

GaussDB

Service Overview

Issue 01
Date 2025-02-17



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

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

Trademarks and Permissions



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

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

Notice

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

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

Huawei Cloud Computing Technologies Co., Ltd.

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

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

Contents

1 What Is GaussDB?	1
2 Scenarios	4
3 Basic Concepts	5
4 Advantages	8
5 DB Instance Description	9
5.1 DB Instance Types	9
5.2 Instance Storage Types	16
5.3 DB Engine Versions	16
5.4 Instance Statuses	18
6 Instance Specifications	20
7 Edition Types	29
8 Security	32
8.1 Shared Responsibilities	32
8.2 Identity Authentication and Access Control	33
8.2.1 Service Access Control	33
8.3 Data Protection	34
8.4 Audit and Logs	35
8.5 Risk Monitoring	35
8.6 Fault Recovery	36
8.7 Certificates	36
9 Permissions Management	38
10 Notes and Constraints	52
11 Related Services	59
12 Compatibility	60

1 What Is GaussDB?

GaussDB is a distributed relational database from Huawei. It supports intra-city cross-AZ deployment with zero data loss. 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.

Progressive Knowledge

You can go to [GaussDB Progressive Knowledge](#) to learn the basic concepts and usage of GaussDB.

Using GaussDB

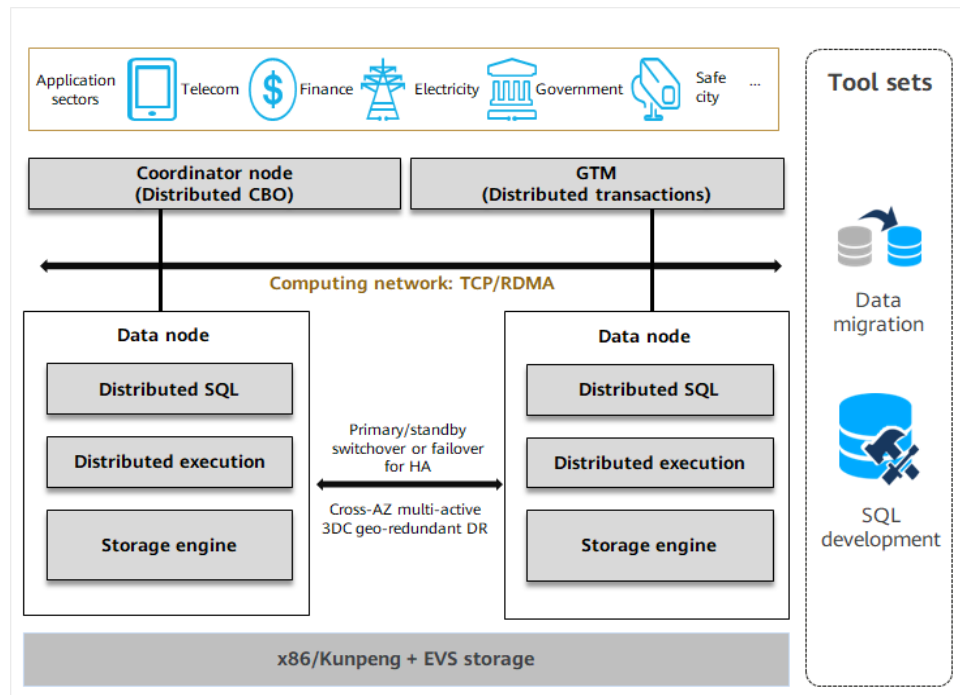
You can use GaussDB in either of the following ways:

- Management console: Create and manage DB instances on the [management console](#).
- API: You can call APIs to use GaussDB. For details, see [GaussDB API Reference](#).

