufficient permission to access the service, **DWS ReadOnlyAccess** has taken effect.

14.3.2 Creating a GaussDB(DWS) Custom Policy

Custom policies can be created as a supplement to the system policies of GaussDB(DWS). For details about the custom policy actions, see [Permissions Policies and Supported Actions](#).

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

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

For details, [Creating a Custom Policy](#). This section provides examples of GaussDB(DWS) custom policies.

Custom Policy Examples

- Example 1: allowing users to create/restore, restart, and delete a cluster, configure security parameters, and reset passwords

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "dws:cluster:create",
        "dws:cluster:restart",
        "dws:cluster:delete",
        "dws:cluster:setSecuritySettings",
        "dws:cluster:resetPassword",
        "dws:*:list*",
        "dws:*:get*",
        "tms:predefineTags:list",
        "ecs:*:get*",
        "ecs:*:list*",
        "elb:*:list*",
        "ecs:*:create*",
        "ecs:*:delete*",
        "vpc:*:get*",

```

```
        "vpc:*:list*",
        "vpc:*:create*",
        "vpc:*:delete*",
        "evs:*:get*",
        "evs:*:list*",
        "evs:*:create*",
        "evs:*:delete*"
    ]
}
]
```

- Example 2: using wildcard character (*)

For example, the following policy has all operation permissions on GaussDB(DWS) snapshots.

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "dws:snapshot:*",
        "dws:cluster:list",
        "dws:openAPISnapshot:detail",
        "dws:cluster:getDetail",
        "ecs:*:get*",
        "ecs:*:list*",
        "vpc:*:get*",
        "vpc:*:list*"
      ]
    }
  ]
}
```

- Example 3: denying cluster deletion

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

The following method can be used if you need to assign permissions of the **GaussDB(DWS) FullAccess** policy to a user but also forbid the user from deleting clusters. Create a custom policy for denying cluster deletion, and assign both policies to the group the user belongs to. Then the user can perform all operations on GaussDB(DWS) except deleting clusters. The following is an example of a deny policy:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "dws:*:list*",
        "dws:*:get*"
      ]
    },
    {
      "Effect": "Deny",
      "Action": [
        "dws:cluster:delete"
      ]
    }
  ]
}
```

- Example 4: defining multiple actions in a policy

A custom policy can contain actions of multiple services that are all of the global or project-level type. The following is an example policy containing actions of multiple services:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "dws:cluster:create",
        "dws:cluster:restart",
        "dws:cluster:setSecuritySettings",
        "dws:*:get*",
        "dws:*:list*",
        "tms:predefineTags:list",
        "elb:*:list*",
        "ecs:*:get*",
        "ecs:*:list*",
        "ecs:*:create*",
        "vpc:*:get*",
        "vpc:*:list*",
        "vpc:*:create*",
        "evs:*:get*",
        "evs:*:list*",
        "evs:*:create*"
      ]
    },
    {
      "Effect": "Deny",
      "Action": [
        "dws:cluster:delete"
      ]
    }
  ]
}
```

14.3.3 Syntax of Fine-Grained Permissions Policies

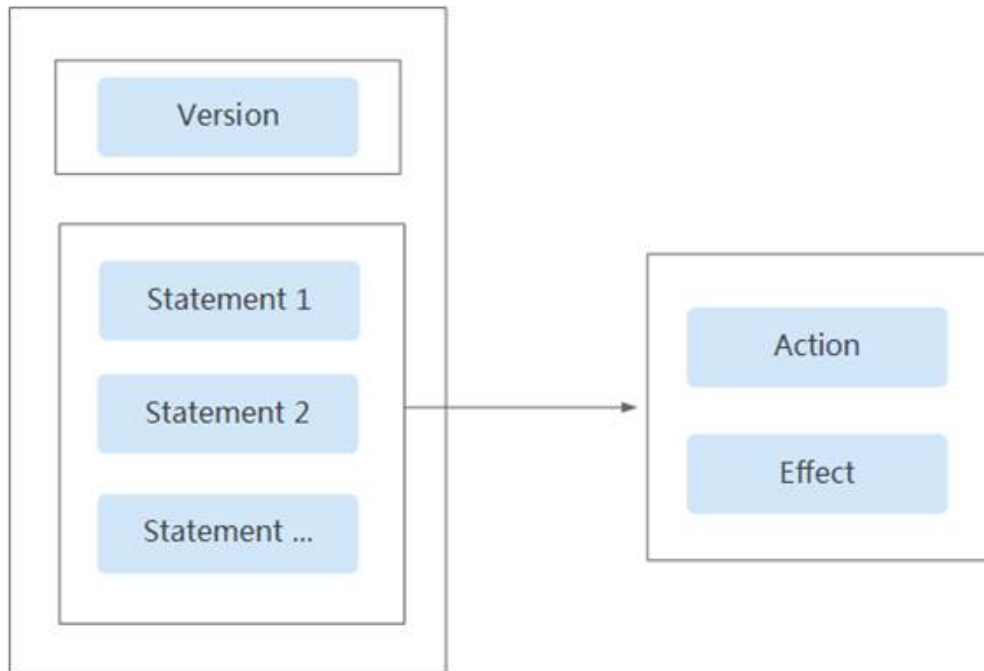
In actual services, you may need to grant different operation permissions on resources to users of different roles. The IAM service provides fine-grained access control. An IAM administrator (a user in the **admin** group) can create a custom policy containing required permissions. After a policy is granted to a user group, users in the group can obtain all permissions defined by the policy. In this way, IAM implements fine-grained permission management.

To control the GaussDB(DWS) operations on resources more precisely, you can use the user management function of IAM to grant different operation permissions to users of different roles for fine-grained permission control.

Policy Structure

A fine-grained policy consists of a Version and a Statement. Each policy can have multiple statements.

Figure 14-2 Policy structure



Policy Syntax

In the navigation pane on the IAM console, click **Policies** and then click the name of a policy to view its details. The **DWS ReadOnlyAccess** policy is used as an example to describe the syntax of fine-grained policies.

Figure 14-3 Setting the policy

Policy Name	DWS Viewer	Scope	Project-level services
Type	System-defined policy	Description	The read-only permissions to all DWS resources.

Policy Content | Visual Editor | **JSON**

```

1 {
2   "Version": "1.1",
3   "Statement": [
4     {
5       "Action": [
6         "dws:*:get*",
7         "dws:*:list*",
8         "ecs:*:get*",
9         "ecs:*:list*",
10        "vpc:*:get*",
11        "vpc:*:list*",
12        "evs:*:get*",
13        "evs:*:list*",
14        "mns:*:get*",
15        "bss:*:list*",
16        "bss:*:get*"
17      ],
18      "Effect": "Allow"
19    }
20  ]
21 }
    
```

```

{
  "Version": "1.1",
    
```

```
"Statement": [
  {
    "Effect": "Allow",
    "Action": [
      "dws:*:get*",
      "dws:*:list*",
      "ecs:*:get*",
      "ecs:*:list*",
      "vpc:*:get*",
      "vpc:*:list*",
      "evs:*:get*",
      "evs:*:list*",
      "mrs:*:get*",
      "bss:*:list*",
      "bss:*:get*"
    ]
  }
]
```

- **Version:** Distinguishes between role-based access control (RBAC) and fine-grained policies.
 - **1.0:** RBAC policies. An RBAC policy consists of permissions for an entire service. Users in a group with such a policy assigned are granted all of the permissions required for that service.
 - **1.1:** Fine-grained policies. A fine-grained policy consists of API-based permissions for operations on specific resource types. Fine-grained policies, as the name suggests, allow for more fine-grained control than RBAC policies. Users granted permissions of such a policy can only perform specific operations on the corresponding service. Fine-grained policies include system and custom policies.
- **Statement:** Permissions defined by a policy, including Effect and Action.
 - **Effect**

The valid values for Effect are Allow and Deny. System policies contain only Allow statements. For custom policies containing both Allow and Deny statements, the Deny statements take precedence.
 - **Action**

Permissions in the format of *Service name:Resource type:Operation*. A policy can contain one or more permissions. The wildcard (*) is allowed to indicate all of the services, resource types, or operations depending on its location in the action.

Example: **dws:cluster:create**, permissions for create data warehouse clusters.

List of Supported Actions

When creating a custom policy on IAM, you can add the operations on GaussDB(DWS) resources or the permissions corresponding to RESTful APIs to the action list of the policy authorization statement so that the policy contains the operation permissions. The following table lists the GaussDB(DWS) permissions.

- **REST API**

For details about RESTful API actions supported by GaussDB(DWS), see [Permissions Policies and Supported Actions](#).
- **Management console operations**

Table 14-4 describes the GaussDB(DWS) operations on resources and corresponding permissions.

 **NOTE**

Some GaussDB(DWS) permissions depend on the actions of ECS, VPC, EVS, ELB, MRS, and OBS. Grant GaussDB(DWS) the required service admin permissions.

Table 14-4 GaussDB(DWS) permissions

Operation	Permission	Dependent Permission	Scope
Creating a cluster	"dws:cluster:create"	"dws::get*", "dws::list*", "ecs::get*", "ecs::list*", "ecs::create*", "vpc::get*", "vpc::list*", "vpc::create*", "vpc:securityGroupRules:delete", "vpc:ports:update", "evs::get*", "evs::list*", "evs::create*",	<ul style="list-style-type: none"> ● Scope: – Project – Enterprise project
Obtaining the cluster list	"dws:cluster:list"	--	<ul style="list-style-type: none"> ● Scope: – Project – Enterprise project
Obtaining the details of a cluster	"dws:cluster:getDetail"	"dws::get*", "dws::list*", "vpc:vpcs:list", "vpc:securityGroups:get"	<ul style="list-style-type: none"> ● Scope: – Project – Enterprise project
Setting automated snapshot policy	"dws:cluster:setAutomatedSnapshot"	"dws:backupPolicy:list"	<ul style="list-style-type: none"> ● Scope: – Project – Enterprise project

Operation	Permission	Dependent Permission	Scope
Setting security parameters/ parameter groups	"dws:cluster:setSecurity Settings"	"dws:*.get*", "dws:*.list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Restarting a Cluster	"dws:cluster:restart"	"dws:*.get*", "dws:*.list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Scaling out clusters	"dws:cluster:scaleOut"	"dws:*.get*", "dws:*.list*", "dws:cluster:scaleOutOrOpenAPIResize", "ecs:*.get*", "ecs:*.list*", "ecs:*.create*", "vpc:*.get*", "vpc:*.list*", "vpc:*.create*", "vpc:*.update*", "evs:*.get*", "evs:*.list*", "evs:*.create*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Scaling out or resizing a cluster via API	"dws:cluster:scaleOutOrOpenAPIResize"	"dws:*.get*", "dws:*.list*", "vpc:vpcs:list", "vpc:ports:create", "vpc:ports:get", "vpc:ports:update", "vpc:subnets:get", "vpc:subnets:update", "vpc:subnets:create", "vpc:routers:get", "vpc:routers:update", "vpc:networks:create", "vpc:networks:get", "vpc:networks:update", "ecs:serverInterfaces:use", "ecs:serverInterfaces:get", "ecs:cloudServerFlavors:get"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Resetting Your Password	"dws:cluster:resetPassword"	"dws:*.get*", "dws:*.list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Deleting a cluster	"dws:cluster:delete"	"dws:*.get*", "dws:*.list*", "ecs:*.get*", "ecs:*.list*", "ecs:*.delete*", "vpc:*.get*", "vpc:*.list*", "vpc:*.delete*", "evs:*.get*", "evs:*.list*", "evs:*.delete*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Configuring maintenance windows	"dws:cluster:setMaintenanceWindow"	"dws:*.get*", "dws:*.list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Binding EIPs	"dws:eip:operate"	"dws:*.get*", "dws:*.list*", "eip:*.get*", "eip:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Unbinding EIPs	"dws:eip:operate"	"dws:*.get*", "dws:*.list*", "eip:*.get*", "eip:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Creating DNS domain names	"dws:dns:create"	"dws:*.get*", "dws:*.list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Releasing DNS domain names	"dws:dns:release"	"dws:*.get*", "dws:*.list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Modifying DNS domain names	"dws:dns:edit"	"dws:*.get*", "dws:*.list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Creating MRS connections	"dws:MRSCONNECTION:create"	"dws:*.get*", "dws:*.list*", "mrs:*.get*", "mrs:*.list*", "mrs:cluster:create" , "ecs:*.get*", "ecs:*.list*", "ecs:*.create*", "vpc:*.get*", "vpc:*.list*", "vpc:*.create*", "evs:*.get*", "evs:*.list*", "evs:*.create*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Updating MRS connections	"dws:MRSCONNECTION:update"	"dws:*.get*", "dws:*.list*", "mrs:*.get*", "mrs:*.list*", "mrs:cluster:create" , "ecs:*.get*", "ecs:*.list*", "ecs:*.create*", "vpc:*.get*", "vpc:*.list*", "vpc:*.create*", "evs:*.get*", "evs:*.list*", "evs:*.create*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Deleting MRS connections	"dws:MRSCONNECTION:delete"	"dws:*:get*", "dws:*:list*", "mrs:*:get*", "mrs:*:list*", "mrs:cluster:create" "ecs:*:get*", "ecs:*:list*", "ecs:*:delete*", "vpc:*:get*", "vpc:*:list*", "vpc:*:delete*", "evs:*:get*", "evs:*:list*", "evs:*:delete*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
MRS data source list	"dws:MRSSOURCE:list"	"mrs:cluster:list", "mrs:tag:listResource", "mrs:tag:list", "dws:*:get*", "dws:*:list*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Adding/Deleting tags	"dws:tag:addAndDelete"	"dws:*:get*", "dws:*:list*", "dws:openAPITag:update", "dws:openAPITag:getResourceTag",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Editing tags	"dws:tag:edit"	"dws:*:get*", "dws:*:list*", "dws:openAPITag:update", "dws:openAPITag:getResourceTag",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Creating a snapshot	"dws:snapshot:create"	"dws:*:get*", "dws:*:list*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Obtaining the snapshot list	"dws:snapshot:list"	--	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Viewing the snapshot list of a cluster	"dws:clusterSnapshot:list"	"dws:cluster:list", "dws:openAPICluster:getDetail"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Deleting snapshots	"dws:snapshot:delete"	"dws:snapshot:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Copying snapshots	"dws:snapshot:copy"	"dws:snapshot:list", "dws:snapshot:create"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Restoring data to a new cluster	"dws:cluster:restore"	"dws:*.get*", "dws:*.list*", "ecs:*.get*", "ecs:*.list*", "ecs:*.create*", "vpc:*.get*", "vpc:*.list*", "vpc:*.create*", "evs:*.get*", "evs:*.list*", "evs:*.create*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Resizing a cluster	"dws:cluster:resize"	"dws:*:get*", "dws:*:list*", "ecs:*:get*", "ecs:*:list*", "ecs:*:create*", "ecs:*:delete*", "vpc:*:get*", "vpc:*:list*", "vpc:*:create*", "vpc:*:delete*", "evs:*:get*", "evs:*:list*", "evs:*:create*", "evs:*:delete"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project
Switchback	"dws:cluster:switchover"	"dws:*:get*", "dws:*:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project
Querying the ELB list	"dws:elb:list"	"dws:*:get*", "dws:*:list*", "elb:*:get*", "elb:*:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project
Associating ELB	"dws:elb:bind"	"dws:*:get*", "dws:*:list*", "ecs:*:get*", "ecs:*:list*", "vpc:*:get*", "vpc:*:list*", "evs:*:get*", "evs:*:list*", "elb:*:get*", "elb:*:list*", "elb:*:delete*", "elb:*:create"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project

Operation	Permission	Dependent Permission	Scope
Disassociating ELB	"dws:elb:unbind"	"dws:*.get*", "dws:*.list*", "ecs:*.get*", "ecs:*.list*", "vpc:*.get*", "vpc:*.list*", "evs:*.get*", "evs:*.list*", "elb:*.get*", "elb:*.list*", "elb:*.delete*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Querying snapshot configurations	"dws:snapshotConfig:list"	"dws:*.get*", "dws:*.list*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Updating a snapshot policy	"dws:backupPolicyDetail:update"	"dws:*.get*", "dws:*.list*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Deleting a snapshot policy	"dws:backupPolicy:delete"	"dws:*.get*", "dws:*.list*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Querying a snapshot policy	"dws:backupPolicy:list"	"dws:cluster:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Querying cluster encryption information	"dws:clusterEncryptInfo:list"	"dws:*.get*", "dws:*.list*", "KMS Administrator"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Creating an agent	"dws:createAgency:create"	"dws:*.get*", "dws:*.list*", "security administrator"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Querying OBS bucket information	"dws:queryBuckets:list"	"dws:*.get*", "dws:*.list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Adding a node	"dws:expandWithExistendNodes:update"	"dws:*.get*", "dws:*.list*", "ecs:*.get*", "ecs:*.list*", "ecs:*.create*", "vpc:*.get*", "vpc:*.list*", "vpc:*.create*", "vpc:*.update*", "evs:*.get*", "evs:*.list*", "evs:*.create*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Deleting a DR backup	"dws:disasterRecovery:delete"	"dws:*.get*", "dws:*.list*", "ecs:*.get*", "ecs:*.list*", "ecs:*.delete*", "vpc:*.get*", "vpc:*.list*", "vpc:*.delete*", "evs:*.get*", "evs:*.list*", "evs:*.delete*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Creating a DR backup	"dws:disasterRecovery:create"	"dws::get*", "dws::list*", "ecs::get*", "ecs::list*", "ecs::create*", "vpc::get*", "vpc::list*", "vpc::create*", "evs::get*", "evs::list*", "evs::create*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project
Other DR and backup operations	"dws:disasterRecovery:otherOperate"	"dws::get*", "dws::list*", "ecs::get*", "ecs::list*", "ecs::create*", "vpc::get*", "vpc::list*", "vpc::create*", "evs::get*", "evs::list*", "evs::create*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project
Querying DR and backup operations	"dws:disasterRecovery:get"	"dws::get*", "dws::list*", "ecs::get*", "ecs::list*", "vpc::get*", "vpc::list*", "evs::get*", "evs::list*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project
Adding a CN	"dws:module:install"	"dws::get*", "dws::list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project

