GaussDB

Service Overview

Issue 01

Date 2025-08-18





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

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

Trademarks and Permissions

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

Notice

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

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

Huawei Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road

Qianzhong Avenue Gui'an New District Gui Zhou 550029

People's Republic of China

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

i

Contents

1 What Is GaussDB?	1
2 Scenarios	3
3 Basic Concepts	
4 Advantages	
5 DB Instance Description	7
5.1 DB Instance Types	
5.2 Instance Storage Types	13
5.3 DB Engine Versions	13
5.4 Instance Statuses	13
6 Instance Specifications	15
7 Permissions Management	16
8 Notes and Constraints	25
9 Related Services	31

1 What Is GaussDB?

GaussDB is a distributed relational database. It supports intra-city cross-AZ deployment. With a distributed architecture, GaussDB supports petabytes of storage and contains more than 1,000 nodes per DB instance. It is highly available, reliable, secure, and scalable and provides services including quick deployment, backup, restoration, monitoring, and alarm reporting for enterprises.

Architecture of a Distributed GaussDB Instance

The overall architecture of a distributed GaussDB instance is as follows.

Application Safe Tool sets sectors GTM Coordinator node (Distributed CBO) (Distributed transactions) Computing network: TCP/RDMA Data migration Data node Data node Distributed SQL Distributed SQL Primary/standby switchover or failover for HA Distributed execution Distributed execution Cross-AZ multi-active 3DC geo-redundant DR development Storage engine Storage engine x86/Kunpeng + EVS storage

Figure 1-1 Overall architecture of a distributed GaussDB instance

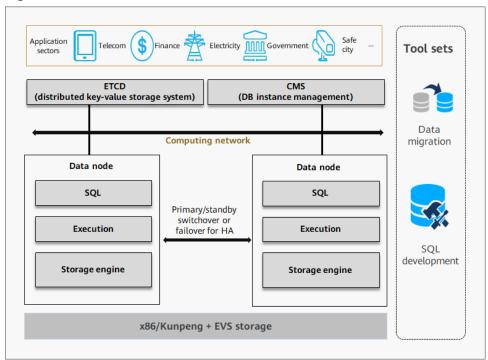
• **Coordinator node**: A coordinator node (CN) receives access requests from applications and returns execution results to clients. It also splits and distributes tasks to different data nodes (DNs) for parallel processing.

- **GTM**: The Global Transaction Manager (GTM) generates and maintains the global transaction IDs, transaction snapshots, timestamps, and sequences that must be unique globally.
- **Data node**: A DN stores business data, performs data queries, and returns execution results to a CN.

Architecture of a Centralized GaussDB Instance

The overall architecture of a centralized GaussDB instance is as follows.

Figure 1-2 Overall architecture of a centralized GaussDB instance



- **ETCD**: The Editable Text Configuration Daemon (ETCD) is used for shared configuration and service discovery (service registry and search).
- **CMS**: The Cluster Management Server (CMS) is a component used to manage cluster status.
- **Data node**: A data node (DN) stores data, executes data query tasks, and returns execution results.

2 Scenarios

Transaction applications

The distributed, highly scalable architecture of GaussDB makes it an ideal fit for highly concurrent online transactions involving a large volume of data from government, finance, e-commerce, O2O, telecom customer relationship management (CRM), and billing. GaussDB supports different deployment models.

CDR query

GaussDB excels in processing petabytes of data. Powered by the memory analysis technology, GaussDB can be used to query massive volumes of data when the data is being written to databases. This makes GaussDB an excellent choice for the Call Detail Record (CDR) query service in the security, telecom, finance, and Internet of things (IoT) sectors.

3 Basic Concepts

Instances

The smallest management unit of GaussDB is the instance. A DB instance is an isolated database environment on the cloud. You can create and manage GaussDB instances on the management console. For details about the storage types, versions, and instance statuses, see **DB Instance Description**.

DB Engine Versions

GaussDB supports the following versions: V2.0-3.226, V2.0-8.102, V2.0-8.103, and V2.0-8.210.