Architecture of a Distributed GaussDB Instance

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

Figure 1-1 Overall architecture of a distributed GaussDB instance

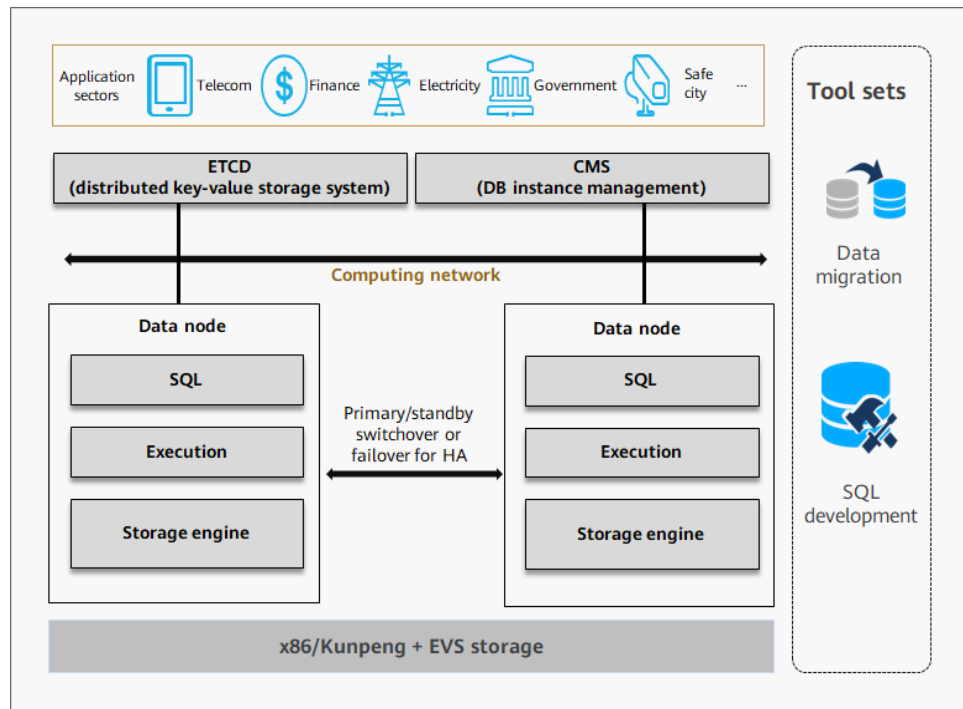


- **Coordinator node:** A coordinator node (CN) receives access requests from applications and returns execution results to clients. It also splits and distributes tasks to different data nodes (DNs) for parallel processing.
- **GTM:** The Global Transaction Manager (GTM) generates and maintains the 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 Manager (CMS) manages and monitors the running status of functional units and physical resources in a distributed system, ensuring stability of the entire system.
- **Data node:** A DN stores business data, performs data queries, 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 instance statuses, instance specifications, storage types, and versions, see [DB Instance Description](#).

DB Engine Versions

GaussDB V2.0-8.201 and V2.0-8.210 are supported.

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 and returns execution results to clients. It also splits and distributes tasks to different data nodes (DNs) for parallel processing.

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.

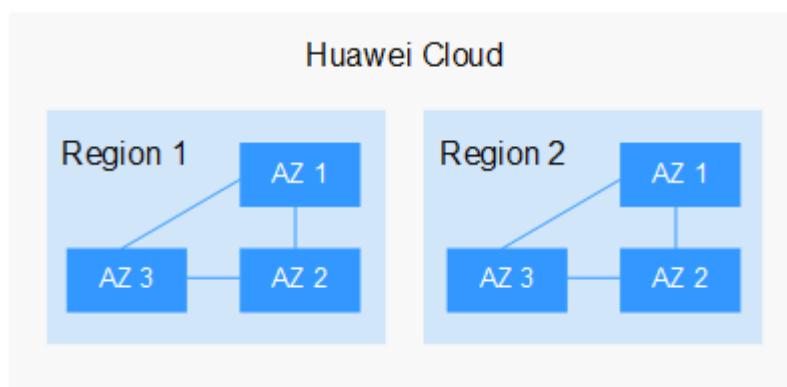
Regions and AZs

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

- Regions are defined by a combination of geographical location and network latency. Public services, such as Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Object Storage Service (OBS), Virtual Private Cloud (VPC), Elastic IP (EIP), and Image Management Service (IMS), are shared within the same region. Regions can be universal or dedicated. A universal region provides all sorts of cloud services for all users. A dedicated region provides only services of a given type or only for specific users.
- An AZ contains one or more physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proofing, and electricity facilities. Within an AZ, compute, network, storage, and other resources are logically divided into multiple instances. AZs within a region are interconnected using high-speed optical fibers to allow you to build highly available systems across AZs.