Operation	Permission	Dependent Permission	Scope
Deleting a CN	"dws:module:uninstall"	"dws*:get*", "dws*:list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Removing nodes	"dws:clusterNodes:operate"	"dws*:get*", "dws*:list*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Updating the node alias	dws:instanceAliasName:update	dws:cluster:list	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Redistributing data	"dws:redistribution:operate"	"dws*:get*", "dws*:list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Querying redistribution	"dws:redistributionInfo:list"	"dws*:get*", "dws*:list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Stopping redistribution	"dws:redistribution:suspend"	"dws*:get*", "dws*:list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Resuming redistribution	"dws:redistribution:recover"	"dws*:get*", "dws*:list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Adding disk capacity	"dws:disk:expand"	"dws:*.get*", "dws:*.list*", "ecs:*.get*", "ecs:*.list*", "ecs:*.create*", "vpc:*.get*", "vpc:*.list*", "vpc:*.create*", "evs:*.get*", "evs:*.list*", "evs:*.create*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Scaling in a cluster	"dws:cluster:shrink"	"dws:*.get*", "dws:*.list*", "dws:createAgency:create", "ecs:*.get*", "ecs:*.list*", "ecs:*.delete*", "vpc:*.get*", "vpc:*.list*", "vpc:*.delete*", "evs:*.get*", "evs:*.list*", "evs:*.delete"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Querying product specifications	"dws:specProduct:list"	"dws:*.get*", "dws:*.list*", "ecs:*.get*", "ecs:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Changing from pay-per-use to yearly/monthly	"dws:ondemandToPeriod:operate"	"dws:*.get*", "dws:*.list*", "ecs:*.get*", "ecs:*.list*", "ecs:*.create*", "vpc:*.get*", "vpc:*.list*", "vpc:*.create*", "vpc:securityGroupRules:delete", "evs:*.get*", "evs:*.list*", "evs:*.create*", "bss:coupon:view", "bss:order:pay", "bss:order:view", "bss:contract:update", "bss:balance:view", "bss:renewal:view", "bss:unsubscribe:update", "bss:renewal:update", "bss:order:update"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Modifying a yearly/monthly cluster	"dws:periodCluster:modify"	"dws:*.get*", "dws:*.list*", "ecs:*.get*", "ecs:*.list*", "ecs:*.delete*", "vpc:*.get*", "vpc:*.list*", "vpc:*.delete*", "evs:*.get*", "evs:*.list*", "evs:*.delete*", "bss:coupon:view", "bss:order:pay", "bss:order:view", "bss:contract:update", "bss:balance:view", "bss:renewal:view", "bss:unsubscribe:update", "bss:renewal:update", "bss:order:update"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Creating a yearly/monthly cluster	"dws:periodCluster:create"	"dws:*.get*", "dws:*.list*", "ecs:*.get*", "ecs:*.list*", "ecs:*.create*", "vpc:*.get*", "vpc:*.list*", "vpc:*.create*", "evs:*.get*", "evs:*.list*", "evs:*.create*", "bss:coupon:view", "bss:order:pay", "bss:order:view", "bss:contract:update", "bss:balance:view", "bss:renewal:view", "bss:unsubscribe:update", "bss:renewal:update", "bss:order:update"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project
Performing a check before adding disk capacity to a yearly/monthly cluster	"dws:periodExpandPre-check:operate"	"dws:*.get*", "dws:*.list*", "ecs:*.get*", "ecs:*.list*", "ecs:*.create*", "vpc:*.get*", "vpc:*.list*", "vpc:*.create*", "evs:*.get*", "evs:*.list*", "evs:*.create*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project

Operation	Permission	Dependent Permission	Scope
Binding the management plane IP address	"dws:bindManagement:operate"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Obtaining user authorization	"dws:checkAuthorize:operate"	"dws:*.get*", "dws:*.list*", "dws:checkSupport:operate"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Authorizing a user	"dws:authorize:operate"	"dws:*.get*", "dws:*.list*", "dws:checkSupport:operate"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Querying user databases	"dws:userDatabase:list"	"dws:*.get*", "dws:*.list*", "dws:checkSupport:operate"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Querying user schemas	"dws:schemas:list"	"dws:*.get*", "dws:*.list*", "dws:checkSupport:operate"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Querying user tables	"dws:tables:list"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Restoring tables	"dws:tableRestore:operate"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Checking the name of the table to be restored	"dws:tableRestoreCheck:operate"	"dws:*.get*", "dws:*.list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Checking whether a cluster supports fine-grained backup	"dws:checkSupport:operate"	"dws:*.get*", "dws:*.list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Querying the list of flavors that can be changed	"dws:supportFlavors:list"	"dws:*.get*", "dws:*.list*",	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Changing the node flavor	"dws:specResize:operate"	"dws:*.get*", "dws:*.list*", "ecs:*.get*", "ecs:*.list*", "ecs:*.create*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Stopping snapshot creation	"dws:snapshot:stop"	"dws:snapshot:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Terminating a session	"dws:dmsSession:terminate"	"dws:dmsGrpcOuter:operation"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Workload report operations	"dws:dmsWorkloadDiagnosisReport:create"	"dws:dmsGrpcOuter:operation"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Modifying an alarm rule	"dws:dmsAlarmRule:update"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> • Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Enabling an alarm rule	"dws:dmsAlarmRule:enable"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> • Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Enabling a cluster alarm	"dws:dmsClusterAlarm:enable"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> • Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Disabling a cluster alarm	"dws:dmsClusterAlarm:disable"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> • Scope: <ul style="list-style-type: none"> - Project - Enterprise project
gRPC external service	"dws:dmsGrpcOuter:operation"	"dws:dmsQuery:list", "dws:cluster:setSecuritySettings", "obs:bucket:ListAllMyBuckets"	<ul style="list-style-type: none"> • Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Adding a SQL probe	"dws:dmsProbe:add"	"dws:dmsGrpcOuter:operation"	<ul style="list-style-type: none"> • Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Modifying a SQL probe	"dws:dmsProbe:update"	"dws:dmsGrpcOuter:operation"	<ul style="list-style-type: none"> • Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Deleting a SQL probe	"dws:dmsProbe:delete"	"dws:dmsGrpcOuter:operation"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Enabling or disabling a SQL probe	"dws:dmsProbe:enable"	"dws:dmsGrpcOuter:operation"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Creating a User panel	"dws:dmsUserBoard:create"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Modifying a user panel	"dws:dmsUserBoard:update"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Deleting a user panel	"dws:dmsUserBoard:delete"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Terminating a query	"dws:dmsQuery:terminate"	"dws:dmsGrpcOuter:operation"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Enabling or disabling DMS	"dws:dmsService:enableOrDisable"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Modifying DMS storage configurations	"dws:dmsStorageConfig:modify"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Obtaining, or creating a DDL review	"dws:dmsDdlExamine:getOrCreate"	"dws:dmsGrpcOuter:operation"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Workload snapshot operations	"dws:dmsWorkloadDiagnosisSnapshot:create"	"dws:dmsGrpcOuter:operation"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Creating an alarm rule	"dws:dmsAlarmRule:add"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Deleting an alarm rule	"dws:dmsAlarmRule:delete"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Executing a SQL probe	"dws:dmsProbe:execute"	"dws:dmsGrpcOuter:operation"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Deleting a monitoring item	"dws:dmsPerformanceMonitor:delete"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Enabling or disabling DMS monitoring metrics	"dws:dmsCollectItem:enableOrDisable"	"dws:dmsGrpcOuter:operation"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Modifying DMS monitoring configurations	"dws:dmsCollectConfig:modify"	"dws:dmsGrpcOuter:operation"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Conditional query	"dws:dmsQuery:list"	"dws:cluster:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
OpenAPI Conditional Query	"dws:dmsOpenapiQuery:list"	"dws:cluster:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Disabling an alarm rule	"dws:dmsAlarmRule:disable"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Deleting an alarm record	"dws:dmsAlarmRecord:delete"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Checking SQL probes	"dws:dmsProbe:check"	"dws:dmsGrpcOuter:operation"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Adding a monitoring item	"dws:dmsPerformanceMonitor:add"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> • Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Modifying monitoring metrics	"dws:dmsPerformanceMonitor:update"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> • Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Downloading historical monitoring trend	"dws:dmsTrendHistory:download"	"dws:dmsQuery:list"	<ul style="list-style-type: none"> • Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Obtaining cluster ring information	"dws:ring:list"	"dws:*.get*", "dws:*.list*"	<ul style="list-style-type: none"> • Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Obtaining the cluster process topology	"dws:processTopo:list"	"dws:*.get*", "dws:*.list*"	<ul style="list-style-type: none"> • Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Querying intelligent O&M information	"dws:operationalTask:get"	"dws:*.get*", "dws:*.list*"	<ul style="list-style-type: none"> • Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Intelligent O&M Operations	"dws:operationalTask:operate"	"dws:*.get*", "dws:*.list*"	<ul style="list-style-type: none"> • Scope: <ul style="list-style-type: none"> - Project - Enterprise project

Operation	Permission	Dependent Permission	Scope
Adding, deleting, and modifying a logical cluster	"dws:logicalCluster:operate"	"dws:*.get*", "dws:*.list*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project
Querying a logical cluster	"dws:logicalCluster:get"	"dws:*.get*", "dws:*.list*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project
Elastic logical cluster planning	"dws:logicalClusterPlan:operate"	"dws:*.get*", "dws:*.list*", "dws:logicalCluster:*", "dws:cluster:scaleOut", "iam:agencies:*", "iam:permissions:*Agency*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project
Creating an endpoint service	"dws:vpcEndpointService:create"	"dws:*.get*", "dws:*.list*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project
Querying the resource management list	"dws:workLoadManager:get"	"dws:*.get*", "dws:*.list*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project
Resource management operations	"dws:workLoadManager:operate"	"dws:*.get*", "dws:*.list*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project
LTS operations	"dws:ltsAccess:operate"	"dws:*.get*", "dws:*.list*"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> – Project – Enterprise project

Operation	Permission	Dependent Permission	Scope
Querying LTS Information	"dws:ltsAccess:get"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project - Enterprise project
Querying events	"dws:event:list"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project
Querying event specifications	"dws:event:list"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project
Querying event subscriptions	"dws:eventSub:list"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project
Creating an event subscription	"dws:eventSub:create"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project
Updating an event subscription	"dws:eventSub:update"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project
Deleting an event subscription	"dws:eventSub:delete"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project
Querying alarm statistics	"dws:alarmStatistic:list"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project
Querying alarm details	"dws:alarmDetail:list"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project
Querying alarm configurations	"dws:alarmConfig:list"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project
Querying alarm subscriptions	"dws:alarmSub:list"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project
Creating an alarm subscription	"dws:alarmSub:create"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project
Updating an alarm subscription	"dws:alarmSub:update"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project
Deleting an alarm subscription	"dws:alarmSub:delete"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> ● Scope: <ul style="list-style-type: none"> - Project

Operation	Permission	Dependent Permission	Scope
Delivering cluster upgrade operations (upgrade, rollback, submission, and retry)	"dws:cluster:doUpdate"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> Scope: <ul style="list-style-type: none"> Project
Querying the available upgrade paths of a cluster	"dws:cluster:getUpgradePaths"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> Scope: <ul style="list-style-type: none"> Project
Querying cluster upgrade records	"dws:cluster:getUpgradeRecords"	"dws:*.get*", "dws:*.list"	<ul style="list-style-type: none"> Scope: <ul style="list-style-type: none"> Project
Starting a cluster	"dws:cluster:startCluster"	"dws:*.get*", "dws:*.list*", "ecs:*.get*", "ecs:*.list*", "ecs:*.start", "ecs:*.stop"	<ul style="list-style-type: none"> Scope: <ul style="list-style-type: none"> Project
Stopping a cluster	"dws:cluster:stopCluster"	"dws:*.get*", "dws:*.list*", "ecs:*.get*", "ecs:*.list*", "ecs:*.start", "ecs:*.stop"	<ul style="list-style-type: none"> Scope: <ul style="list-style-type: none"> Project

Authorization Using the Fine-Grained Permission Policy

Step 1 Log in to the IAM console and create a custom policy.

Refer to the following to create the policy:

- Use the IAM administrator account, that is, the user in the admin user group, because only the IAM administrator has the permissions to create users and user groups and modify user group permissions.
- GaussDB(DWS) is a project-level service, so its **Scope** must be set to **Project-level services**. If this policy is required to take effect for multiple projects, authorization is required to each project.
- Two GaussDB(DWS) policy templates are preconfigured on IAM. When creating a custom policy, you can select either of the following templates and modify the policy authorization statement based on the template:

- **DWS Admin:** has all execution permissions on GaussDB(DWS).
- **DWS Viewer:** has the read-only permission on GaussDB(DWS).
- You can add permissions corresponding to GaussDB(DWS) operations or RESTful APIs listed in [List of Supported Actions](#) to the action list in the policy authorization statement, so that the policy can obtain the permissions.
For example, if **dws:cluster:create** is added to the action list of a policy statement, the policy has the permission to create or restore clusters.
- If you want to use other services, grant related operation permissions on these services. For details, see the help documents of related services.
For example, when creating a data warehouse cluster, you need to configure the VPC to which the cluster belongs. To obtain the VPC list, add permission **vpc:*:get*** to the policy statement.

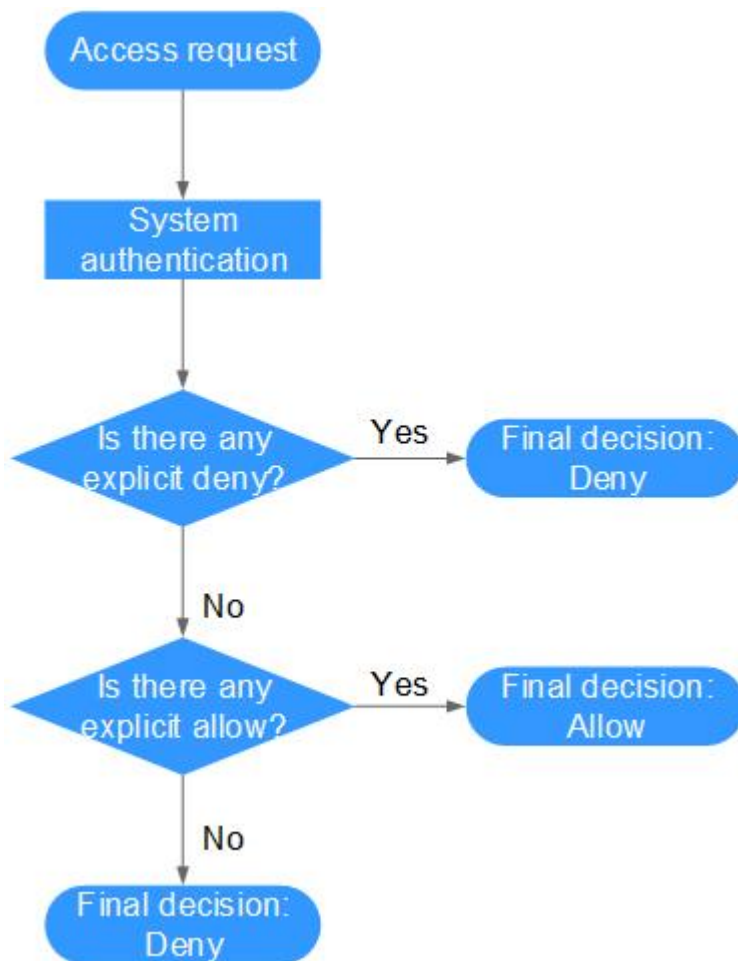
Step 2 Create a user group.

Step 3 Add users to the user group and grant the new custom policy to the user group so that users in it can obtain the permissions defined by the policy.

----End

Authentication Logic

If a user is granted permissions of multiple policies or of only one policy containing both Allow and Deny statements, then authentication starts from the Deny statements. The following figure shows the authentication logic for resource access.

Figure 14-4 Authentication logic**NOTE**

The actions in each policy bear the OR relationship.

1. A user accesses the system and makes an operation request.
2. The system evaluates all the permissions policies assigned to the user.
3. In these policies, the system looks for explicit deny permissions. If the system finds an explicit deny that applies, it returns a decision of Deny, and the authentication ends.
4. If no explicit deny is found, the system looks for allow permissions that would apply to the request. If the system finds an explicit allow permission that applies, it returns a decision of Allow, and the authentication ends.
5. If no explicit allow permission is found, IAM returns a decision of Deny, and the authentication ends.

14.3.4 RBAC Syntax of RBAC Policies

Policy Structure

An RBAC policy consists of a Version, a Statement, and Depends.

Figure 14-5 RBAC policy structure



Policy Syntax


When selecting a policy for a user group, click  below the policy to view the details of the policy. The **DWS Administrator** policy is used as an example to describe the syntax of RBAC policies.

Figure 14-6 Syntax of RBAC Policies

Assign Permissions

If the policies listed here do not contain the permissions you need, [modify existing policies](#) or [create new policies](#).