Instance Types

GaussDB supports distributed instances and centralized instances.

- Distributed instances allow you to add nodes as needed to handle large volumes of concurrent requests.
- Centralized instances are suitable for scenarios with small and stable volumes of data, where data reliability and service availability are extremely important.

Instance Specifications

The instance specifications determine the computation (vCPUs) and memory capacity (in GB) of an instance. For details, see **Instance Specifications**.

Coordinator Nodes

A coordinator node (CN) receives access requests from applications, distributes tasks to DNs for parallel processing, and returns execution results to clients.

Data Nodes

A data node (DN) stores data, executes data query tasks, and returns execution results to CNs.

Shards

A shard contains multiple DNs. The number of DNs in a shard depends on the value of **Replicas**, for example, if **Replicas** is set to **3**, there are three DNs (one primary and two standby DNs) in a shard.

Automated Backups

When you buy a GaussDB instance, an automated backup policy is enabled by default with the retention period set to seven days. You can **modify** it as required. GaussDB automatically creates full backups for the instance based on your configuration.

Manual Backups

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

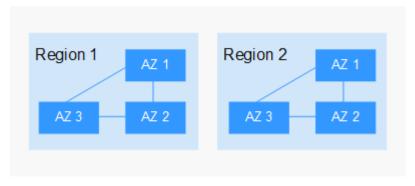
Regions and AZs

A region and availability zone (AZ) identify the location of a data center. You can create resources in a specific region and AZ.

- A region is a physical data center. Each region is isolated from the other regions, improving fault tolerance and stability. The region that is selected during resource creation cannot be changed after the resource is created.
- An AZ is a physical location using independent power supplies and networks.
 Faults in an AZ do not affect other AZs. A region contains one or more AZs that are physically isolated but interconnected through internal networks. This ensures the independence of AZs and provides low-cost, low-latency network connections.

Figure 3-1 shows the relationship between regions and AZs.

Figure 3-1 Regions and AZs



4 Advantages

High Security

GaussDB provides features to let you enjoy the security of top-level commercial databases. These features meet the core data security requirements of enterprises, government, and finance institutions.

Supported security features: dynamic data masking, row-level access control, and cryptographic computing

• Comprehensive Tools and Service-oriented Capabilities

GaussDB can be provided for general availability and can work with ecosystem tools such as Data Replication Service (DRS) and Data Admin Service (DAS) to make database development, O&M, tuning, monitoring, and migration easy.

• In-House, Full-Stack Development

GaussDB provides guaranteed business continuity with full-stack, independent and controllable capabilities. GaussDB performance is always improved to meet ever-increasing demands in different scenarios.

• Open-Source Ecosystem

GaussDB provides an open-source version.

5 DB Instance Description

5.1 DB Instance Types

Table 5-1 lists the instance types of GaussDB.

Table 5-1 Instance types

Inst anc e Typ e	Deployment Model	Scale- out Support ed	Use Case	Comp onent	Service Processing
Dist rib ute d	Independent: Database components are deployed on different nodes. This model is suitable for where high availability and stability are required and the instance scale is large.	Yes	The data volume is large, and large data capacity and high concurrenc y are required.	• O M • CM • ETC D C D D D D D D D D D D D D D D D D D	An application sends a SQL query request to a CN. The CN uses the optimizer of the database to generate an execution plan and sends the plan to DNs. Each DN processes data based on the execution plan. After the processing is complete, DNs return the result sets to the CN for summary. Finally, the CN returns the summary result to the application.

Inst anc e Typ e	Deployment Model	Scale- out Support ed	Use Case	Comp onent	Service Processing
Cen tral ize d	1 primary + 2 standby: 3-node deployment where there is a shard. The shard contains one primary DN and two standby DNs. Single: single-node deployment where there is only one CMS component and one DN. This deployment model is not suitable for production environments. To create a single-replica instance, ensure that the instance version is V2.0-2.2 or later. The availability (SLA) cannot be guaranteed because the instance is deployed on a single server.	No	The data volume is small and stable, and data reliability and service availability are important.	O M CM ETC D D D	An application sends a task directly to DNs, and the DNs then process and return the results to the application.