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

Figure 3-1 Regions and AZs



Huawei Cloud provides services in many regions around the world. You can select a region and an AZ as needed. For more information, see [Global Products and Services](#).

Projects

Projects are used to group and isolate OpenStack resources (compute, storage, and network resources). A project can be a department or a project team. Multiple projects can be created for one account.

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 in Huawei Cloud 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. Based on the Kunpeng ecosystem, GaussDB performance is always improved to meet ever-increasing demands in different scenarios.
- **Open-Source Ecosystem**
The centralized version of GaussDB is available for you to download from the open source community.

5 DB Instance Description

5.1 DB Instance Types

[Table 5-1](#) lists the instance types of GaussDB.

Table 5-1 Instance types

Instance Type	Deployment Model	Scale-out Supported	Scenario	Component	Service Processing
Distributed	<p>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.</p> <p>Combined: Different database components are deployed on the same node. This model is only available for instances of the basic edition.</p>	Yes	The data volume is large, and large data capacity and high concurrency are required.	<ul style="list-style-type: none"> • OM • CM • GTM • ETD • CN • DN 	<p>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 set to the CN for summary. Finally, the CN returns the summary result to the application.</p>

Instance Type	Deployment Model	Scale-out Supported	Scenario	Component	Service Processing
<p>Centralized</p>	<p>1 primary + 2 standby: 3-node deployment where there is a shard. The shard contains one primary DN and two standby DNs.</p> <p>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.</p> <p>1 primary + 1 standby + 1 log: 3-node deployment where there is a shard. The shard contains one primary DN, one standby DN, and one log-dedicated DN. The log node is used only to store logs and does not incur fees, which reduces costs and resource consumption.</p>	<p>No</p>	<p>The data volume is small and stable, and data reliability and service availability are extremely important.</p>	<ul style="list-style-type: none"> • OM • CM • ETC • D • DN 	<p>An application sends a task directly to the DN, and the DN returns the result to the application after processing the task.</p>

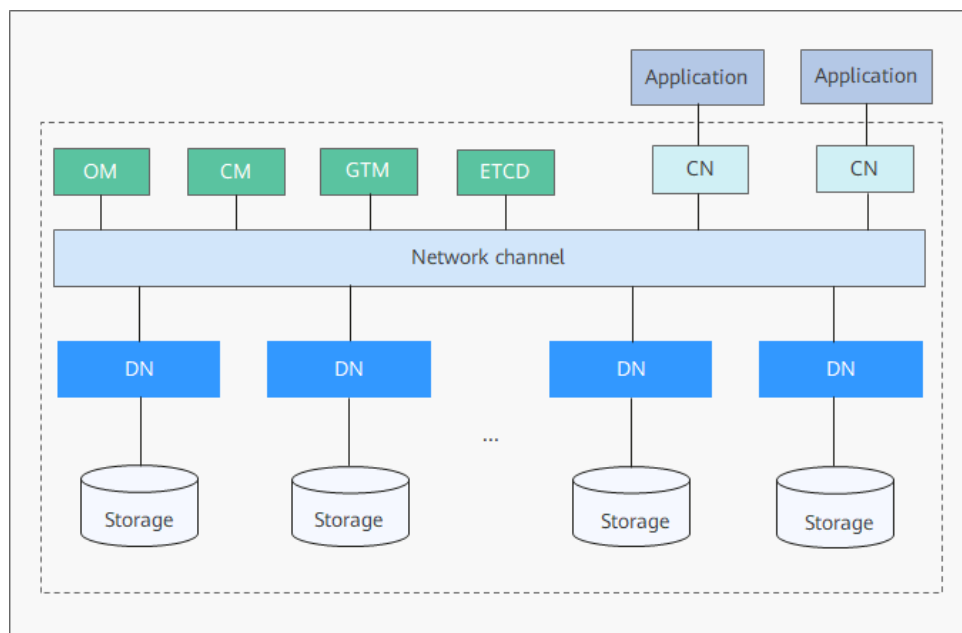
NOTE

- 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. 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 CN receives access requests from applications and returns execution results to clients. It also splits and distributes tasks to different DNs for parallel processing. Each CN can connect to all DNs, and all CNs share the same metadata.
- DN: A DN stores business data, performs data queries, and returns execution results.