View Selected (0) All policies DWS X Q C View by Policy View by Project

<input type="checkbox"/>	Policy Name	Description	Project [Region]
<input checked="" type="checkbox"/>	DWS Administrator	Data Warehouse Service Administrator	All projects

JSON

```

1- {
2-   "Version": "1.0",
3-   "Statement": [
4-     {
5-       "Effect": "Allow",
6-       "Action": [
7-         "DWS:DWS:*"
8-       ]
9-     }
10-  ],
11-   "Depends": [
12-     {
13-       "catalog": "BASE",
14-       "display_name": "Server Administrator"
15-     },
16-     {
17-       "catalog": "BASE",
18-       "display_name": "Tenant Guest"
19-     }
20-  ]
21- }
    
```

```

{
  "Version": "1.0",
  "Statement": [
    {
      "Effect": "Allow",
    }
  ]
}
    
```

```

        "Action": [
            "dws:dws:*"
        ]
    },
    "Depends": [
        {
            "catalog": "BASE",
            "display_name": "Server Administrator"
        },
        {
            "catalog": "BASE",
            "display_name": "Tenant Guest"
        }
    ]
}

```

Parameter		Meaning	Value
Version		Policy version	The value is fixed to 1.0 .
Statement	Action	Operations to be performed on GaussDB(DWS)	Format: <i>Service name:Resource type:Operation</i> . dws:dws:* : Permissions for performing all operations on all resource types in GaussDB(DWS).
	Effect	Whether the operation defined in an action is allowed	<ul style="list-style-type: none"> Allow Deny
Depends	catalog	Name of the service to which dependencies of a policy belong	Service name Example: BASE
	display_name	Name of a dependent policy	Policy name Example: Server Administrator

 **NOTE**

When using RBAC for authentication, pay attention to the **Depends** parameter and grant other dependent permissions at the same time.

For example, the **DWS Administrator** permission depends on the **Server Administrator** and **Tenant Guest** permissions. When granting the **DWS Administrator** permission to users, you also need to grant the two dependent permissions to the users.

14.4 Protection for Mission-Critical Operations

Scenario

GaussDB(DWS) protects mission-critical operations. If you want to perform a mission-critical operation on the management console, you must enter a

credential for identity verification. You can perform the operation only after your identity is verified. For account security, it is a good practice to enable operation protection. The setting will take effect for both the account and users under the account.

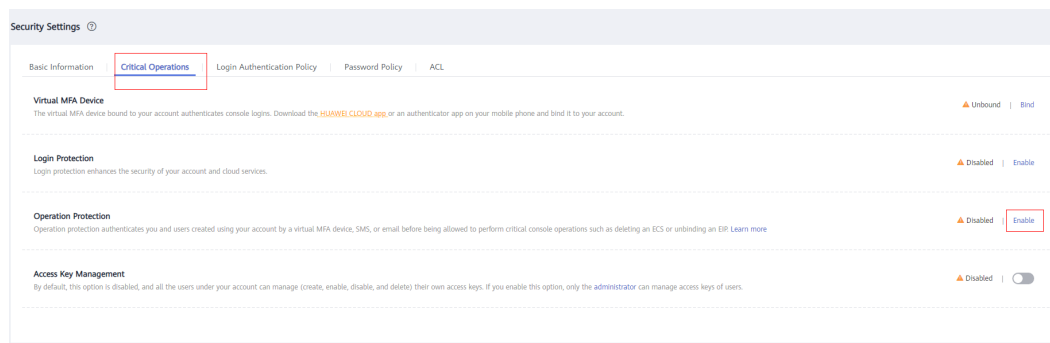
Currently, the following operations are supported: scaling out a cluster, deleting a cluster, restarting a cluster, adding a CN, and deleting a CN.

Enabling Critical Operation Protection

Operation protection is disabled by default. To enable it, perform the following steps:

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Move the cursor to the username in the upper right corner of the page and click **Security Settings** from the drop-down list.
- Step 3** On the **Security Settings** page, click the **Critical Operations** tab. Click **Enable** in the **Operation Protection** area.

Figure 14-7 Critical Operations



- Step 4** On the **Operation Protection** page, select **Enable** to enable operation protection.

NOTE

- When IAM users created using your account perform a critical operation, they will be prompted to choose a verification method from email, SMS, and virtual MFA device.
 - If a user is only associated with a mobile number, only SMS verification will be available.
 - If a user is only associated with an email address, only email verification will be available.
 - If the user has not bound an email address, a mobile number, or a virtual MFA device, the user needs to bind one to continue with the critical operation.
- Change your phone number or email address for verification in **My Account** on the management console.

- Step 5** After operation protection is enabled, when you perform a mission-critical operation, the system will protect the operation.

For example, when you delete a cluster, a verification dialog box for mission-critical operation protection is displayed. You need to select a mode to perform verification. This helps avoid risks and losses caused by misoperations.

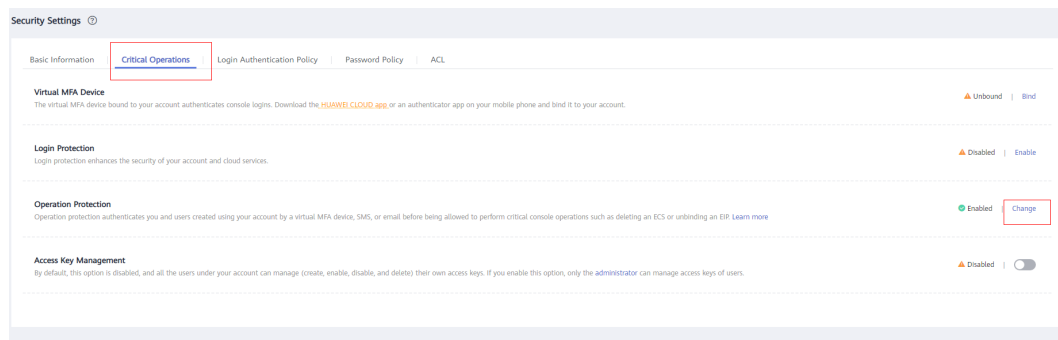
----End

Disabling Operation Protection

To disable operation protection, perform the following steps:

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Move the cursor to the username in the upper right corner of the page and click **Security Settings** from the drop-down list.
- Step 3** On the **Security Settings** page, click the **Critical Operations** tab. Click **Change** in the **Operation Protection** area.

Figure 14-8 Modifying operation protection settings



- Step 4** On the **Operation Protection** page, select **Disable** and click **OK**.

----End

15 Resource Management

15.1 Overview

The system resources (CPU, memory, I/O, and storage resources) of a database are limited. When multiple types of services (such as data loading, batch analysis, and real-time query) are running at the same time, they may compete for resources and hinder operations. As a result, the throughput decreases and the overall query performance deteriorates. To avoid this problem, resources must be properly allocated.

GaussDB(DWS) provides the resource management function. You can put resources into different resource pools, which are isolated from each other. Then, you can associate database users with these resource pools. When a user starts a SQL query, the query will be transferred to the resource pool associated with the user. You can specify the number of queries that can be concurrently executed in a resource pool, the upper limit of memory used for a single query, and the memory and CPU resources that can be used by a resource pool. In this way, you can limit and isolate the resources occupied by different workloads, properly utilizing resources to process hybrid database loads and achieve high query performance. After a cluster is converted into a logical cluster, you can create, modify, or delete a resource pool in the logical cluster.

NOTICE

- This feature is supported only by clusters of version 8.0 or later.
 - Resources cannot be managed during offline scale-out. If a resource management plan is enabled, stop it before performing offline scale-out.
 - The hybrid data warehouse (standalone) does not support resource management.
-

Enabling or Disabling Resource Management

You can enable or disable resource management, and configure the maximum global concurrency. **Max. Concurrent Queries** refers to the maximum concurrent

queries on a single CN. If you disable **Resource Management**, all resource management functions will be unavailable.

Resource Management Functions

The resource management functions of GaussDB(DWS) can be classified into the following types based on managed resources:

- **Computing resource management.** It is implemented using resource pools. Computing resources are isolated and controlled to prevent cluster-level issues caused by abnormal SQL queries. Computing resource management includes concurrency management, memory management, CPU management, and exception rules. For details, see [Resource Pool](#).
- **Storage space management:** Storage is managed at user and schema level to prevent disk exhaustion, which makes the database read only. For details, see [Workspace Management](#).
- **Resource management plan:** Resources are managed automatically based on a preconfigured plan, which can flexibly cope with complex scenarios. For details, see [Importing or Exporting a Resource Management Plan](#).

The resource management functions of GaussDB(DWS) can be classified into the following types based on when they are implemented:

- **Management before a query**

The service checks whether there are sufficient resources for a query. If there are, the query can be executed. If there are not, the query waits in a queue, and can be executed only after resources are released by other queries. Concurrency and memory are managed in this phase.
- **Management during a query**

During query execution, resources used by the query are managed and controlled to prevent cluster exceptions caused by time-consuming SQL statements. Memory, CPU, storage space, and exception rules are managed in this phase.

Simple and Complex Queries

GaussDB(DWS) supports fine-grained resource management. Before workload management is implemented, queries are classified into complex queries (with long execution time and high resource consumption) and simple queries (with short execution time and low resource consumption). Simple and complex queries also differ in their estimated memory usage.

- The estimated memory usage of a simple query is less than 32 MB.
- The estimated memory usage of a complex query is 32 MB or higher.

In a hybrid load database, complex queries often occupy a large number of resources for a long time. A simple query queued after a complex query is time consuming, because it has to wait for the complex query to complete and resources to be freed up. To improve execution efficiency and system throughput, GaussDB(DWS) provides the short query acceleration function, managing simple queries separately.

- If short query acceleration is enabled, simple queries and complex queries are managed separately. Simple queries do not need to compete with complex queries for resources.
- If short query acceleration is disabled, simple and complex queries are under the same resource management rules.

To prevent a large number of simple queries from consuming too many resources during acceleration, concurrency management is performed on the queries. Resource management is not performed, because it may affect query performance and system throughput.

NOTE

Queries are categorized based on estimated memory usage, but the estimation does not equal the actual usage, nor does it reflect the query duration or CPU usage. In resource pools that are insensitive to performance and only run specific services, you can disable short query acceleration to manage resources and handle exceptions for simple queries.

15.2 Resource Pool

15.2.1 Feature Description

GaussDB(DWS) resource pools provide concurrency management, memory management, CPU management, and exception rules.

Concurrency Management

Concurrency represents the maximum number of concurrent queries in a resource pool. Concurrency management can limit the number of concurrent queries to reduce resource contention and improve resource utilization.

In the **Short Query Configuration** area, you can enable or disable the short query acceleration function. To change the number of simple statements (-1 by default. 0 or -1 indicates that the concurrent short queries are not controlled), you can enable short query acceleration.

The concurrency management rules are as follows:

- If short query acceleration is enabled, complex queries are under resource pool concurrency control, and simple queries are under short query concurrency control.
- If short query acceleration is disabled, complex and simple queries are both under resource pool concurrency control. Short query concurrency control is invalid.

Memory Management

Each resource pool occupies a certain percentage of memory.

Memory management aims to prevent out of memory (OOM) in a database, isolate the memory of different resource pools, and to control memory usage. Memory is managed from the following aspects:

- Global memory management

To prevent OOM, set the global memory upper limit (**max_process_memory**) to a proper value. Global memory management before a query controls memory usage to prevent OOM management. Global memory management during a query prevents errors during query execution.

 - Management before a query

The service checks the estimated memory usage of a query in the slow queue, and compares it with the actual usage. The estimation will be adjusted if it is smaller than the actual usage. Before a query is executed, the service checks whether the available memory is sufficient for the query. If yes, the query can be executed directly. If no, the query needs to be queued and executed after other queries release resources.
 - Management during a query

During a query, the service checks whether the requested memory exceeds a certain limit. If yes, an error will be reported, and memory occupied by the query will be released.
- Resource pool memory management

Resource pool memory management puts a limit on dedicated quotas. A workload queue can only use the memory allocated to it, and cannot use idle memory in other resource pools.

The resource pool memory is allocated in percentage. The value range is 0 to 100. The value **0** indicates that the resource pool does not perform memory management. The value **100** indicates that the resource pool performs memory management and can use all the global memory.

The sum of memory percentages allocated to all resource pools cannot exceed 100. Resource pool memory management is performed only before a query in the slow queue starts. It works in a way similar to the global memory management before a query. Before a query in the slow queue in a resource pool is executed, its memory usage is estimated. If the estimation is greater than the resource pool memory, the query needs to be queued and can be executed only after earlier queries in the pool are complete and resources released.

CPU Management

CPU share and CPU limit can be managed.

- CPU share: If the system is heavily loaded, CPU resources are allocated to resource pools based on the specific CPU shares. If the system not busy, this configuration does not take effect.
- CPU limit: It specifies the maximum number of CPU cores used by a resource pool. The resource usage of jobs in the resource pool cannot exceed this limit no matter whether the system is busy or not.

In the Resource Configuration area, you can modify the CPU time limit and CPU usage limit.

Choose either of the preceding management methods as needed. In CPU share management, CPUs can be shared and fully utilized, but resource pools are not isolated and may affect the query performance of each other. In CPU limit management, the CPUs of different resource pools are isolated, but this may result in the waste of idle resources.

 NOTE

The CPU limit is supported only by clusters of version 8.1.3 or later.

Exception Rules

To avoid query blocking or performance deterioration, you can configure exception rules to let the service automatically identify and handle abnormal queries, preventing slow SQL statements from occupying too many resources for a long time.

In the **Associated Exception Rules** area, you can view the exception rules bound to the current resource pool, bind new exception rules, and unbind existing exception rules. For more information, see [Exception parameters](#).

 NOTE

- The cluster version 8.2.1 and later supports downgradation of exception rules. All exception rules support downgradation behaviors. After downgradation, only network resource preemption is downgraded to a low priority. Downgraded network queries are scheduled only when there is no normal queries.
- Only clusters of 8.2.0 and later versions support associating exception rules. GaussDB(DWS) 3.0 does not support this function. For cluster versions earlier than 8.2.0, see [Step 7.3](#).

Table 15-1 Exception rule parameters

Parameter	Description	Value Range (0 Means No Limit)	Operation
Blocking Time	Job blocking time. It refers to the total time spent in global and local concurrent queuing. The unit is second. For example, if the blocking time is set to 300s, a job executed by a user in the resource pool will be terminated after being blocked for 300 seconds.	An integer in the range 1 to 2,147,483,647. The value 0 indicates no limit.	Terminated, Downgraded, or Not limited
Execution Time	Time that has been spent in executing the job, in seconds. For example, if Time required for execution is set to 100s, a job executed by a user in the resource pool will be terminated after being executed for more than 100 seconds.	An integer in the range 1 to 2,147,483,647. The value 0 indicates no limit.	Terminated, Downgraded, or Not limited

Parameter	Description	Value Range (0 Means No Limit)	Operation
Total CPU time on all DNs.	Total CPU time spent in executing a job on all DNs, in seconds.	An integer in the range 1 to 2,147,483,647 . The value 0 indicates no limit.	Terminated, Downgraded, or Not limited
Interval for Checking CPU Skew Rate	Interval for checking the CPU skew, in seconds. This parameter must be set together with Total CPU Time on All DNs .	An integer in the range 1 to 2,147,483,647 . The value 0 indicates no limit.	Terminated, Downgraded, or Not limited
Total CPU Time Skew Rate on All DNs	CPU time skew rate of a job executed on DNs. The value depends on the setting of Interval for Checking CPU Skew Rate .	An integer in the range 1 to 100. The value 0 indicates no limit.	Terminated, Downgraded, or Not limited
Data Spilled to Disk Per DN	Allowed maximum job data spilled to disks on a DN. The unit is MB. NOTE This rule is supported only by clusters of version 8.2.0 or later.	An integer in the range 1 to 2,147,483,647 . The value 0 indicates no limit.	Terminated, Downgraded, or Not limited
Average CPU Usage Per DN	Average CPU usage of a job on each DN. If Interval for Checking CPU Skew Rate is configured, the interval takes effect for this parameter. If the interval is not configured, the check interval is 30 seconds by default. NOTE This rule is supported only by clusters of version 8.2.0 or later.	An integer in the range 1 to 100. The value 0 indicates no limit.	Terminated, Downgraded, or Not limited
Maximum Bandwidth on a Single DN	Maximum network bandwidth (MB) for a job on a single DN. NOTE This rule is supported only by clusters of version 8.2.1 or later.	An integer in the range 1 to 2,147,483,647 . The value 0 indicates no limit.	Terminated, Downgraded, or Not limited

15.2.2 Creating a Resource Pool

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters**. Click the name of a cluster.

Step 3 Choose **Resource Management Configurations**.

Step 4 Click **Add Resource Pool**.

 **NOTE**

Up to 63 resource pools can be created.

Step 5 Configure the resource pool. For more information, see [Table 15-2](#).

Table 15-2 Resource pool parameters

Parameter	Description	Default Value
Name	Resource pool name	-
CPU Resource (%)	<ul style="list-style-type: none">• CPU share: Percentage of CPU time that can be used by users associated with the current resource pool to execute jobs. The value is an integer ranging from 1 to 99.• CPU limit: Maximum percentage of CPU cores used by a database user in a resource pool. The value is an integer ranging from 0 to 100. 0 indicates no limit. <p>NOTE</p> <ul style="list-style-type: none">• The sum of the parameter values of all the resource pools cannot exceed 99%. If there is only one resource pool, the CPU share parameter does not take effect.• The CPU share parameter takes effect only when CPU contention occurs. For example, resource pools A and B are bound to CPU 1. If A and B are both running, the parameter takes effect. If there is only A running, the parameter does not take effect.• The sum of the CPU limits of all the resource pools cannot exceed 100%. The default value is 0.• The CPU limit is supported only by clusters of version 8.1.3 or later.	-
Memory Resource (%)	Percentage of the memory that can be used by a resource pool. You can manage memory and query concurrency separately or jointly. Under joint management, jobs can be delivered only when both the memory and concurrency conditions are met.	0 (not limited)

Parameter	Description	Default Value
Storage Resource (MB)	Size of the available space for permanent tables. This parameter indicates the total tablespace of all DNs in a resource pool. Available space of a single DN = Configured value/Number of DNs.	-1 (not limited)
Complex Statement Concurrency	Maximum number of concurrent queries in a resource pool. You can manage memory and query concurrency separately or jointly. Under joint management, jobs can be delivered only when both the memory and concurrency conditions are met.	10
Network Bandwidth Weight	Weight for network scheduling. The value is an integer ranging from 1 to 2147483647. The default value is -1. CAUTION Only cluster 8.2.1 and later versions support the network bandwidth weight.	-1 (not limited)

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

----End

15.2.3 Modifying a Resource Pool

You can modify the parameters of a resource pool on the resource management page.