Ⅲ NOTE

- OM: The Operation Manager (OM) provides management APIs and tools for routine maintenance and configuration management of the cluster.
- CM: The Cluster Manager (CM) manages and monitors the running status of functional units and physical resources in a distributed system, ensuring stable running of the entire system. The CM acts as the overseer of the system's health, regularly assessing the status of DNs to ensure optimal performance. Once an issue arises, the CM takes corrective measures, like restarting a DN or promoting a standby DN to primary.
- GTM: The Global Transaction Manager (GTM) generates and maintains the global transaction IDs, transaction snapshots, timestamps, and sequences that must be unique globally.
- ETCD: The Editable Text Configuration Daemon (ETCD) is used for shared configuration and service discovery (service registry and search). ETCD stores metadata information, for example, the status of each DN. The CM checks in DN statuses with ETCD to stay informed and make management decisions.
- CN: A coordinator node (CN) receives access requests from applications, distributes
 tasks to DNs for parallel processing, and returns execution results to clients. Each CN
 can connect to all DNs, and all CNs share the same metadata.
- DN: A data node (DN) stores data, executes data query tasks, and returns execution results to CNs.

Distributed

The following figure shows the logical architecture of a distributed GaussDB instance.

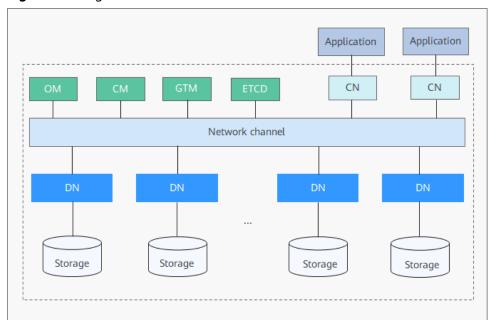


Figure 5-1 Logical architecture of a distributed instance

Deployment models for distributed instances

Independent deployment

In the independent deployment model, database components are independently deployed on different servers. That is, each CN or DN occupies a VM and resources are isolated from each other. Take the default configuration (3 CNs, 3 replicas, and 3 shards) as an example. In the independent deployment model, such configuration requires 12 nodes (VMs) in total, including three CNs and nine DNs, and the cost of independent deployment is relatively high. Figure 5-2 shows the architecture. This model is suitable for where high availability and stability are required and the instance scale is large.

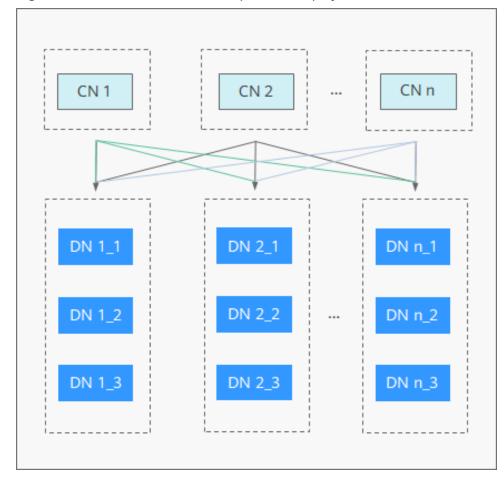


Figure 5-2 Architecture of the independent deployment model

Centralized

The following figure shows the logical architecture of a centralized GaussDB instance.