Distributed

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

Figure 5-1 Logical architecture of a distributed instance

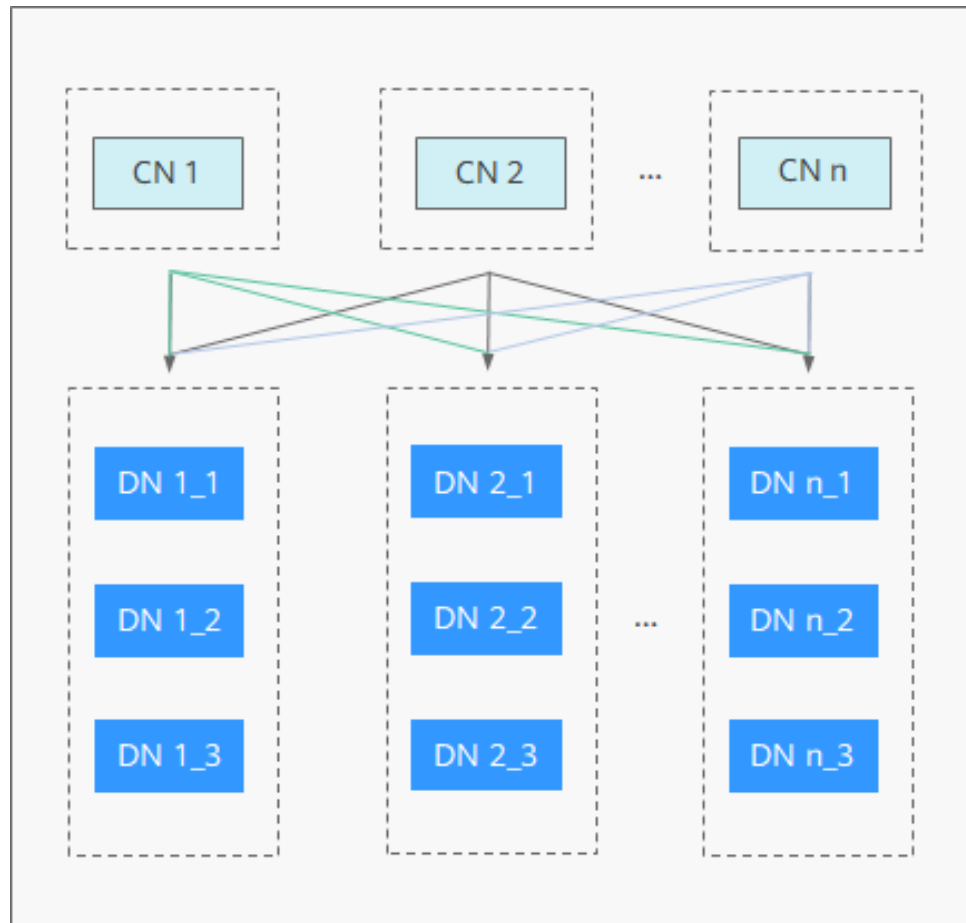


Deployment models for distributed instances

- 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.