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters**. Click the name of a cluster.

Step 3 Choose **Resource Management Configurations**.

Step 4 In the **Resource Pools** drop-down list, click the name of a resource pool. The following configuration areas are displayed, including **Short Query Configuration**, **Resource Configuration**, **Exception Rule**, and **Associated User**.

Step 5 Modify the short query configuration. Set the parameters as required and click **Save** on the right.

Parameter	Description	Value
Short Query Acceleration	Whether to enable short query acceleration. This function is enabled by default.	Enable

Parameter	Description	Value
Concurrent Short Queries	A short query is a job whose estimated memory used for execution is less than 32 MB. The default value -1 indicates that the job is not controlled.	10

Step 6 Modify the resource configuration.

1. Click **Edit** on the right and modify the parameters according to [Table 15-3](#).

Table 15-3 Resource pool parameters

Parameter	Description	Default Value
Name	Resource pool name	-
CPU Resource (%)	<ul style="list-style-type: none">- CPU share: Percentage of CPU time that can be used by users associated with the current resource pool to execute jobs. The value is an integer ranging from 1 to 99.- CPU limit: Maximum percentage of CPU cores used by a database user in a resource pool. The value is an integer ranging from 0 to 100. 0 indicates no limit. NOTE <ul style="list-style-type: none">- The sum of the parameter values of all the resource pools cannot exceed 99%. If there is only one resource pool, the CPU share parameter does not take effect.- The CPU share parameter takes effect only when CPU contention occurs. For example, resource pools A and B are bound to CPU 1. If A and B are both running, the parameter takes effect. If there is only A running, the parameter does not take effect.- The sum of the CPU limits of all the resource pools cannot exceed 100%. The default value is 0.- The CPU limit is supported only by clusters of version 8.1.3 or later.	-
Memory Resource (%)	Percentage of the memory that can be used by a resource pool. You can manage memory and query concurrency separately or jointly. Under joint management, jobs can be delivered only when both the memory and concurrency conditions are met.	0 (not limited)

Parameter	Description	Default Value
Storage Resource (MB)	Size of the available space for permanent tables. This parameter indicates the total tablespace of all DNs in a resource pool. Available space of a single DN = Configured value/Number of DNs.	-1 (not limited)
Complex Statement Concurrency	Maximum number of concurrent queries in a resource pool. You can manage memory and query concurrency separately or jointly. Under joint management, jobs can be delivered only when both the memory and concurrency conditions are met.	10
Network Bandwidth Weight	Weight for network scheduling. The value is an integer ranging from 1 to 2147483647. The default value is -1. CAUTION Only cluster 8.2.1 and later versions support the network bandwidth weight.	-1 (not limited)

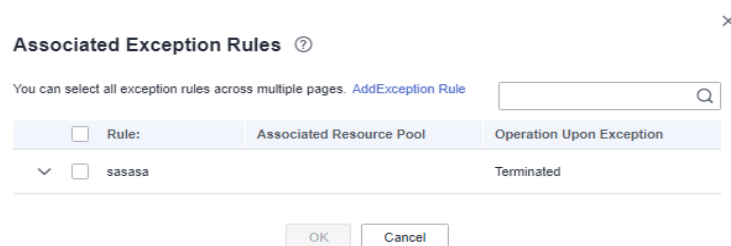
 **NOTE**

The CPU limit is supported only by clusters of version 8.1.3 or later.

2. Click **OK**.

Step 7 Associate exception rules.

1. Associate exception rules.



2. Unbind exception rules.

 NOTE

- Only clusters of 8.2.0 and later versions support associating exception rules. DWS 3.0 does not support this function. For cluster versions earlier than 8.2.0, see [Step 7.3](#).
- The default exception rules take effect for users not associated with any resource pools, and for users whose resource pools do not have any exception rules configured. If a user-defined rule is associated with a resource pool, this rule prevails in the pool.
 - The default exception rules are supported only by clusters of version 8.2.0 or later. After a cluster of an earlier version is upgraded to version 8.2.0 or later, the default exception rules do not take effect. You can create exception rules as needed.
 - The cluster version 8.2.1 supports downgradation of exception rules. All exception rules support downgradation behaviors. After downgradation, only network resource preemption is downgraded to a low priority. Downgraded network queries are scheduled only when there is no normal queries.
 - A resource pool can be associated with up to 16 exception rules.
- A resource pool can be associated with multiple groups of exception rules, which work in an OR way. One group of exception rules works if all its conditions are met. For example, a resource pool is associated with two groups of rules. One group specifies **elapsedtime=2400**, and the other group specifies **elapsedtime=1200** and **memsize=2000**. If the execution time of a job reaches 1200 seconds and the memory usage reaches 2000 MB, or if the execution time reaches 2400 seconds, the job will be terminated.

3. Modify the exception rules.

Modify rule parameters. See the following table for more information.

Table 15-4 Exception rule parameters

Parameter	Description	Value Range (0 Means No Limit)	Operation
Blocking Time	Job blocking time. It refers to the total time spent in global and local concurrent queuing. The unit is second. For example, if the blocking time is set to 300s, a job executed by a user in the resource pool will be terminated after being blocked for 300 seconds.	An integer in the range 1 to 2,147,483,647. The value 0 indicates no limit.	Terminated, Downgraded, or Not limited

Parameter	Description	Value Range (0 Means No Limit)	Operation
Execution Time	Time that has been spent in executing the job, in seconds. For example, if Time required for execution is set to 100s, a job executed by a user in the resource pool will be terminated after being executed for more than 100 seconds.	An integer in the range 1 to 2,147,483,647 . The value 0 indicates no limit.	Terminated, Downgraded, or Not limited
Total CPU time on all DNs.	Total CPU time spent in executing a job on all DNs, in seconds.	An integer in the range 1 to 2,147,483,647 . The value 0 indicates no limit.	Terminated, Downgraded, or Not limited
Interval for Checking CPU Skew Rate	Interval for checking the CPU skew, in seconds. This parameter must be set together with Total CPU Time on All DNs .	An integer in the range 1 to 2,147,483,647 . The value 0 indicates no limit.	Terminated, Downgraded, or Not limited
Total CPU Time Skew Rate on All DNs	CPU time skew rate of a job executed on DNs. The value depends on the setting of Interval for Checking CPU Skew Rate .	An integer in the range 1 to 100. The value 0 indicates no limit.	Terminated, Downgraded, or Not limited
Data Spilled to Disk Per DN	Allowed maximum job data spilled to disks on a DN. The unit is MB. NOTE This rule is supported only by clusters of version 8.2.0 or later.	An integer in the range 1 to 2,147,483,647 . The value 0 indicates no limit.	Terminated, Downgraded, or Not limited
Average CPU Usage Per DN	Average CPU usage of a job on each DN. If Interval for Checking CPU Skew Rate is configured, the interval takes effect for this parameter. If the interval is not configured, the check interval is 30 seconds by default. NOTE This rule is supported only by clusters of version 8.2.0 or later.	An integer in the range 1 to 100. The value 0 indicates no limit.	Terminated, Downgraded, or Not limited

Parameter	Description	Value Range (0 Means No Limit)	Operation
Maximum Bandwidth on a Single DN	Maximum network bandwidth (MB) for a job on a single DN. NOTE This rule is supported only by clusters of version 8.2.1 or later.	An integer in the range 1 to 2,147,483,647. The value 0 indicates no limit.	Terminated, Downgraded, or Not limited

NOTE

Exception rules allow you to control exceptions of jobs executed by users in a resource pool. Currently, you can configure the parameters listed in [Table 15-4](#).

- If you select **Terminate** or **Downgrade**, you need to set the corresponding time or percentage.
- If you select **Not limited**, the corresponding execution rule does not take effect.
- Only cluster 8.2.0 and later versions support the function of modifying exception rules.

Step 8 Associate users.**NOTE**

- The resources used by a user to run jobs can be controlled only after the user is added to a resource pool.
- A database user can be added to only one resource pool. Users removed from a resource pool can be added to another pool.
- Database administrators cannot be associated.
- If no resource pools are associated with a user, the user will be associated with **default_pool** by default, and its resource usage will be restricted by **default_pool**. The **default_pool** will be automatically created after resource management is enabled.

1. Click **Add**.
2. Select the users to be added from the current user list. You can select multiple users at a time.
3. Click **OK**.
4. To remove a user, click **Disassociate User** in the **Operation** column of the user.

----End

15.2.4 Deleting a Resource Pool

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters**. Click the name of a cluster.

Step 3 Choose **Resource Management Configurations**.

Step 4 In the **Resource Pools** area on the left, click the name of a resource pool.

Step 5 Click **Delete Resource Pool**.

 NOTE

After a resource pool is deleted, the users (if any) associated with this pool will be associated with the default resource pool instead.

----End

15.3 Resource Management Plan

15.3.1 Managing Resource Management Plans

Overview

The resource management plan is an advanced resource management feature provided by GaussDB(DWS). You can create a resource management plan, add multiple stages to the plan, and configure different queue resource ratios for the stages. After a plan is started, it automatically changes the resource configurations in different stages as scheduled. If you need to run services in different stages with different proportions of resources, you can create a resource management plan to automatically change resource configurations in different stages.

 NOTE

Resource management plans are supported only in version 8.1.0.100 or later.

Creating a Resource Management Plan

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Clusters**. Click the name of a cluster.
- Step 3** Go to the **Basic Information** page and click the **Resource Management** tab in the navigation pane on the left.
- Step 4** Click to the **Resource Management Plans** tab and click **Add**.
- Step 5** Enter a plan name and click **OK**.

NOTICE

- Before creating a resource management plan, you must design and create a resource pool. For details, see [Creating a Resource Pool](#).
 - You can create up to 10 resource management plans.
-

----End

Starting a Resource Management Plan

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Clusters**. Click the name of a cluster.

- Step 3** Go to the **Basic Information** page and click the **Resource Management** tab in the navigation pane on the left.
- Step 4** Switch to the **Resource Management Plans** tab page and click **Start/Stop** to start or stop the plan.

NOTICE

- Only one plan can be started for each cluster.
 - A plan must have at least two stages before it can be started.
-

----End

Viewing the Execution Logs of a Resource Management Plan

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Clusters**. Click the name of a cluster.
- Step 3** Go to the **Basic Information** page and click the **Resource Management** tab in the navigation pane on the left.
- Step 4** Switch to the **Resource Management Plans** tab page. In the plan execution logs area, click **View** to view the plan execution logs.

----End

Deleting a Resource Management Plan

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Clusters**. Click the name of a cluster.
- Step 3** Go to the **Basic Information** page and click the **Resource Management** tab in the navigation pane on the left.
- Step 4** Switch to the **Resource Management Plans** tab page and click **Delete** to delete the current resource management plan.

NOTICE

You cannot delete a running resource management plan.

----End

15.3.2 Managing Resource Management Plan Stages

Prerequisites

The following conditions must be met when you add or modify a resource management plan:

- The total CPU share of all resource pools does not exceed 99%.
- The total CPU limit of all resource pools does not exceed 100%.

NOTE

- The CPU usage limit can be configured only in 8.1.3 and later versions.
- The default start time is the UTC time. The next execution time is your local time.

Adding a Resource Management Plan Stage

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Choose **Clusters**. Click the name of a cluster.

Step 3 Go to the **Basic Information** page and click the **Resource Management** tab in the navigation pane on the left.

Step 4 Go to the plan details page and click **Add** in the **Plan stage** area. On the **Add Stage** page, enter the stage name and configure the resource information. Confirm the configuration and click **OK**.

NOTICE

- Stages cannot be added to a running resource management plan.
- You can add a maximum of 48 stages for each plan.
- The switchover time of all phases in a plan cannot be the same.
- Configure the time, date, and month. Do not set an invalid date, for example, February 30.

Add Stage [X]

* Stage

* Month

1	2	3	4	5	6	7	8	9	10
11	12								
All									

* Day

1	2	3	4	5	6	7	8	9	10
11	12	13	14	15	16	17	18	19	20
21	22	23	24	25	26	27	28	29	30
31									
All									

* Start Time (UTC)
Note: The UTC time is used by default. Set the policy based on the time zone and time difference as required [Learn more](#)

Selected Resource Pools [Q]

<input type="checkbox"/>	Resource Pool Name	Share	Limit	Memory (%)	Concurrency	Simple State...
--------------------------	--------------------	-------	-------	------------	-------------	-----------------

[OK] [Cancel]

----End

Modifying a Resource Management Plan Stage

Step 1 Log in to the GaussDB(DWS) console.

- Step 2** Choose **Clusters**. Click the name of a cluster.
- Step 3** Go to the **Basic Information** page and click the **Resource Management** tab in the navigation pane on the left.
- Step 4** Switch to the **Resource Management Plans** tab page and click **Modify** in the **Operation** column of the plan stage.
- Step 5** Modify parameters, such as the stage changing time and resource configurations.

Modify Stage stage1
×

11	12	13	14	15	16	17	18	19	20
21	22	23	24	25	26	27	28	29	30
31									

All

★ Start Time (UTC)

Note: The UTC time is used by default. Set the policy based on the time zone.

Selected Resource Pools Enter a resource pool name. Q

<input type="checkbox"/> Resource Po...	Shared Quota	Limit	Memory (%)	Concurrency	Simple State...	Network band...
<input type="checkbox"/> pool1	<input style="width: 40px;" type="text" value="20"/>	<input style="width: 40px;" type="text" value="0"/>	<input style="width: 40px;" type="text" value="0"/>	<input style="width: 40px;" type="text" value="10"/>	<input style="width: 40px;" type="text" value="-1"/>	<input style="width: 40px;" type="text" value="4"/>

Available Queues Enter a resource pool name. Q

<input type="checkbox"/> Resource Po...	Shared Quota	Limit	Memory (%)	Concurrency	Simple State...	Network band...
<input type="checkbox"/> pool2	24	0	0	10	-1	-1

NOTE

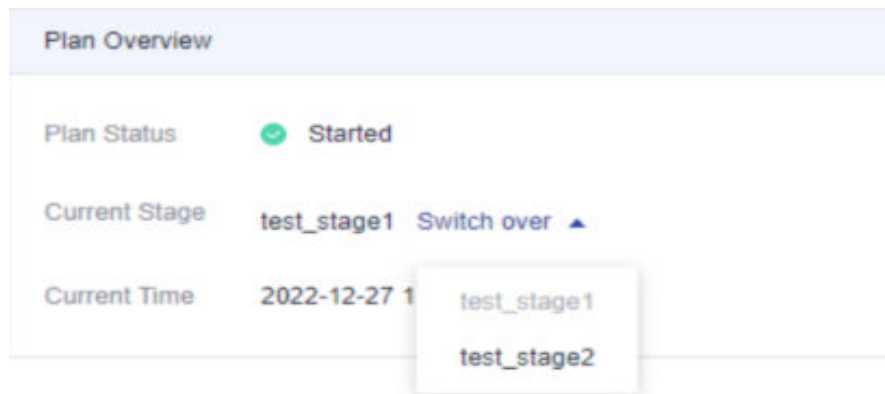
Only clusters of the version 8.2.1 and later support the network bandwidth weight.

----End

Manually Changing the Resource Management Plan Stage

If a running plan needs to be switched to a stage in advance, you can manually do it.

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Clusters**. Click the name of a cluster.
- Step 3** Go to the **Basic Information** page and click the **Resource Management** tab in the navigation pane on the left.
- Step 4** Switch to the **Resource Management Plans** tab page, click the switch button in the plan overview area, and select the target stage.



----End

Importing/Exporting Resource Management Plan Stages

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Clusters**. Click the name of a cluster.
- Step 3** Go to the **Basic Information** page and click the **Resource Management** tab in the navigation pane on the left.
- Step 4** Switch to the **Resource Management Plans** tab page. In the plan stages area, click **Import/Export** to import or export a resource management plan stage.

NOTICE

- Configurations cannot be imported to a running resource management plan.
 - Ensure there is a resource pool before import.
-

----End

Deleting a Resource Management Plan Stage

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Clusters**. Click the name of a cluster.
- Step 3** Go to the **Basic Information** page and click the **Resource Management** tab in the navigation pane on the left.
- Step 4** Switch to the **Resource Management Plans** tab page and click **Delete** in the **Operation** column of the plan stage.

----End

NOTICE

Stages in a running resource management plan cannot be deleted.

15.4 Workspace Management

Overview

Your cluster may run out of space if disk usage is not controlled, resulting in cluster exceptions and service interruption. Once disks are full, it takes long and huge efforts to recover workloads. Setting a database to read-only can reduce disk usage, but it also interrupts services. To solve this problem, GaussDB(DWS) provides multi-dimensional storage management. You can limit the permanent space that can be occupied by a schema; and can limit the usage of permanent space, temporary space, and operator space for a user.

- Schema level: Schema space management allows you to query database and schema space information in a cluster and modify the total schema space.
- User level: User space management allows you to limit users' space usage, preventing task execution from being blocked due to insufficient storage space. When you create a user in GaussDB(DWS), you can specify the space available to the user. The following types of storage space can be managed:
 - Permanent space (**PREM SPACE**)
Space occupied by permanent tables (non-temporary tables) created by users
 - Temporary space (**TEMP SPACE**)
Space occupied by temporary tables created by users
 - Operator spill space (**SPILL SPACE**)
During query execution, if the actual memory usage is greater than estimated, the query may be spilled to disks. The storage space occupied in this case is called operator spill space. You can control a user's operator spill space usage during query execution.

NOTE

- This feature is supported only in cluster version 8.1.1 or later.
- Currently, the GaussDB(DWS) management plane only supports schema space management.