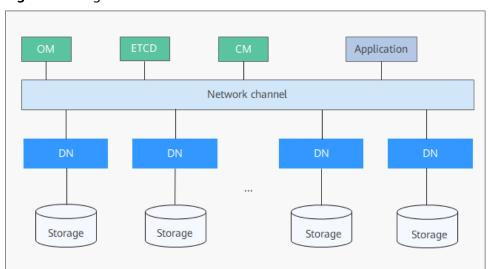


Figure 5-3 Logical architecture of a centralized instance

Deployment models for centralized instances

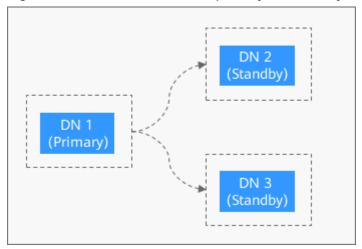
• 1 primary + 2 standby

The 1 primary + 2 standby deployment model is a 3-node deployment where there is one shard. The shard contains one primary DN and two standby DNs. **Figure 5-4** shows the architecture. This model is only available for centralized instances. Instances deployed using this model are not scalable. This model is suitable for scenarios with small and stable volumes of data, where data reliability and service availability are important.

◯ NOTE

- If any of primary and standby DNs is faulty, services will not be affected.
- 1 primary + 2 standby instances can be deployed in either of the following modes:
 - 3 replicas, single AZ
 - 3 replicas, 3 AZs

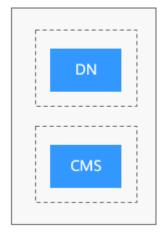
Figure 5-4 Architecture of the 1 primary + 2 standby deployment model



• Single (also known as single-replica)

The single-replica deployment model is a single-node deployment where there is only one CMS and one DN. This model is only available for centralized instances. Figure 5-5 shows the architecture.

Figure 5-5 Architecture of the single-replica deployment model



5.2 Instance Storage Types

Database systems are one of the most important IT systems and have high requirements on storage I/O performance. GaussDB provides ultra-high I/O storage with up to 350 MB/s of throughput.

5.3 DB Engine Versions

GaussDB supports the following versions: V2.0-3.226, V2.0-8.102, V2.0-8.103, and V2.0-8.210.

GaussDB Engine Version Description

GaussDB uses version numbers in dot-decimal notation.

The format is **V2.0-***A.BCD*, where:

- V2.0: Versions starting with V2.0 use the second-generation architecture of GaussDB.
- A: A designates an annual requirement baseline version.
- *BCD*: The first digit (*B*) represents a semi-annual requirement baseline version, and C and D indicate verified patch versions for third-party adaptation.

5.4 Instance Statuses

Instance Statuses

The status of a DB instance reflects the health of the instance. You can use the management console to view the status of a DB instance.

Table 5-2 Instance statuses

Status	Description
Available	The DB instance is available.
Abnormal	The DB instance is unavailable.
Creating	The DB instance is being created.
Creation failed	The DB instance failed to be created.
Rebooting	The DB instance is being rebooted.
Starting	The DB instance is being started.
Starting node	The DB instance node is being started.
Stopping	The DB instance is being stopped.

Status	Description	
Stopping node	The DB instance node is being stopped.	
Stopped	The DB instance or node is stopped.	
Scaling up	The storage space of the DB instance is being scaled up.	
Adding nodes	Nodes are being added to the DB instance.	
Backing up	A backup is being created.	
Restoring	The instance is being restored from a backup.	
Restore failed	The instance failed to be restored.	
Storage full	The storage space of the instance is full. No more data can be written to the databases on this instance.	
Deleted	The instance has been deleted. Deleted instances will not be displayed in the instance list.	
Upgrading	The engine version is being upgraded.	
Parameters change. Pending reboot	A modification to a database parameter is waiting for a DB instance reboot before it can take effect.	

Backup Statuses

Table 5-3 Backup statuses

Status	Description
Completed	The backup was successfully created.
Failed	The backup failed to be created.
Creating	The backup is being created.

6 Instance Specifications

Table 6-1 Instance specifications

Specificati on Type	vCPUs	Memory (GB)	Default Maximum Connections (Per CN)
Dedicated (1:4) NOTE Dedicated (1:4) is based on the x86 architectur	4	16 NOTE This specification is not available for production environments.	N/A
	8	32	N/A
e.	16	64	N/A
	32	128	N/A
	64	256	N/A
Dedicated (1:8) NOTE Dedicated (1:8) is based on the x86 architectur e.	8	64	2,000
	16	128	4,000
	32	256	9,000
	64	512	18,000

Permissions Management

If you need to assign different permissions to employees in your company to access your GaussDB resources, IAM is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you securely manage access to your resources.