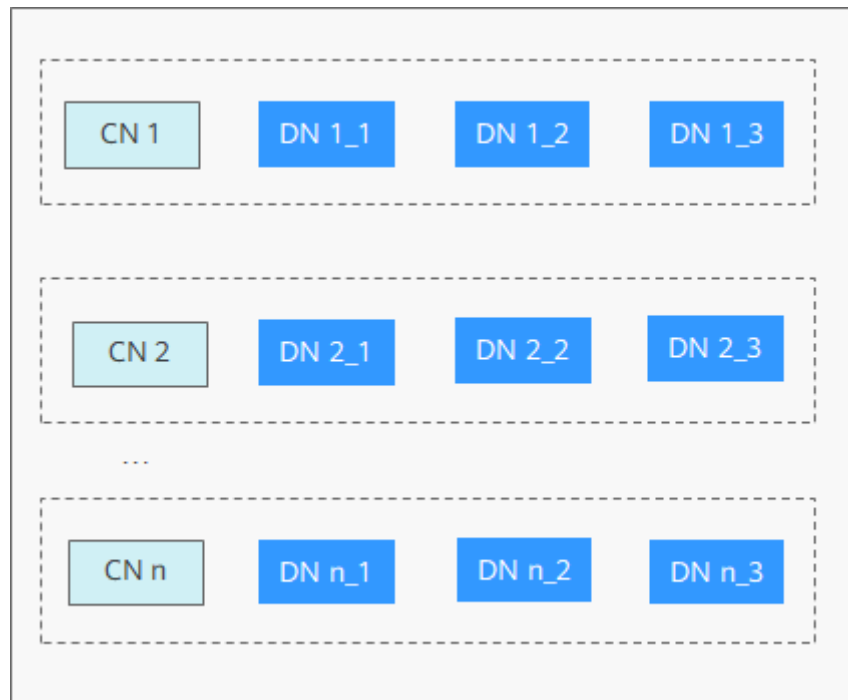
Figure 5-2 Architecture of the independent deployment model



- Combined deployment

In the combined deployment model, multiple database components are deployed on the same server. For example, the default configuration (3 CNs and 3 DN shards) in this model requires three nodes (VMs) in total, and one CN and three DNs share one VM. **Figure 5-3** shows the architecture. This model is only available for distributed instances of the basic edition. Compared with the enterprise edition, the basic edition offers a lower price and is suitable for users who do not require high performance but are sensitive to costs.

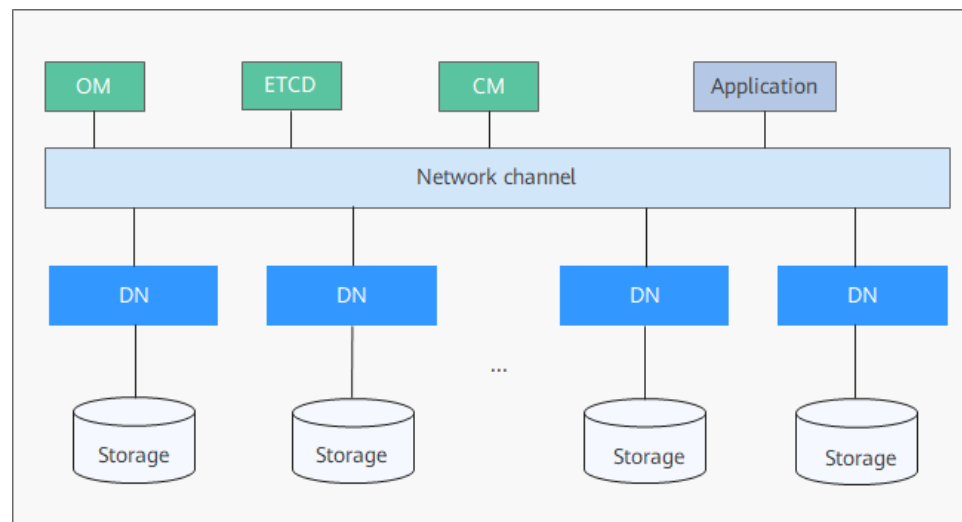
Figure 5-3 Architecture of the combined deployment model