Procedure

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Choose **Clusters**. Click the name of a cluster.
- Step 3** Go to the **Basic Information** page and click the **Resource Management** tab in the navigation pane on the left.
- Step 4** On the **Schema Space Manage** page, select a database.
- Step 5** In the row where the scheme to be edited resides, click **Edit** and modify the space limit.

Schema Name	Used Space	Space Allowance	Skew Percent	Max Value	Min Value	Operation
gs_logical_cluster	0	Unlimited	0	0	0	Edit
public	0	Unlimited	0	0	0	Edit
scheduler	0.02 MB	Unlimited	0	0.01 MB	0.01 MB	Edit

Step 6 Click **OK**.

NOTE

- The space quota limits only common users but not database administrators. Therefore, when the used space is equal to the space limit, the actual used space may exceed the specified value.
- Quota for a single DN = Total quota/Number of DNs. Therefore, the configured value may be slightly different from the displayed value.

----End

15.5 Exception Rules

Feature Description

Some complex statements may consume a large number of resources for computing, leading to performance deterioration of the entire GaussDB(DWS) database. To maintain system stability, GaussDB(DWS) allows you to customize exception rules, and terminate/downgrade the tasks that hit the rules. You can use SQL syntax to configure exception rules based on your resource and workload conditions, and associate the rules with resource pools. If there are no user-defined rules, two default exception rules will take effect.

Figure 15-1 Exception rules

Rule	Associated Resource Pool	Operation Upon Exception	Operation
default_cpu_percent		Terminated	Edit Delete
default_memory		Terminated	Edit Delete

NOTICE

- This feature is supported only by clusters of version 8.2.0 or later. The exception rule "Maximum Bandwidth on a Single DN" is supported only by clusters of version 8.2.1 or later.
- The cluster version 8.2.1 supports downgradation of exception rules. All exception rules support downgradation behaviors. After downgradation, only network resource preemption is downgraded to a low priority. Downgraded network queries are scheduled only when there is no normal queries.
- A resource pool can be associated with multiple groups of exception rules, which work in an OR way. One group of exception rules works if all its conditions are met. For example, a resource pool is associated with two groups of rules. One group specifies **elapsedtime=2400**, and the other group specifies **elapsedtime=1200** and **memsize=2000**. If the execution time of a job reaches 1200 seconds and the memory usage reaches 2,000 MB, or if the execution time reaches 2,400 seconds, the job will be terminated.
- Rules in the same exception rule group take effect only if all the conditions are met. For example, if you set **elapsedtime=1000** and **memsize=500**, it indicates that a job is terminated only if its execution time reaches 1,000 seconds and its memory usage reaches 500 MB. If only one of the thresholds is reached, the job will not be terminated.
- The default exception rules are supported only by clusters of version 8.2.0 or later. After a cluster of an earlier version is upgraded to version 8.2.0 or later, the default exception rules do not take effect. You can create exception rules as needed.
- The default exception rules take effect for users not associated with any resource pools, and for users whose resource pools do not have any exception rules configured. If a user-defined rule is associated with a resource pool, this rule prevails in the pool.

User-defined Exception Rules and Default Exception Rules

The following table describes the user-defined exception rules and default exception rules supported by the current GaussDB (DWS) version.

Table 15-5 User-defined exception rules

Exception Threshold Type	Description	Value Range (-1 disables a parameter. 0 is not supported.)	Operation upon Exception
Blocking Time	Job blocking duration, in seconds. The time includes the total time spent in global and local concurrent queuing. The queuing time of each substatement (if any) in a statement is also counted.	-1 or 1 to INT64_MAX-1	Terminate / Downgrade

Exception Threshold Type	Description	Value Range (-1 disables a parameter. 0 is not supported.)	Operation upon Exception
Execution Time	Execution duration of a job, in seconds. The time indicates the duration from the start point of execution to the current time point. The execution time of each substatement (if any) in a statement is also counted.	-1 or 1 to INT64_MAX-1	Terminate / Downgrade
Total CPU time on all DNs.	Total CPU time spent in executing a job on all DNs, in seconds.	-1 or 1 to INT64_MAX-1	Terminate / Downgrade
Total CPU Time Skew Rate on All DNs	CPU time skew of a job executed on DNs. The value depends on the setting of elapsedtime . The system starts to check the CPU time skew of a job every 5 seconds after the job execution time reaches elapsedtime .	-1, or 1 to 100	Terminate / Downgrade
Average CPU Usage Per DN	Average CPU usage of a job executed across all DNs.	-1, or 1 to 100	Terminate / Downgrade
Data Spilled to Disk Per DN	Allowed maximum job data spilled to disks on a DN. The unit is MB.	-1 or 1 to INT64_MAX-1	Terminate / Downgrade
Maximum Bandwidth on a Single DN	Maximum network bandwidth (MB) for a job on a single DN.	-1 or 1 to INT64_MAX-1	Terminate / Downgrade

Table 15-6 Default exception rules

Rule Name	Description	Operation upon Exception
default_cpu_percent	This rule is triggered if multiple jobs are running in a cluster, and the CPU usage of a resource pool reaches 90%. (If no resource pools are configured, the total CPU usage of the cluster is checked). This rule terminates the job whose execution time reached 15 minutes and average CPU usage exceeded 50%.	Terminate
default_spill_size	This rule is triggered if the size of data spilled to disk on a single DN reaches 1/10 of the instance space during job execution in the cluster.	Terminate

Creating an Exception Rule

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Click the name of the target cluster. The **Basic Information** page is displayed.
- Step 3** In the navigation pane on the left, choose **Resource Management** and switch to the **Exception Rules** tab page.
- Step 4** Click **Add** to add an exception rule.
- Step 5** Click **OK**.

NOTE

- After an exception rule is created, it does not take effect immediately. You need to associate it to a resource pool. For details, see [1. Associate exception rules.](#)
- The cluster version 8.2.1 supports downgradation of exception rules. All exception rules support downgradation behaviors. After downgradation, only network resource preemption is downgraded to a low priority. Downgraded network queries are scheduled only when there are no normal queries.

----End

Editing an Exception Rule

- Step 1** Log in to the GaussDB(DWS) console.
- Step 2** Click the name of the target cluster. The **Basic Information** page is displayed.
- Step 3** In the navigation pane on the left, choose **Resource Management** and switch to the **Exception Rules** tab page.
- Step 4** Locate the row that contains the target exception rule and click **Edit** in the **Operation** column to edit the exception rule.

 **NOTE**

- When editing an exception rule, if you want to delete an exception rule threshold, clear the value or set it to **-1**.
- If the exception threshold is changed during job execution, the new threshold will take effect for the statement being executed.

----End

Deleting an Exception Rule

Step 1 Log in to the GaussDB(DWS) console.

Step 2 Click the name of the target cluster. The **Basic Information** page is displayed.

Step 3 In the navigation pane on the left, choose **Resource Management** and switch to the **Exception Rules** tab page.

Step 4 Locate the row that contains the target exception rule and click **Delete** in the **Operation** column to delete the rule.

 **NOTE**

If an exception rule has been associated to a resource pool, the exception rule cannot be deleted. You need to disassociate the exception rule from the resource pool before deleting it.

Step 5 Click **OK**.

----End

16 Data Source Management

16.1 MRS Data Sources

16.1.1 MRS Data Source Usage Overview

MRS Cluster Overview

MRS is a big data cluster running based on the open-source Hadoop ecosystem. It provides the industry's latest cutting-edge storage and analysis capabilities of massive volumes of data, satisfying your data storage and processing requirements.

You can use Hive/Spark (analysis cluster of MRS) to store massive volumes of service data. Hive/Spark data files are stored in HDFS. On GaussDB(DWS), you can connect a data warehouse cluster to MRS clusters, read data from HDFS files, and write the data to GaussDB(DWS) when the clusters are on the same network.

NOTE

Currently, the hybrid data warehouse (standalone mode) cannot import data from MRS.

Operation Process

Perform the following operations to import data from MRS to a data warehouse cluster:

1. Prerequisites
 - a. Create an MRS cluster. For details, see [Buying a Custom Cluster](#).
 - b. Create an HDFS foreign table for querying data from the MRS cluster over APIs of a foreign server.

For details, see [Importing Data from MRS to a Data Warehouse Cluster](#) in *Data Warehouse Service (DWS) Data Migration and Synchronization*.

 NOTE

- Multiple MRS data sources can exist on the same network, but one GaussDB(DWS) cluster can connect to only one MRS cluster at a time.
2. In the data warehouse cluster, create an MRS data source connection according to [Creating an MRS Data Source Connection](#).
 3. Import data from an MRS data source to the cluster. For details, see [Importing Data from MRS to a Cluster](#).
 4. (Optional) When the HDFS configuration of the MRS cluster changes, update the MRS data source configuration on GaussDB(DWS). For details, see [Updating the MRS Data Source Configuration](#).

16.1.2 Creating an MRS Data Source Connection

Scenario

Before GaussDB(DWS) reads data from MRS HDFS, you need to create an MRS data source connection that functions as a channel of transporting data warehouse cluster data and MRS cluster data.

Impact on the System

- You can create only one MRS data source connection in the data warehouse cluster at a time.
- When an MRS data source connection is being created, the system automatically adds inbound and outbound rules to security groups of the data warehouse cluster and MRS cluster. Nodes in the same subnet can be accessed.
- For the MRS cluster with Kerberos authentication enabled, the system automatically adds a **Machine-Machine** user that belongs to user group **supergroup** to the MRS cluster.

Prerequisites

- You have created a data warehouse cluster and recorded the VPC and subnet where the cluster resides.
- An MRS cluster of the analysis type has been created.

Procedure

Step 1 Log in to the Huawei Cloud management console.

Step 2 Choose **Service List > Analytics > MapReduce Service** to enter the MRS management console and create a cluster.

Configure parameters as required. For details, see "Cluster Operation Guide > Custom Creation of a Cluster" in the *MapReduce Service User Guide*.

- The VPC of the MRS cluster must be the same as that of the data warehouse cluster.
- MRS cluster versions 1.9.2, 2.1.0, 3.0.2-LTS, and 3.1.2-LTS are recommended.

 **NOTE**

- For clusters of version 8.1.1.300 and later, MRS clusters support versions 1.6.*; 1.7.*; 1.8.*; 1.9.*; 2.0.*; 3.0.*; 3.1.*; and later (* indicates a number).
 - For clusters earlier than version 8.1.1.300, MRS clusters support versions 1.6.*; 1.7.*; 1.8.*; 1.9.*; and 2.0.* (* indicates a number).
- Select the Hadoop component.

If you already have a qualified MRS cluster, skip this step.

Step 3 Choose **Service List > Analytics > GaussDB(DWS)**.

Step 4 On the GaussDB(DWS) console, choose **Cluster > Dedicated Clusters**.

Step 5 In the cluster list, click the name of a cluster. The **Cluster Information** page is displayed.

Step 6 In the navigation tree on the left, choose **Data Sources > MRS Data Sources**.

Step 7 Click **Create MRS Cluster Connection** and configure parameters.

Table 16-1 MRS common connection parameters

Parameter	Description
Data Source	GaussDB(DWS) database server name. It can contain 3 to 63 characters, including lowercase letters, numbers, and underscores (_), and must start with a lowercase letter.
Configuration Mode	The way in which the system obtains files. The options are as follows: <ul style="list-style-type: none">• MRS Account: Configure the username and password of the Manager of the MRS cluster. The system will log in to the Manager and automatically download configuration and verification files. For more information, see Table 16-2.• File upload: Download the configuration file from the Manager of the MRS cluster and manually upload it. You can use this method for Kerberos authentication. For more information, see Table 16-3. NOTE If you select File upload , ensure that MRS can communicate with the GaussDB(DWS) cluster.
Database	Database where the data source is located.
Description	Description of the connection.

Table 16-2 Parameters of the MRS Account mode

Parameter	Description
MRS Data Source	<p>Select an MRS cluster that can be connected to GaussDB(DWS) from the drop-down list box. By default, the custom, hybrid, and analytical MRS clusters that are in the same VPC and subnet as the current GaussDB(DWS) cluster and available to the current user are displayed.</p> <p>After you select an MRS cluster, the system automatically displays whether Kerberos authentication is enabled for the selected cluster. Click View MRS Cluster to view its detailed information.</p> <p>If the MRS Data Source drop-down list is empty, click Create MRS Cluster to create an MRS cluster.</p>
MRS Account	Account used when a GaussDB(DWS) cluster connects to an MRS cluster.
Password	<p>Password of the connection user. If you change the password, you need to create a connection again.</p> <p>NOTICE</p> <p>Ensure the account has been used for logging in to MRS Manager. If you use a new account, you will be asked to change your password when you first log in. In this case, the MRS data source will fail to be configured.</p>
Use a Machine-Machine Account	<p>Creates a machine-machine account named dws in MRS and uses it for interaction with MRS. This account is in the supergroup group and has all permissions. If the switch is toggled off, the configured man-machine account will be used. Ensure this account has the permission to access data, or a message will be displayed during data source access, indicating the required file does not exist.</p>

Table 16-3 Parameters of the File upload mode

Parameter	Description
Authentication Credential	<p>Keytab file of a user A credential file downloaded from Manager of the MRS cluster. File name format: Username_Timestamp_keytab.tar</p> <ul style="list-style-type: none">• For MRS 2.x or earlier, choose System > Manage User. In the Operation column of a user, choose More > Download authentication credential.• For MRS 3.x or later, choose System > Permission > User. In the Operation column of a user, choose More > Download Authentication Credential.

Parameter	Description
Client Profile	<p>Client configuration files of HDFS, Hive, and hosts. When downloading the client, set Select Client Type to Configuration Files Only.</p> <ul style="list-style-type: none">• For MRS 2.x or earlier, choose Services and click Download Client.• For MRS 3.x or later, choose Homepage. Click the More icon and choose Download Client.

Step 8 Click **OK** to save the connection.

Configuration Status turns to **Creating**. You can view the connection that is successfully created in the MRS data source list and the connection status is **Available**.

 **NOTE**

- In the **Operation** column, you can click **Update Configurations** to update **MRS Cluster Status** and **Configuration Status**. During configuration update, you cannot create a connection. The system checks whether the security group rule is correct. If the rule is incorrect, the system rectifies the fault. For details, see [Updating the MRS Data Source Configuration](#).
- In the **Operation** column, you can click **Delete** to delete the unnecessary connection. When deleting a connection, you need to manually delete the security group rule.
- If the security group rules are not deleted, nodes in the data warehouse cluster can still communicate with nodes in the MRS cluster. If you have strict requirements on network security, manually delete the rules.

----End

16.1.3 Updating the MRS Data Source Configuration

Scenario

For MRS, if the following parameter configurations of the HDFS cluster change, data may fail to be imported to the data warehouse cluster from the HDFS cluster. Before importing data using the HDFS cluster, you must update the MRS data source configuration.

Prerequisites

You have created an MRS data source connection for the data warehouse cluster.

Impact on the System

When you are updating an MRS data source connection, the data warehouse cluster will automatically restart and cannot provide services.

Procedure

Step 1 On the GaussDB(DWS) console, choose **Clusters** > **Dedicated Clusters**.

- Step 2** In the cluster list, click the name of a cluster. On the page that is displayed, click **MRS Data Sources**.
- Step 3** In the MRS data source list, select the MRS data source that you want to update. In the **Operation** column, click **Update Configurations**.

MRS Cluster Status and **Configuration Status** of the current connection will be updated. During configuration update, you cannot create a connection. The system checks whether the security group rule is correct. If the rule is incorrect, the system rectifies the fault. The following table describes the parameters.

Table 16-4 Parameter description

Parameter	Description
dfs.client.read.shortcircuit	Specifies whether to enable the local read function.
dfs.client.read.shortcircuit.skip.checksum	Specifies whether to skip data verification during the local read.
dfs.client.block.write.replace-datanode-on-failure.enable	Specifies whether to replace the location storing copies with the new node when data blocks fail to be written to HDFS.
dfs.encrypt.data.transfer	Specifies whether to enable data encryption. The value true indicates that the channels are encrypted. The channels are not encrypted by default. NOTE <ul style="list-style-type: none">This parameter is available only for clusters with Kerberos authentication enabled.This parameter is valid only when hadoop.rpc.protection is set to privacy.
dfs.encrypt.data.transfer.algorithm	Specifies the encryption and decryption algorithm for key transmission. This parameter is valid only when dfs.encrypt.data.transfer is set to true . The default value is 3des , indicating that the 3DES algorithm is used for encryption.
dfs.encrypt.data.transfer.cipher.suites	Specifies the encryption and decryption algorithm for the transmission of actually stored data. If this parameter is not specified, the cryptographic algorithm specified by dfs.encrypt.data.transfer.algorithm is used for data encryption. The default value is AES/CTR/NoPadding .
dfs.replication	Specifies the default number of data copies.
dfs.blocksize	Specifies the default size of a data block.

Parameter	Description
hadoop.security.authentication	Specifies the security authentication mode.
hadoop.rpc.protection	Specifies the RPC communication protection mode. Default value: <ul style="list-style-type: none">• Security mode (Kerberos authentication enabled): privacy• Common mode (Kerberos authentication disabled): authentication NOTE <ul style="list-style-type: none">• authentication: indicates that only authentication is required.• integrity: indicates that authentication and consistency check need to be performed.• privacy: indicates that authentication, consistency check, and encryption need to be performed.
dfs.domain.socket.path	Specifies the locally used Domain socket path.

----End

16.2 Managing OBS Data Sources

GaussDB(DWS) allows you to access data on OBS by using an agency. You can create a GaussDB(DWS) agency, grant the OBS OperateAccess or OBS Administrator permission to the agency, and bind the agency to an OBS data source you created. In this way, you can access data on OBS by using OBS foreign tables.

NOTE

- This feature is supported only in 8.2.0 or later.
- For the OBS data source of a cluster, only one of the creation, modification, and deletion operations can be performed at a time.

Creating an OBS Agency

Scenario

Before creating an OBS data source, create an agency that grants GaussDB(DWS) the OBS OperateAccess or OBS Administrator permission.