If your account does not need individual IAM users for permissions management, you can skip this section.

With IAM, you can use your account to create IAM users for your employees, and assign specific permissions to different users to control their access to specific resource types. For example, you can grant software developers in your company permissions to use GaussDB resources but not the permissions needed to delete them or perform any high-risk operations.

GaussDB Permissions

By default, new IAM users do not have any permissions assigned. You need to add a user to one or more groups, and attach permission policies or roles to these groups. Users inherit permissions from the groups to which they are added and can perform specified operations on cloud services.

GaussDB is a project-level service deployed for specific regions. To assign GaussDB permissions to a user group, specify the scope as region-specific projects and select the project for the permissions to take effect. If **All projects** is selected, the permissions will be granted to the user group in all region-specific projects. When accessing GaussDB, the users need to switch to the authorized region.

You can use roles and policies to manage user permissions.

- Roles: A coarse-grained way of granting permissions related to users responsibilities. Only a limited number of service-level roles for authorization are available. When using roles to grant permissions, you may need to assign additional roles because of the different dependencies involved with rolebased permissions. Roles are not ideal for fine-grained authorization and least privilege access.
- Policies: A more fine-grained system. Policies let you define permissions required to perform operations on specific cloud resources under certain conditions. This mechanism allows for more flexible policy-based authorization on a principle of least privilege (PoLP) basis. For example, you

can grant IAM users only the permissions needed to manage a certain type of GaussDB resources. Most policies define permissions based on APIs.

Table 7-1 lists all the system-defined policies supported by GaussDB.

Table 7-1 System-defined permissions for GaussDB

Policy Name	Description	Category
GaussDB FullAccess	Full permissions for GaussDB	System-defined policy
GaussDB ReadOnlyAccess	Read-only permissions for GaussDB	System-defined policy

Table 7-2 lists the common operations supported by each system policy of GaussDB. Choose appropriate system policies based on this table.

Table 7-2 Common operations supported by the GaussDB system policies

Operation	GaussDB FullAccess	GaussDB ReadOnlyAccess
Creating a GaussDB instance	√	х
Deleting a GaussDB instance	√	х
Querying GaussDB instances	√	√

Table 7-3 Common operations and supported actions

Operation	Action	Remarks
Creating a DB instance	gaussdb:instance:create gaussdb:param:list	To select a VPC, subnet, and security group, configure the following actions:
		vpc:vpcs:list
		vpc:vpcs:get
		vpc:subnets:get
		vpc:securityGroups:get
		To create an encrypted instance, configure the following actions for the project:
		kms:cmk:get
		kms:cmk:list
		To report event monitoring of a failed operation, configure the following actions:
		"ces:alarmsOnOff:put"
		"ces:alarms:create"
Changing instance specifications	gaussdb:instance:modifySpec	To report event monitoring of a failed operation, configure the following actions: "ces:alarmsOnOff:put"
		"ces:alarms:create"
Adding nodes	gaussdb:instance:modifySpec	To report event monitoring of a failed operation, configure the following actions: "ces:alarmsOnOff:put" "ces:alarms:create"
Scaling up storage space	gaussdb:instance:modifySpec	To report event monitoring of a failed operation, configure the following actions: "ces:alarmsOnOff:put" "ces:alarms:create"

Operation	Action	Remarks
Rebooting a DB instance	gaussdb:instance:restart	To report event monitoring of a failed operation, configure the following actions: "ces:alarmsOnOff:put" "ces:alarms:create"
Deleting a DB instance	gaussdb:instance:delete	To report event monitoring of a failed operation, configure the following actions: "ces:alarmsOnOff:put" "ces:alarms:create"
Querying instances	gaussdb:instance:list	None
Querying instance details	gaussdb:instance:list	If the VPC, subnet, and security group are displayed in the instance list, configure vpc:*:get and vpc:*:list. If the used disk is displayed, configure ces:*:list.
Changing a DB instance password	gaussdb:instance:modify	To report event monitoring of a failed operation, configure the following actions: "ces:alarmsOnOff:put" "ces:alarms:create"
Changing a DB instance name	gaussdb:instance:modify	None
Binding or unbinding an EIP	gaussdb:instance:modify	To display EIPs on the console, configure the following actions: vpc:publicIps:get vpc:publicIps:list To report event monitoring of a failed operation, configure the following actions: "ces:alarmsOnOff:put" "ces:alarms:create"
Creating a parameter template	gaussdb:param:create gaussdb:param:list	None

Operation	Action	Remarks
Modifying a parameter template	gaussdb:param:modify	None
Obtaining parameter templates	gaussdb:param:list	None
Applying a parameter template	gaussdb:param:apply	To report event monitoring of a failed operation, configure the following actions: "ces:alarmsOnOff:put" "ces:alarms:create
Deleting a parameter template	gaussdb:param:delete	None
Creating a manual backup	gaussdb:backup:create	To report event monitoring of a failed operation, configure the following actions: "ces:alarmsOnOff:put" "ces:alarms:create"
Obtaining backups	gaussdb:backup:list	None
Modifying the backup policy	gaussdb:instance:modifyBacku pPolicy	None
Deleting a manual backup	gaussdb:backup:delete	To report event monitoring of a failed operation, configure the following actions: "ces:alarmsOnOff:put" "ces:alarms:create"
Restoring data to a new DB instance	gaussdb:instance:create	To select a VPC, subnet, and security group, configure the following actions: vpc:vpcs:list vpc:vpcs:get vpc:subnets:get vpc:securityGroups:get To report event monitoring of a failed operation, configure the following actions: "ces:alarmsOnOff:put" "ces:alarms:create

Operation	Action	Remarks
Querying project tags	gaussdb:tag:list	None
Adding or deleting project tags in batches	gaussdb:instance:dealTag	None
Modifying quotas	gaussdb:quota:modify	None
Querying predefined tags	gaussdb:instance:list	To query predefined tags, configure the following action: tms:resourceTags:list
Querying configured log groups	-	To query configured log groups, configure the following action: lts:groups:get
Querying configured log streams	-	To query configured log streams, configure the following action: lts:topics:get
Enabling autoscaling	gaussdb:instance:autoEnlargeP olicy	To enable autoscaling, configure the following actions for the IAM users instead of your Huawei account:
		iam:agencies:listAge ncies
		• iam:agencies:create Agency
		 iam:permissions:list RolesForAgencyOn- Project
		 iam:permissions:gra ntRoleToGroupOn- Project
		iam:roles:listRoles
Querying advanced features	gaussdb:instance:listFeatures	None
Enabling advanced features	gaussdb:instance:updateFeatur es	None

Table 7-4 DR operations and supported actions

Operation	Action	Remarks
Querying instances that can establish a DR relationship with a primary instance	gaussdb:disasterRecovery:list gaussdb:instance:listAll	None
Querying the real- time DR monitoring status of an instance	gaussdb:disasterRecovery:list gaussdb:disasterRecovery:get	None
Querying instances that have established a DR relationship	gaussdb:disasterRecovery:list gaussdb:disasterRecovery:list All	None
Resetting the DR relationship	gaussdb:disasterRecovery:con struct gaussdb:disasterRecovery:cre ate	To enable the DR operation, configure the following actions for the IAM users instead of your Huawei account: • iam:agencies:listAgenc ies • iam:agencies:createAg ency • iam:permissions:listRo lesForAgencyOnProject • iam:permissions:grant RoleToGroupOnProject • iam:roles:listRoles An agency RDSAccessProjectResource (including the policy DBS AgencyPolicy) will be automatically created in the current region.