Procedure

- Step 1** Click your account in the upper right corner of the page and choose **Identity and Access Management**.
- Step 2** In the navigation pane on the left, choose **Agency**. In the upper right corner, click **Create Agency**.

- Step 3** Select **Cloud Service** and set **Cloud Service** to **DWS**.
 - Step 4** Click **Next** to grant the OBS OperateAccess or OBS Administrator permission to the agency.
 - Step 5** Click **Next**. Select **All resources** or specific resources, confirm the information, and click **Submit**.
- End

Creating an OBS Data Source

Prerequisites

An agency has been created to grant GaussDB(DWS) the OBS OperateAccess permission.

Procedure

- Step 1** On the GaussDB(DWS) console, choose **Clusters** > **Dedicated Clusters**.
- Step 2** In the cluster list, click the name of a cluster. On the page that is displayed, choose **Data Sources** > **OBS Data Source**.
- Step 3** Click **Create OBS Cluster Connection** and configure parameters.

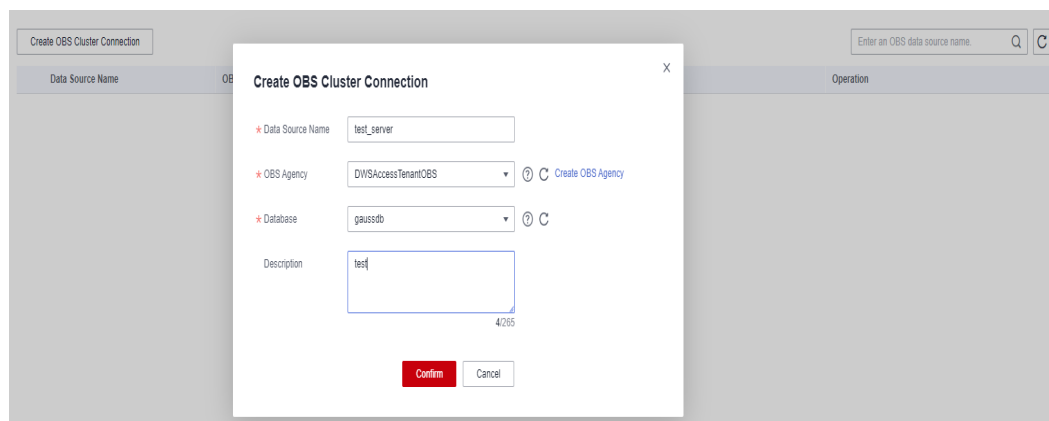


Table 16-5 OBS data source connection parameters

Parameter	Description
Data Source Name	Name of the OBS data source connection to be created. You can assign a personalized value to this parameter. The data source name is used as the server name specified in the statement for creating an OBS foreign table.
OBS Agency	Agency with the OBS OperateAccess permission to be granted to GaussDB(DWS)
Database	Database where the OBS data source connection is to be created

Parameter	Description
Description	Description about the OBS data source connection

Step 4 Confirm the settings and click **OK**. The creation takes about 10 seconds.

----End

Updating the OBS Data Source Configuration

Scenario

After an OBS data source connection is created, GaussDB(DWS) periodically updates the temporary agency information used by the data source. If the automatic update fails for 24 hours, the data source connection will be unavailable. To solve this problem, manually update the information on the console.

Procedure

Step 1 On the GaussDB(DWS) console, choose **Clusters > Dedicated Clusters**.

Step 2 In the cluster list, click the name of a cluster. On the page that is displayed, choose **Data Sources > OBS Data Source**.

Step 3 In the **Operation** column of an OBS data source, click **Update Configuration**.

Step 4 Confirm the settings and click **OK**. The update takes about 10 seconds.

----End

Changing the OBS Data Source Agency

Scenario

You can change the agency bound to the OBS data source.

Procedure

Step 1 On the GaussDB(DWS) console, choose **Clusters > Dedicated Clusters**.

Step 2 In the cluster list, click the name of a cluster. On the page that is displayed, choose **Data Sources > OBS Data Source**.

Step 3 In the **Operation** column of a data source, click **Manage Agency**. In the dialog box that is displayed, select a new agency.

Step 4 Confirm the settings and click **OK**. The change takes about 10 seconds.

----End

Deleting an OBS Data Source

Step 1 On the GaussDB(DWS) console, choose **Clusters > Dedicated Clusters**.

Step 2 In the cluster list, click the name of a cluster. On the page that is displayed, choose **Data Sources > OBS Data Source**.

Step 3 In the **Operation** column of an OBS data source, click **Delete**.

Step 4 Confirm the settings and click **OK**. The deletion takes about 10 seconds.

----End

Using an OBS Data Source

GaussDB(DWS) uses foreign tables to access data on OBS. The **SERVER** parameters specified for accesses with and without an agency are different.

If you access OBS without an agency, the **SERVER** provided on the console contains parameters **access_key** and **secret_access_key**, which are the AK and SK of the OBS access protocol, respectively.

If you access OBS with an agency, the **SERVER** provided on the console contains the **access_key**, **secret_access_key**, and **security_token** parameters, which are the temporary AK, temporary SK, and the **SecurityToken** value of the temporary security credential in IAM, respectively.

After the OBS agency and OBS data source are created, you can obtain the **SERVER** information, for example, **obs_server**, on the console. The way users create and use foreign tables with an agency is the same as the way they do without an agency. For how to use the OBS data source, see [Importing Data from OBS](#).

The following example shows how common user **jim** reads data from OBS through a foreign table.

1. Repeat the preceding steps to create an OBS data source named **obs_server**.
2. Connect to the database as system administrator dbadmin, create a common user, and grant the common user the permission to use OBS servers and OBS foreign tables. Replace **{Password}** with the actual password and **obs_server** with the actual OBS data source name.

```
CREATE USER jim PASSWORD '{Password}';
ALTER USER jim USEFT;
GRANT USAGE ON FOREIGN SERVER obs_server TO jim;
```
3. Connect to the database as common user **jim** and create an OBS foreign table **customer_address** that does not contain partition columns.

In the following command, replace **obs_server** with the name of the created OBS data source. Replace **/user/obs/region_orc11_64stripe1/** with the actual OBS directory for storing data files. **user** indicates the OBS bucket name.

```
CREATE FOREIGN TABLE customer_address
(
  ca_address_sk          integer          not null,
  ca_address_id         char(16)          not null,
  ca_street_number      char(10)          ,
  ca_street_name        varchar(60)       ,
  ca_street_type        char(15)          ,
  ca_suite_number       char(10)          ,
  ca_city               varchar(60)       ,
  ca_county             varchar(30)       ,
  ca_state              char(2)           ,
  ca_zip               char(10)           ,
  ca_country            varchar(20)       ,
  ca_gmt_offset         decimal(36,33)    ,
  ca_location_type      char(20)         ,
)
SERVER obs_server OPTIONS (
  FOLDERNAME '/user/obs/region_orc11_64stripe1/');
```

```
FORMAT 'ORC',  
ENCODING 'utf8',  
TOTALROWS '20'  
)  
DISTRIBUTE BY roundrobin;
```

4. Query data stored in OBS by using a foreign table.

```
SELECT COUNT(*) FROM customer_address;  
count  
-----  
20  
(1row)
```

16.3 Managing LakeFormation Data Sources

On the GaussDB(DWS) console, you can create a LakeFormation data source on the console to access metadata on LakeFormation.

NOTE

- This feature is for limited commercial use only. It only supports DWS 3.0 clusters in version 9.0.1 and later, and DWS 2.0 clusters in version 8.2.1.300 and later.
- LakeFormation is connected through VPC Endpoint. When you create a LakeFormation data source, if there is no VPC endpoint for your GaussDB(DWS) cluster, it will be automatically created. VPC endpoints incur extra fees.

Prerequisites

- A LakeFormation instance is available. For details, see section "Creating a LakeFormation Instance" in *LakeFormation Usage Guide*.
- Create an agency with LakeFormation permissions (including the minimum permissions). For details about how to configure the permissions, see section "Data Permission Authorization" in *LakeFormation Usage Guide*. If the permissions are not correctly configured, an error will be reported.
- For an IAM user to use GaussDB(DWS) to call APIs on the LakeFormation management plane, the user must have LakeFormation permissions (at least **lakeformation:instance:access** and **lakeformation:instance:describe**).

Creating a LakeFormation Data Source

- Step 1** On the GaussDB(DWS) console, choose **Clusters > Dedicated Clusters**.
- Step 2** In the cluster list, click the name of a cluster. On the page that is displayed, choose **Data Sources > LakeFormation Data Sources**.
- Step 3** Click **Create LakeFormation Data Source Connection** and configure parameters.

Table 16-6 LakeFormation data source connection parameters

Parameter	Description
Data Source	Name of the LakeFormation data source connection to be created
LakeFormation Instance	LakeFormation cluster instance to be bound

Parameter	Description
Database	Database where the LakeFormation data source connection is to be created
Agency	An agency authorized by LakeFormation. GaussDB(DWS) interacts with LakeFormation through this agency token o obtain metadata.
Description	Description about the LakeFormation data source connection

Step 4 Confirm the settings and click **OK**. The creation takes about 1 minute.

----End

Updating a Configuration

Scenario

- After a data source connection is created, its VPC endpoint is deleted by mistake and the data source cannot be used.
- The agency needs to be changed.
- A token fails to be updated. After this issue is fixed, the token needs to be updated immediately.

Procedure

Step 1 On the GaussDB(DWS) console, choose **Clusters > Dedicated Clusters**.

Step 2 In the cluster list, click the name of a cluster. On the page that is displayed, choose **Data Sources > LakeFormation Data Sources**.

Step 3 In the **Operation** column of a LakeFormation data source, click **Update Configuration**.

Figure 16-1 Updating a LakeFormation data source connection

Create LakeFormation Data Source Connection

* Data Source

* LakeFormation Instance [Create Instance](#)

* Database

* Agency

Description 0/256

OK Cancel

Step 4 During the update, you can only change the agency. After confirming that the agency is correct, click **OK** to submit the update. The update takes about 1 minute.

----End

Deleting a LakeFormation Data Source

Step 1 On the GaussDB(DWS) console, choose **Clusters > Dedicated Clusters**.

Step 2 In the cluster list, click the name of a cluster. On the page that is displayed, choose **Data Sources > LakeFormation Data Sources**.

Step 3 In the **Operation** column of a LakeFormation data source, click **Delete**.

Step 4 Confirm the settings and click **OK**. The deletion takes about 10 seconds.

----End

Using a LakeFormation Data Source

For details about how to use the LakeFormation data source, see "Data Migration > Data Import > Using LakeFormation to Import Data" in *Data Warehouse Service Developer Guide*

17 Managing Logical Clusters

17.1 Logical Cluster Overview

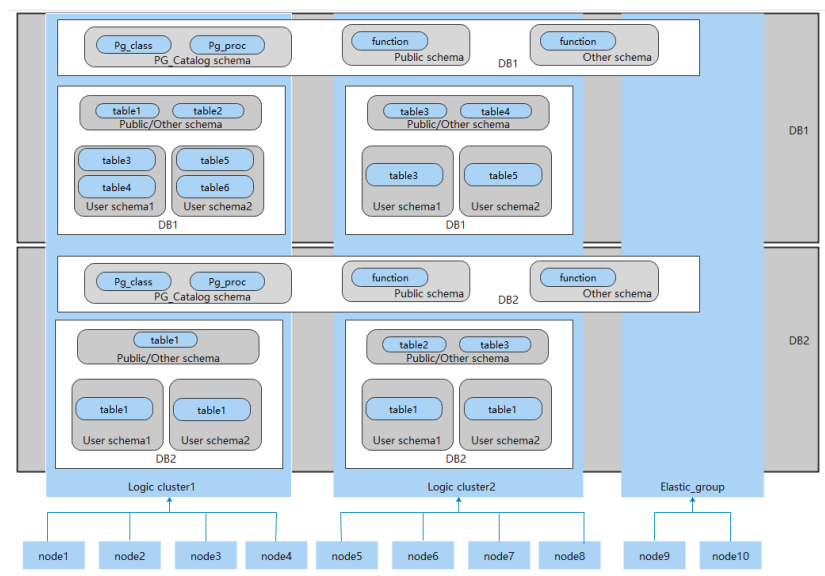
Concepts

A physical cluster can be divided into Node Groups, which are logical clusters. All physical nodes in a physical cluster are divided into multiple logical clusters. A logical cluster is essentially a node group that contains one or more physical nodes. Each physical node belongs to only one logical cluster, and user data tables can only be distributed within the same logical cluster. The data of each logical cluster is isolated from the others. The physical resources allocated to a logical cluster are mainly used for operations on its own data tables, but also for interactive queries with other logical clusters. An enterprise can deploy services on different logical clusters to implement unified service management, and meanwhile isolate the data and resources of services.

Logical clusters are created by dividing nodes of a physical cluster. Tables in a database can be allocated to different physical nodes by logical cluster. A logical cluster can contain tables from multiple databases. [Figure 17-1](#) shows the relationships between logical clusters, databases, and tables.

An elastic cluster is a cluster that always exists in logical cluster mode and consists of nodes that are not part of any logical cluster. It is a special node group that can have multiple or zero DNs. An elastic cluster cannot be manually created. When the first logical cluster is created in a physical cluster, an elastic cluster is also automatically created and all physical nodes not belonging to the logical cluster are automatically added to the elastic cluster. DNs in the elastic cluster will be used for logical clusters created later. To create a logical cluster, ensure that your logical cluster has DNs. (DNs are not required only when you create the first logical cluster in physical cluster mode.) You can add new physical nodes to the elastic cluster through scale-out.

Figure 17-1 Relationships between logical clusters, databases, and tables



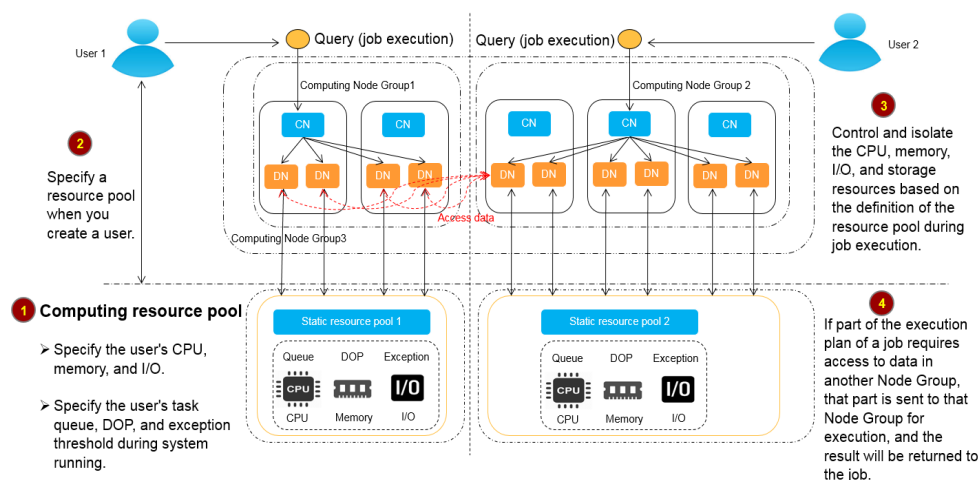
NOTE

- Logical clusters are supported only by clusters of 8.1.0.100 or later.
- You are advised to allocate tables in a database to the same logical cluster.
- A logical cluster is not an independent sub-cluster. It can isolate data, resource, and permissions, but cannot be independently operated or maintained.
- The **Change all specifications** option does not support logical clusters.
- The logical cluster of a DWS 3.0 cluster cannot be switched if the original physical cluster contains data. Ensure the original physical cluster is empty during the switchover.

Logical Cluster Architecture

Figure 17-2 shows the architecture of a physical cluster divided into multiple logical clusters. Nodes in the physical cluster are divided into Node Groups. The jobs of users 1 and 2 are executed in different Node Groups. The two users can define resource pools within their own logical cluster to control resources (CPU, memory, and I/O) used for different jobs. If some jobs of user 1 need to access the data of user 2, they can access data across Node Groups after being authorized. For a logical cluster, you can configure resources accessible across logical clusters to ensure its resources are sufficient.

Figure 17-2 Logical Cluster Architecture



A physical cluster is divided into multiple logical ones. You can define a resource pool for each of them based on service requirements. User tables are not distributed across logical clusters. If services do not access data across logical clusters, they will not compete for resources. Resources can be allocated to jobs in the same logical cluster by using resource pools. If necessary, you can let services access data across logical clusters, and control the resources used for such access to reduce resource competition between jobs within and outside a logical cluster.

After creating a physical cluster, you need to decide whether to divide it into logical clusters. You cannot divide it into logical clusters if you have already created user tables before, because these user tables are distributed on all physical nodes. For more information about the limitations, see [Constraints and Limitations](#). If you want to manage an existing cluster (for example, a database cluster built in a version earlier than 8.1.0.100) as a logical cluster, you can upgrade the cluster to 8.1.0.100 or later and then convert all the nodes in the cluster into a single logical cluster. Then, add nodes to the physical cluster and create another logical cluster on the new nodes.

Operations on logical clusters include:

- **Adding/Deleting a Logical Cluster:**
 - Creating a logical cluster: After converting a physical cluster into a logical cluster, you can group some physical nodes into a logical cluster by specifying the name and the nodes of the logical cluster.
 - Deleting a logical cluster: You can delete a logical cluster with a specified name. After the logical cluster is deleted, the released physical nodes are removed from the physical cluster.
- **Managing Logical Clusters:**
 - Modifying a logical cluster: You can add or remove nodes from a logical cluster as needed.
 - Resource management (logical cluster mode): You can manage resources in a specified logical cluster (supported only by 8.1.3.101 and later versions).
 - Scaling out a logical cluster: This operation increases the number of physical nodes in the logical cluster and redistributes tables in the logical cluster to the new physical nodes.