Operation	Action	Remarks
Establishing a DR relationship	gaussdb:disasterRecovery:con struct gaussdb:disasterRecovery:cre ate	To enable the DR operation, configure the following actions for the IAM users instead of your Huawei account: • iam:agencies:listAgenc ies • iam:agencies:createAg ency • iam:permissions:listRo lesForAgencyOnProject • iam:permissions:grant RoleToGroupOnProject • iam:roles:listRoles An agency RDSAccessProjectResource (including the policy DBS AgencyPolicy) will be automatically created in
		the current region.
Promoting the DR instance to primary	gaussdb:disasterRecovery:fail over	The RDSAccessProjec- tResource agency is required.
Removing a DR relationship	gaussdb:disasterRecovery:rele ase	The RDSAccessProjec- tResource agency is required.
Switching roles of primary and DR instances	gaussdb:disasterRecovery:swi tchover	The RDSAccessProjec- tResource agency is required.
Re-establishing a DR relationship	gaussdb:disasterRecovery:con struct gaussdb:disasterRecovery:cre ate	The RDSAccessProjec- tResource agency is required.
Performing a DR drill	gaussdb:disasterRecovery:sim ulation	The RDSAccessProjec- tResource agency is required.
Caching logs for DR	gaussdb:disasterRecovery:kee plog	None
Querying operation records	gaussdb:instance:listRecord	None

□ NOTE

In DR scenarios, you also need to configure permissions and actions on the cloud where the DR instance resides before performing DR-related operations.

8 Notes and Constraints

To ensure the stability and security of GaussDB, certain constraints are put in place for access or permissions control.

There is no SLA commitment, so GaussDB single-replica instances cannot be used in production environments. For function constraints of single-replicas, see **Table 8-4**.

Naming

Table 8-1 Naming constraints

Item	Description
Instance name	 4 to 64 characters long Must start with a letter (case-sensitive). Only letters, digits, hyphens (-), and underscores (_) are allowed.
Backup name	 4 to 64 characters long Must start with a letter (case-sensitive). Only letters, digits, hyphens (-), and underscores (_) are allowed.
Parameter template name	 1 to 64 characters long Only letters (case-sensitive), digits, hyphens (-), underscores (_), and periods (.) are allowed.

Security

Table 8-2 Security constraints

Item	Description
root permissions	Only the root user permissions are available on the instance creation page.

Item	Description
root password	 8 to 32 characters long Must contain at least three types of the following: uppercase letters, lowercase letters, digits, and special characters. Supported characters: ~!@#%^*=+?, For more information, see .
Database port	The port ranges from 1024 to 39989, but the following ports that are reserved for system use cannot be used: 2378 to 2380, 2400, 4999 to 5001, 5100, 5500, 5999 to 6001, 6009 to 6010, 6500, 8015, 8097, 8098, 8181, 9090, 9100, 9180, 9187, 9200, 12016, 12017, 20049, 20050, 21731, 21732, 32122 to 32126, and 39001.
Disk encryption	If you enable disk encryption during instance creation, the disk encryption status and the key cannot be changed later.
VPC	The VPC where a GaussDB instance is located cannot be changed after the instance is created.
Security group	 By default, each tenant can create up to 500 security group rules. It is recommended that a security group contain a maximum of 50 rules.
System accounts	To provide O&M services, the system automatically creates system accounts when you create GaussDB instances. These accounts cannot be used by users.
	• rdsAdmin: a management account with the highest permission.
	rdsBackup: a backup account. It is used for backend backup.
	• rdsRepl : an account used to synchronize data from the primary instance to the standby instance.
	• root : the system administrator account.
	• rdsMetric: a metric monitoring account. It is used to obtain GaussDB performance metrics and report them to Cloud Eye for the tenant to view the status of GaussDB.