- Restarting a logical cluster: This operation restarts all DNs in the logical cluster. Considering the impact on the entire physical cluster, the DNs in a logical cluster cannot be stopped or started individually.
- Scaling in a logical cluster: Select and scale in a host ring from an elastic cluster.
- **Scheduling GaussDB(DWS) 3.0 Logical Cluster Creation and Deletion:** Computing logical clusters can be created and deleted during the scheduled period of time to dynamically scale computing resources.

Constraints and Limitations

- The smallest unit of the creation, scale-out, and scale-in of a logical cluster is a ring. A ring consists of at least three hosts, where the primary, standby, and secondary DNs are deployed.
- During the logical cluster switchover, if the original physical cluster has data, the cluster will be locked. You can run simple DML statements, such as adding, deleting, modifying, and querying data. However, running complex DDL statements, such as operating database objects, will block services and report errors. Exercise caution when performing this operation.
- A logical cluster cannot be independently backed up or restored.
- A logical cluster cannot be independently upgraded.
- A physical cluster cannot be rolled back to a physical cluster after it is converted to a logical cluster.
- In logical cluster mode, only logical clusters can be created, and Node Groups cannot be created. In addition, Node Groups cannot be created in a logical cluster.
- O&M operations (creation, deletion, editing, scale-out, scale-in, and restart) of logical clusters cannot be performed concurrently.
- Public database objects (excluding system catalogs, foreign tables, and views) are distributed on all nodes in a physical cluster. After a node of the logical cluster is restarted, the DDL operations performed by other logical clusters on the objects will be interrupted.
- In logical cluster mode, each DN only contains the tables in the logical cluster that the DN belongs to. User-defined functions need to be created on all DNs. Therefore, **%type** cannot be used to reference table field types in the function body.
- In logical cluster mode, the **WITH RECURSIVE** statement cannot be pushed down.
- In logical cluster mode, partitions can be swapped only in the same logical cluster. Partitioned tables and common tables in different logical clusters cannot be swapped.
- In logical cluster mode, if the function parameters or return values contain table types, these table types must belong to the same logical cluster.
- In logical cluster mode, when you create a foreign table using **CREATE TABLE... LIKE**, the source table and the foreign table to be created must be in the same logical cluster.
- In logical cluster mode, tables cannot be created schemas (by using **CREATE SCHEMA... CREATE TABLE** statements). Create a schema, and then create tables in the schema.

- A logical cluster does not support the architecture of one primary node and multiple standby nodes. A logical cluster takes effect only in the architecture of one primary node, one standby node, and one secondary node.
- A logical cluster user cannot access the global temporary tables created by another logical cluster user.

Required permissions on tools

The following describes user permissions for database objects in logical clusters:

- The **CREATE ON NODE GROUP** permission can be granted to any user or role for performing operations such as creating tables in a logical cluster.
 - If the schema specified for a created table is a private schema of a user (that is, the schema has the same name as the user and the owner of the schema is the user), the owner of the created table defaults to the user. You do not need to associate the table with a logical cluster.
 - When a user associated with a logical cluster creates a table, if the **to group** clause is not specified, the table will be created in that logical cluster. The logical cluster associated with the user can be changed.
 - If a user is not associated with any logical cluster, when the user creates a table, the table will be created in the logical cluster specified by **default_storage_nodegroup**. If **default_storage_nodegroup** is set to **installation**, the table will be created in the first logical cluster. In logical cluster mode, the logical cluster with the smallest OID is set as the first logical cluster. If **default_storage_nodegroup** is not set, its value is **installation** by default.
 - DWS 3.0 supports the creation of read-only logical clusters. If a user is associated with a read-only logical cluster, session-level temporary tables (local temporary tables and volatile temporary tables, excluding global temporary tables) can be created only in the read-only logical cluster. If the user creates other common and foreign tables, the tables will be created in the logical cluster specified by **default_storage_nodegroup**. If **default_storage_nodegroup** is set to **installation**, the table will be created in the first logical cluster.
 - The system administrator can run the **ALTER ROLE** command to set **default_storage_nodegroup** for each user. For details about the syntax, see [ALTER ROLE](#).
- Table creation rules
 - If **to group** is not specified for a user table but **default_storage_nodegroup** is set, tables will be created in the specified logical cluster.
 - If **default_storage_nodegroup** is set to **installation**, tables will be created in the first logical cluster, that is, the logical cluster with the smallest OID.
- The owner of a table can be changed to any user. However, you need to check the schema and node group permissions when performing operations on the table.
- A system administrator can be associated with a logical cluster and can create tables in multiple logical clusters.

- If the system administrator is associated with a logical cluster and **to group** is not specified when you create a table, the table will be created in the associated logical cluster by default. If **to group** is specified, the table is created in the specified logical cluster.
- If the system administrator is not associated with a logical cluster and **to group** is not specified, tables are created in the logical cluster of **default_storage_nodegroup**. For details, see the [table creation rules](#).
- System administrator permissions can be granted to a user associated with a logical cluster, but the [table creation rules](#) also apply.
- The logical cluster permission for accessing non-table objects (such as schemas/sequences/functions/triggers) will not be checked.
- A resource pool must be associated with a logical cluster.
 - A logical cluster can be associated with multiple resource pools but a resource pool can be associated with only one logical cluster.
 - Jobs executed by logical cluster users associated with a resource pool can only use resources in the resource pool.
 - You do not need to create a workload group to define the number of concurrent jobs in a logical cluster. Therefore, workload groups are not required for logical clusters.
- When a logical cluster is deleted, only the table, foreign table, and resource pool objects are deleted.
 - Objects dependent on the tables (including the partly dependent sequences/functions/triggers) in the logical cluster will also be deleted.
 - Logical cluster associations with its users and parent-child tenants will be removed during the process. As a result, the users will be associated with the default **installation** node group and with the default global resource pool.
- A logical cluster user can create a database if granted the permission.

Replication Table Node Group

A replication table node group is a special node group in logical cluster mode. It can contain one or more logical clusters, but can only create replication tables. One typical scenario is to create public dimension tables. If multiple logical clusters require some common dimension tables, create a replication table node group and add the common dimension tables to it. The logical clusters contained in the replication table node group can access these dimension tables on the local DNs, with no need to access the tables on other DNs. If a logical cluster is scaled in, the replication table node group will be scaled accordingly. If the logical cluster is deleted, the replication table node group will be scaled in. However, if the replication table node group contains only one logical cluster and the logical cluster is deleted, the replication table node group will also be deleted. In this case, create tables in a logical cluster instead.

Create a replication table node group using the **CREATE NODE GROUP** SQL statement and delete one using **DROP NODE GROUP**. Before deleting a replication table node group, delete all table objects in the node group.


NOTE

Creation of replication table node groups is supported in 8.1.2 or later.

Application Scenarios

Scenario 1: Isolating data with different resource requirements

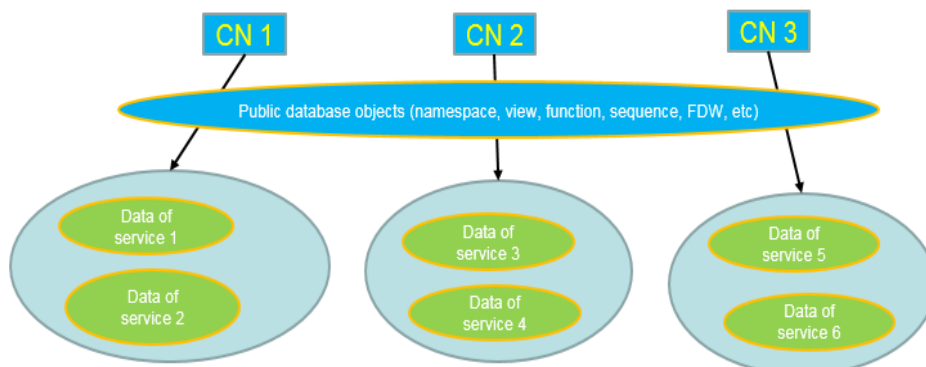
Figure 17-3 Logical cluster division based on resource requirements



As shown in the preceding figure, data with different resource requirements is stored in different logical clusters, and different logical clusters also support mutual access. This ensures that functions are not affected while resources are isolated.

- Tables T1 and T2 are used to calculate a large amount of data and generate report data (for example, bank batch processing). This process involves large batch import and big data query, which consume a lot of memory and I/O resources of nodes and take a long time. However, such a query does not require high real-time performance. Therefore, the data of these two tables can be separated into a different logical cluster.
- Tables T3 and T4 contain some computing data and real-time data, which are mainly used for service point query and real-time query. These queries need high real-time performance. To prevent the interference of other high-load operations, the data of these two tables can be separated into a different logical cluster.
- Tables T5 and T6 are mainly used for OLTP operations with high concurrency. Data in these tables is frequently updated and sensitive to I/O. To prevent the impact of big data query on I/O, the data of these two tables can be separated into a different logical cluster.

Scenario 2: Isolating data for different services and enhancing the multi-tenancy of a data cluster

Figure 17-4 Logical cluster-based multi-service data and multi-tenant management

A large database cluster often stores data for various services. Each service has its own data tables. To allocate resources for different services, you can create multiple tenants. Specifically, assign different service users to different tenants to minimize resource contention among services. As the service scale grows continuously, the number of services in the cluster system also increases. Creating multiple tenants becomes less effective in controlling resource competition. Since each table is distributed across all DN's of a database cluster, every data table operation may involve all DN's, which increases network load and system resource consumption. Simply scaling up the cluster is not enough to solve this problem. Therefore, multiple logical clusters can be created to handle the increasing number of services, as shown in the figure above.

You can create a separate logical cluster and assign new services to it. This way, new services have little impact on existing services. Also, if the service scale in existing logical clusters grows, you can scale out the existing logical clusters.

NOTE

A logical cluster is not suitable for managing multiple independent database systems. An independent database system requires independent O&M and needs to be managed, monitored, backed up, and upgraded separately. Moreover, faults must be isolated between clusters. Logical clusters cannot achieve independent O&M and complete fault isolation.

17.2 Adding/Deleting a Logical Cluster

Adding a Logical Cluster

Precautions

- If you access the **Logical Clusters** page for the first time, the metadata of the logical cluster created at the backend is synchronized to the frontend. After the synchronization is complete, you can view information about the logical clusters at the frontend. The logical cluster name is case sensitive. For example, metadata of **lc1** and **LC1** cannot be synchronized.
- The original resource pool configuration is cleared when the cluster is converted from physical to logical. The resource pool information configured

after the cluster is converted to a logical cluster will be bound to the logical cluster.

Procedure

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane, choose **Clusters > Dedicated Clusters**.
- Step 2** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 3** Enable **Logical Clusters**. The **Logical Clusters** menu item will be displayed in the navigation pane on the left.
- Step 4** Go to the **Logical Clusters** tab and click **Add Logical Cluster**.
- Step 5** Move the ring you want to add from the right to the left panel, enter the logical cluster name, and click **OK**.

----End

Deleting Logical Clusters

Precautions

- The first manually added logical cluster cannot be deleted.
- Nodes of the deleted logical cluster are added to the elastic cluster.

Procedure

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane, choose **Clusters > Dedicated Clusters**.
- Step 2** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 3** In the navigation pane on the left, switch to the **Logical Clusters** page, locate the row that contains the logical cluster to be deleted, and click **Delete** in the **Operation** column.
- Step 4** Confirm that all information is correct and click **OK**.

----End

17.3 Managing Logical Clusters

Editing a Logical Cluster

Precautions

- Nodes are added to or removed from a logical cluster by ring.
- At least one ring must be reserved in a logical cluster.
- The ring removed from the logical cluster will be added to the elastic cluster.

- Logical clusters of version 8.1.3 and later support online scale-out. Clusters of version 8.2.1.100 and later support job termination.

Procedure

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane, choose **Clusters > Dedicated Clusters**.
- Step 2** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 3** In the navigation pane, choose **Logical Clusters** and click **Edit** in the **Operation** column of the target cluster.
- Step 4** Add a node to the logical cluster by moving the selected ring from the right to the left, or remove a node from the logical cluster by moving the selected ring from the left to the right, and click **OK**.
- Step 5** When adding a node, select online or offline scale-out as needed.
- Step 6** If you select online scale-out, you can configure job termination. When job termination is enabled and there is congestion during online scale-out, the system will wait for the specified duration before terminating the congested jobs. The duration value must be an integer ranging from **30** to **1200**.

----End

Managing Resources (in a Logical Cluster)

Precautions

The original resource pool configuration is cleared when the cluster is converted from physical to logical. You have to add the resource pool again if you want to configure it after the conversion.

Procedure

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane, choose **Clusters > Dedicated Clusters**.
- Step 2** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.
- Step 3** In the navigation pane, choose **Logical Cluster Management**. In the **Operation** column of a logical cluster, click **Resource Management Configurations**. On the displayed page, you can manage resources in a logical cluster. For details, see [Resource Management](#).

----End

Restarting Logical Clusters

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane, choose **Clusters > Dedicated Clusters**.
- Step 2** In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.

Step 3 In the navigation pane on the left, switch to the **Logical Clusters** page, locate the row that contains the logical cluster to be restarted, and click **More > Restart** in the **Operation** column.

Step 4 Confirm that all information is correct and click **OK**.

----End

Scaling Out a Logical Cluster

Prerequisites

- Logical clusters of version 8.1.3 and later support online scale-out.
- Before a scale-out, you need to enable the logical cluster mode and add a logical cluster.
- After scaling out or scaling in a logical cluster, you need to reconfigure the backup policy for full backup. For details, see [Configuring an Automated Snapshot Policy](#).

Procedure

Step 1 Log in to the GaussDB(DWS) console. In the navigation pane, choose **Clusters > Dedicated Clusters**.

Step 2 On the displayed **Clusters** page, choose **More > Scale Node > Scale Out**.

Step 3 On the scale-out page, select a logical or elastic cluster. Enabling online scale-out also allows you to enable congested job termination. If congestion occurs during online scale-out and job termination is enabled, the system will wait for the specified duration before terminating congested jobs. The duration value must be an integer between **30** and **1200**.

NOTE

Clusters of version 8.2.1.100 and later support job termination.

----End

Scaling In a Logical Cluster

Constraints

- A host ring with CN nodes cannot be scaled in.
- A host ring with GTM and CM nodes cannot be scaled in.

Procedure

Step 1 Log in to the GaussDB(DWS) console. In the navigation pane, choose **Clusters > Dedicated Clusters**.

Step 2 In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.

Step 3 In the navigation pane on the left, switch to the **Logical Clusters** page. Locate the row that contains the target elastic pool and click **Edit** in the **Operation** column to delete nodes from the elastic pool (move the selected ring from the left to the right).

Step 4 Click **OK**.

----End

17.4 Scheduling GaussDB(DWS) 3.0 Logical Cluster Creation and Deletion

Context

You can schedule the creation and deletion of logical clusters. Computing logical clusters can be created and deleted during the scheduled period of time to dynamically scale computing resources.

NOTE

- This feature supports only GaussDB(DWS) 3.0 clusters. For an earlier version, contact technical support to upgrade it first.
- By default, a logical cluster created using this feature is used to provide computing power. After a user is associated with the logical cluster, the user's queries are processed by this logical cluster, but table creation statements are still processed in the original logical cluster.
- A user can be bound to only one computing logical cluster.
- If a user associated with the read-only logical cluster has workloads in progress when the cluster is deleted, an error may be reported.

Procedure

Step 1 Log in to the GaussDB(DWS) console. In the navigation pane, choose **Clusters > Dedicated Clusters**.

Step 2 In the cluster list, click the name of the target cluster. The **Cluster Information** page is displayed.

Step 3 After GaussDB(DWS) 3.0 clusters are created, they are in logical cluster mode.

- GaussDB(DWS) 3.0 cluster: In the navigation pane on the left, switch to the **Logical Clusters** page, click **Add Plan** to configure a proper scheduling plan.

Step 4 Select a plan type. It can be:

- **Periodically**: The plan is executed once in every specified period (week or month). A logical cluster is created or deleted as scheduled as long as it does not conflict with other O&M operations.
- **One-time**: The plan is executed only once in the specified period.

Step 5 Click **OK**. In the scheduled plan list, you can view the plan details and next execution time.

NOTE

To avoid affecting services and ensure resources are available at the scheduled time, the plan may be skipped if it conflict with O&M operations, and may be executed about 20 minutes earlier than planned if the cluster creation is time-consuming.

----End

17.5 Tutorial: Converting a Physical Cluster That Contains Data into a Logical Cluster

Scenario

A large database cluster usually contains a large amount of data put in different tables. With the **resource management** feature, you can create resource pools to isolate the resources of different services. Different service users can be allocated to different resource pools to reduce resource (CPU, memory, I/O, and storage) competition between services.

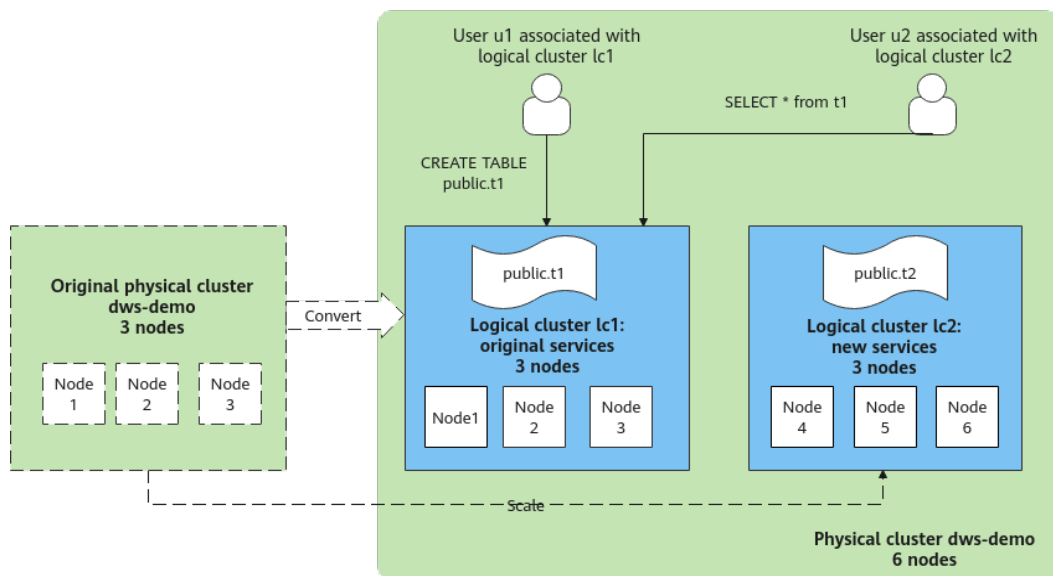
As the service scale grows, the number of services in the cluster system also increases. Creating multiple resource pools becomes less effective in controlling resource competition. GaussDB(DWS) uses the distributed architecture, and its data is distributed on multiple nodes. Each table is distributed across all DN nodes in the cluster, an operation on a data table may involve all DN nodes, which increases network loads and system resource consumption. To solve this problem, scale-out is not effective. You are advised to divide a GaussDB(DWS) cluster into multiple logical clusters.