Instance Operations

Table 8-3 Constraints on instance operations

Function	Constraints	
Database access	 The ECSs must be allowed by the security group to access the target GaussDB instance. If a GaussDB instance and the ECSs belong to different security groups, no communication between them is established by default. To allow it, you must add an inbound rule to the GaussDB security group. The default port number of GaussDB instances is 8000. 	
Deployment	The servers where instances are deployed are not directly visible to you. You can only access the instances through IP addresses and database ports.	
Rebooting a GaussDB instance	GaussDB instances cannot be rebooted through commands. They must be rebooted on the management console.	
Querying GaussDB backups	GaussDB backup files are stored in OBS buckets and are not visible to you.	
Specification changes	 Before you change the instance specifications, ensure that the instance is available. If the instance or node is abnormal, or the storage space is full, you cannot perform this operation. During the specification change for HA (1 primary + 2 standby) instances, a primary/standby failover is triggered. During the failover, services are interrupted for about 1 minute. For a single-replica instance, changing instance 	
	 specifications will reboot the instance and interrupt services for 5 to 10 minutes. Changing the CPU and memory specifications will cause the instance to reboot, and services will be temporarily interrupted during the reboot. 	
Failover	For centralized instances, services are unavailable for about 10 seconds when the primary node is being switched to the standby node.	
Data restoration	To prevent data loss, you are advised to back up key data before data restoration.	
Storage space	If the storage space of an instance is full, data cannot be written to databases. You are advised to periodically check the storage space.	
Performance tuning	Performance tuning may need to reboot the instance and interrupt services.	

Function	Constraints
Recycle bin management	GaussDB allows you to move unsubscribed yearly/ monthly instances and deleted pay-per-use instances to the recycle bin. You can restore an instance that was deleted up to 7 days ago from the recycle bin.
	The recycle bin is enabled by default and cannot be disabled. The deleted instances can be retained for 7 days by default.

Table 8-4 Function constraints of single-replica centralized instances

Function	Supported by Versions Earlier Than V2.0-3.0	Supported by V2.0-3.0 and Later
Creating an instance	Yes	Yes
Rebooting an instance	Yes	Yes
Modifying parameters	Yes	Yes
Applying parameters	Yes	Yes
Resetting a password	Yes	Yes
Creating a full backup	No	Yes
Creating a differential backup	No	Yes
Deleting a backup	No	Yes
Modifying the backup policy	No	Yes
Restoring data to the original instance	No	Yes

Function	Supported by Versions Earlier Than V2.0-3.0	Supported by V2.0-3.0 and Later
Restoring data to a new instance	No	Yes
Scaling up storage	Yes	Yes
Changing instance specifications	Yes	Yes
Hot patch update	No	Yes
In-place upgrade	Yes (The version can be upgraded only to V2.0-3.0 or later.)	Yes
Gray upgrade	No	Yes (only supported in version V2.0-3.207 or later)
Viewing monitoring metrics	Yes	Yes
Deleting an instance	Yes	Yes
Rebuilding a deleted instance from the recycle bin	No	Yes
Querying the disk usage	Yes	Yes
Creating a database	Yes	Yes
Querying a database	Yes	Yes
Creating a schema and user	Yes	Yes

Function	Supported by Versions Earlier Than V2.0-3.0	Supported by V2.0-3.0 and Later
Deleting a schema and user	Yes	Yes
Performing database operations	Yes	Yes

9 Related Services

Table 9-1 shows the relationship between GaussDB and other services.

Table 9-1 Related services

Service	Description
Elastic Cloud Service (ECS)	Enables you to access ECS instances through an ECS to reduce application response time and public network traffic fees.
Virtual Private Cloud (VPC)	Isolates your network and controls access to your GaussDB instances.
Object Storage Service (OBS)	Stores automated and manual backups of your GaussDB instances.
Cloud Eye	Acts as an open monitoring platform that monitors GaussDB resources in real time. It reports alarms and issues warnings promptly to ensure you are notified immediately if your services are not running properly.
Cloud Trace Service (CTS)	Provides records of operations on cloud service resources, enabling you to query, audit, and backtrack operations.
Data Admin Service (DAS)	Provides a visualized GUI interface for you to connect and manage cloud databases.