You can create a separate logical cluster and assign new services to it. This way, new services have little impact on existing services. Also, if the service scale in existing logical clusters grows, you can scale out the existing logical clusters.

Figure 17-5 shows an example. The original service data tables of a company are stored in the original physical cluster **dws-demo** (in green). After services are switched over to the logical cluster **lc1** (in blue), a new logical cluster **lc2** is added to the physical cluster through scale-out. The original service data tables are switched to logical cluster **lc1**, and new service data tables are written to logical cluster **lc2**. In this way, the data of old and new services is isolated. User **u2** associated with logical cluster **lc2** can access the tables of logical cluster **lc1** across logical clusters after authorization.

- **Cluster scale:** Scale out the original physical cluster from three nodes to six nodes and split it into two logical clusters.
- **Service isolation:** New and old service data is isolated in different logical clusters.

Figure 17-5 Accessing data across logical clusters



Creating a Cluster and Preparing Table Data

- Step 1** Create a cluster. For details, see [Creating a DWS 2.0 Cluster](#).
- Step 2** After connecting to the database, create table **name** as the system administrator **dbadmin** and insert two data records into the table.

```
CREATE TABLE name (id int, name varchar(20));
INSERT INTO name VALUES (1,'joy'),(2,'lily');
```

----End

Converting to Logical Cluster lc1


NOTICE

During the conversion, you can run simple DML statements, such as adding, deleting, modifying, and querying data. Complex DDL statements, such as operations on database objects, will block services. You are advised to perform the conversion during off-peak hours.

- Step 1** Log in to the GaussDB(DWS) console. In the navigation pane, choose **Clusters > Dedicated Cluster**. Click the name of a cluster to go to the **Cluster Information** page.
- Step 2** Toggle on the **Logical Cluster** switch.
- Step 3** In the navigation pane on the left, choose **Logical Clusters**.
- Step 4** Click **Add Logical Cluster** in the upper right corner, enter the logical cluster name **lc1**, and click **OK**.

During the switchover, the current cluster is unavailable. Wait for about 2 minutes (the conversion time varies depending on the service data volume). If **lc1** is displayed on the logical cluster page, the conversion is successful.

Figure 17-6 Adding a logical cluster

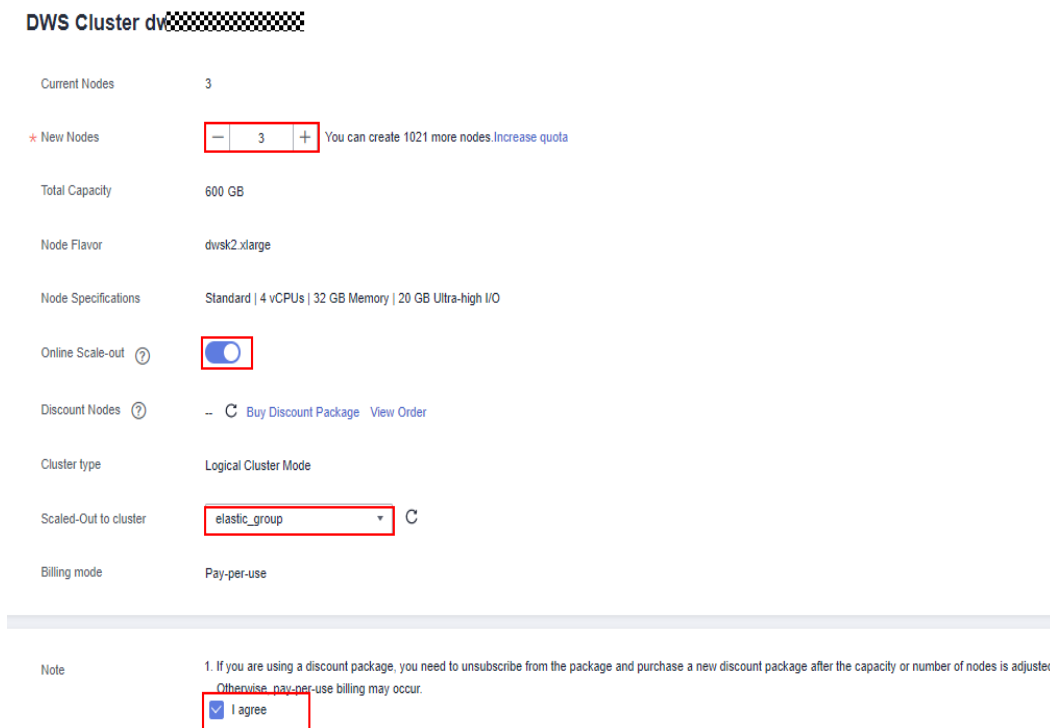


----End

Adding Nodes to the elastic_group Cluster

- Step 1** Return to the **Cluster Management** page. In the **Operation** column of the cluster, choose **More > Scale Node > Scale Out**.
- Step 2** Set **New Nodes** to **3**. Enable **Online Scale-out**. Set **elastic_group** as the target logical cluster. Confirm the settings, select the confirmation check box, and click **Next: Confirm**.

Figure 17-7 Scale-out process



- Step 3** Click **Next: Confirm**, and then click **OK**.
Wait for about 10 minutes until the scale-out is successful.

----End

Adding Logical Cluster lc2

- Step 1** On the **Cluster Management** page, click the name of a cluster to go to the cluster details page. In the navigation pane, choose **Logical Clusters**.

Step 2 Click **Add Logical Cluster** in the upper right corner, select three nodes from the right pane to add to the left pane, enter the logical cluster name **lc2**, and click **OK**.

After about 2 minutes, the logical cluster is successfully added.

Figure 17-8 Adding a logical cluster

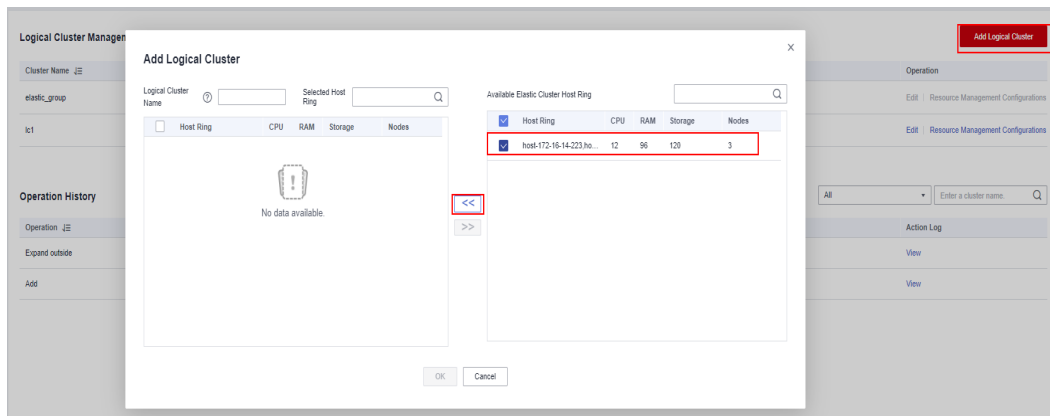
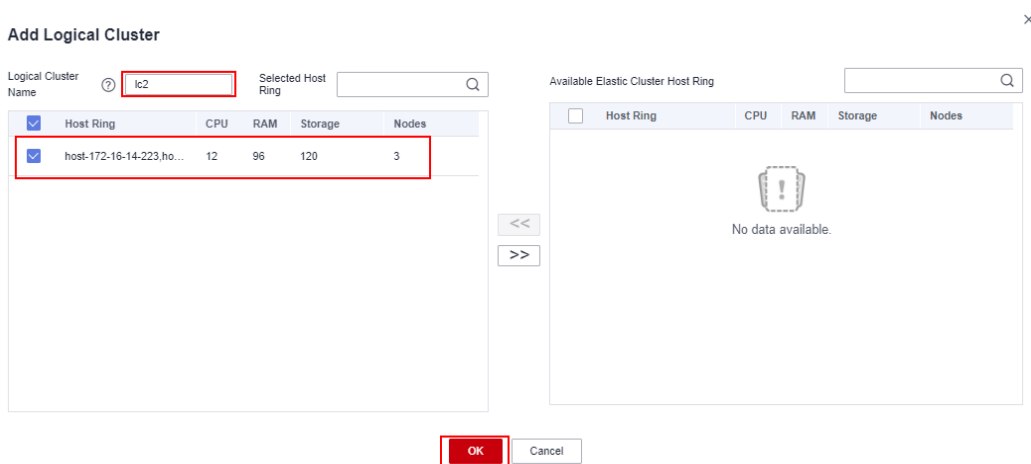


Figure 17-9 Selecting a host ring



----End

Creating Logical Clusters, Associating Them with Users, and Querying Data Across Logical Clusters

Step 1 Connect to the database as the system administrator and run the following SQL statement to query the original service table **name**.

Verify that service data can be queried after the conversion.

```
SELECT * FROM name;
```

Step 2 Create logical clusters **lc1** and **lc2** for **u1** and **u2**, respectively.

```
CREATE USER u1 NODE GROUP "lc1" password '{password}';  
CREATE USER u2 NODE GROUP "lc2" password '{password}';
```

Step 3 Log in to the database as user **u1**, create table **u1.t1**, insert two data records into the table, and grant user **u2** the permission to access the table.

```
CREATE TABLE u1.t1 (id int, name varchar(20));
INSERT INTO u1.t1 VALUES (1,'joy'),(2,'lily');
GRANT USAGE ON SCHEMA u1 TO u2;
GRANT SELECT ON TABLE u1.t1 TO u2;
```

- Step 4** Log in to the database as user **u2** and query data in the original service table **t1**. A message is displayed, indicating that you do not have the permission to access logical cluster **lc1**. The result shows that even if user **u1** has authorized user **u2** to access the table, the table cannot be accessed because it is in different logical clusters. This proves that data is isolated between logical clusters.

```
SELECT * FROM u1.t1;
```

Result Information	SQL Details	Status	Times
<p>Notice</p> <p>error_code="DWS_S0010001",error_msg="sql error STATE: 42501,message: ERROR: USAGE permission denied to user 'u2' for node group 'lc1'")</p>	SELECT * FROM u1.t1	Running failed	6ms

Total Records: 1

- Step 5** Switch back to system administrator **dbadmin** and grant the access permission of logical cluster **lc1** to user **u2**.

```
GRANT USAGE ON NODE GROUP lc1 TO u2;
```

- Step 6** Switch to user **u2** and query the **t1** table. This proves that the user bound to logical cluster **lc2** can query the original service table **t1** across logical clusters. In this way, data is shared between logical clusters.

```
SELECT * FROM u1.t1;
```

table	Chart	Export	Copy Column	Columns
id	name			id × name ×
1	joy			
2	lily			

Total Records: 2

----End

17.6 Tutorial: Dividing a New Physical Cluster into Logical Clusters

Scenario

This section describes how to divide a new six-node physical cluster (having no service data) into two logical clusters. If your physical cluster already has service data, perform operations by referring to [Tutorial: Converting a Physical Cluster That Contains Data into a Logical Cluster](#).

Prerequisites

Create a six-node cluster. For details, see [Creating a DWS 2.0 Cluster](#).

Dividing a Cluster into Logical Clusters

- Step 1** On the **Cluster Management** page, click the name of a cluster to go to the cluster details page. In the navigation pane, choose **Logical Clusters**.
- Step 2** Click **Add Logical Cluster** in the upper right corner, select a host ring (three nodes) on the right, add it to the list on the left, enter the logical cluster name **lc1**, and click **OK**.
- After about 2 minutes, the logical cluster is added.
- Step 3** Repeat the preceding steps to create the second logical cluster **lc2**.

----End

Creating Logical Clusters, Associating Them with Users, and Querying Data Across Logical Clusters

- Step 1** Connect to the database as system administrator **dbadmin** and run the following SQL statement to check whether the logical cluster is created:

```
SELECT group_name FROM PGXC_GROUP;
```

	id	name
1	1	Jcy
2	2	lly

- Step 2** Create users **u1** and **u2** and associate them with logical clusters **lc1** and **lc2**, respectively.

```
CREATE USER u1 NODE GROUP "lc1" password '{password}';  
CREATE USER u2 NODE GROUP "lc2" password '{password}';
```

- Step 3** Switch to user **u1**, create table **t1**, and insert data into the table.

```
SET ROLE u1 PASSWORD '{password}';  
CREATE TABLE u1.t1 (id int);  
INSERT INTO u1.t1 VALUES (1),(2);
```

- Step 4** Switch to user **u2**, create table **t2**, and insert data into the table.

```
SET ROLE u2 PASSWORD '{password}';  
CREATE TABLE u2.t2 (id int);  
INSERT INTO u2.t2 VALUES (1),(2);
```

- Step 5** Query the **u1.t1** table as user **u2**. The command output indicates that the user does not have the permission.

```
SELECT * FROM u1.t1;
```

```
ERROR: permission denied for schema u1
```

- Step 6** Switch back to the system administrator **dbadmin** and query the **u1.t1** and **u2.t2** tables, which are created in clusters **lc1** and **lc2**, respectively, corresponding to two services. In this way, data is isolated based on logical clusters.

```
SET ROLE dbadmin PASSWORD '{password}';  
SELECT p.oid,relname,pgroup,nodeoids FROM pg_class p LEFT JOIN pgxc_class pg ON p.oid = pg.pcrelid  
WHERE p.relname = 't1';  
SELECT p.oid,relname,pgroup,nodeoids FROM pg_class p LEFT JOIN pgxc_class pg ON p.oid = pg.pcrelid  
WHERE p.relname = 't2';
```

oid	relname	pgroup	nodeoids
25374	t1	lc1	16718 16719 16720

oid	reiname	pgroup	nodeoids
25377	t2	lc2	16676 16713 16717

Step 7 Grant user **u2** the permissions to access logical cluster **lc1**, schema **u1**, and table **u1.t1**.

```
GRANT usage ON NODE GROUP lc1 TO u2;  
GRANT usage ON SCHEMA u1 TO u2;  
GRANT select ON TABLE u1.t1 TO u2;
```

 **NOTE**

Logical clusters implement permission isolation (by node groups) based on physical clusters. To let a user access data across logical clusters, you need to grant the logical cluster (node-group layer) permissions, schema permissions, and table permissions to the user in sequence. If no logical cluster permissions are granted, the error message "permission denied for node group xx" will be displayed.

Step 8 Switch to user **u2** and query the **u1.t1** table. The query is successful. The logical cluster implements data isolation and allows cross-logical cluster access after user authorization.

```
SET ROLE u2 PASSWORD '{password}';  
SELECT * FROM u1.t1;
```

	id
1	2
2	1

----End

17.7 Tutorial: Setting a Read-Only Logical Cluster and Binding It to a User

Scenario

If your workloads vary greatly in different periods of time, a three-node cluster may be unable to handle all the throughput during peak hours; but a six-node cluster may be too large, wasting resources and increasing costs. In this case, you can follow this tutorial and the instructions in [Scheduling DWS 3.0 Logical Cluster Creation and Deletion](#) to use only three nodes during off-peak hours at night, six nodes during daytime, and nine nodes during peak hours.

This tutorial describes how to configure a new logical cluster (without service data) as read-only and switch some users to the cluster. In this way, tables created by those users are still in the original Node Group, but the computing logic is switched to the read-only logical cluster.

Prerequisites

A six-node cluster has been created and divided into two logical clusters: **v3_logical** and **lc1**. The **lc1** cluster has no service data. For details, see [Creating a DWS 3.0 Cluster](#).

Configuring a Read-Only Logical Cluster and Switching Users to the Cluster

- Step 1** Connect to the database as system administrator **dbadmin** and run the following SQL statement to check whether the logical cluster is created:

```
SELECT group_name FROM PGXC_GROUP;
```

- Step 2** Set logical cluster **lc1** to be read-only.

```
SET xc_maintenance_mode=on;  
ALTER NODE GROUP lc1 SET READ ONLY;  
SET xc_maintenance_mode=off;
```

- Step 3** Create a user.

```
create user testuser password 'testuser12#$$%';
```

- Step 4** Bind the user to the logical cluster **lc1**. Replace variables in the following statements (such as **testuser** and **lc1**) as needed.

Find the NodeGroup of the user. If a record can be found, set the record to the **default_storage_nodegroup** of the user so that the tables created by the user will still be in the original Node Group. If no records are found, directly run the two ALTER statements in the end.

```
SELECT nodegroup FROM pg_user WHERE username='testuser';  
ALTER USER testuser SET default_storage_nodegroup='nodegroup'; // Replace nodegroup with the node  
group name obtained in the preceding SQL statement.
```

Bind the user to the new read-only logical cluster. In this way, the computing logic of the user is switched to the read-only logical cluster for execution.

```
ALTER USER testuser NODE GROUP lc1;  
ALTER USER testuser SET enable_cudesc_streaming=ON;
```

----End