Centralized

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

Figure 5-4 Logical architecture of a centralized instance



Deployment models for centralized instances

- 1 primary + 2 standby

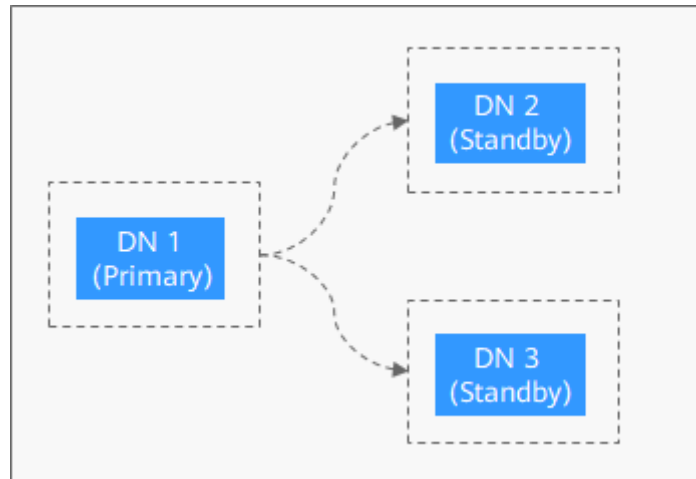
The 1 primary + 2 standby deployment model is a 3-node deployment where there is one shard. The shard contains one primary DN and two standby DNs. [Figure 5-5](#) shows the architecture. This model is only available for centralized instances. Instances deployed using this model are not scalable. This model is

suitable for scenarios with small and stable volumes of data, where data reliability and service availability are important.

NOTE

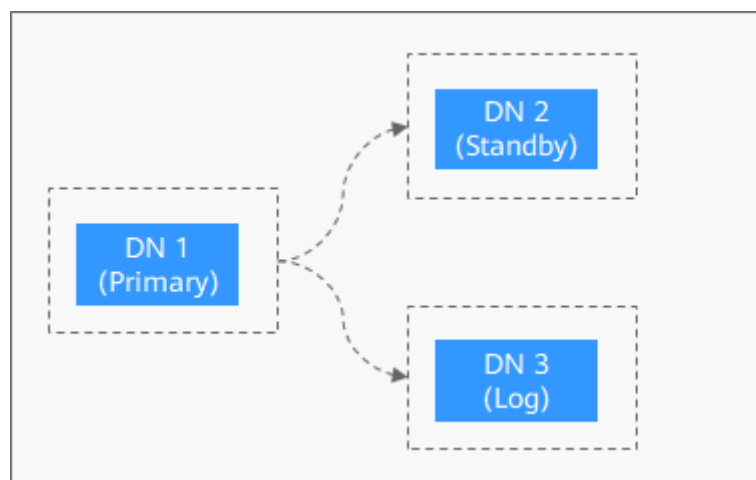
- If any of primary and standby DNs is faulty, services will not be affected.
- 1 primary + 2 standby instances can be deployed in either of the following modes:
 - 3 replicas, single AZ
 - 3 replicas, 3 AZs

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




- 1 primary + 1 standby + 1 log
The 1 primary + 1 standby + 1 log deployment model is a 3-node deployment where there is one shard with three replicas. The shard contains one primary DN, one standby DN, and one log-dedicated DN. This model is only available for centralized instances. [Figure 5-6](#) shows the architecture.

Figure 5-6 Architecture of the 1 primary + 1 standby + 1 log deployment model



- Single (also known as single-replica)
The single-replica deployment model is a single-node deployment where there is only one CMS component and one DN. This model is only available for centralized instances. [Figure 5-7](#) shows the architecture.

Figure 5-7 Architecture of the single-replica deployment model



NOTE

The single-replica deployment model is only available to whitelisted users. You can submit a service ticket to request it at [Service Tickets > Create Service Ticket](#) in the upper right corner of the management console.

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 supports ultra-high I/O and extreme SSD storage types.

- Ultra-high I/O: Ultra-high performance cloud disks excellent for enterprise mission-critical services as well as workloads demanding high throughput and low latency. For details, see the descriptions of "Ultra-high I/O" in [Disk Types and Performance](#).
- Extreme SSD: Superfast disks ideal for workloads demanding ultra-high bandwidth and ultra-low latency. For details, see the descriptions of "Extreme SSD" in [Disk Types and Performance](#).

NOTE

To apply for the permissions needed for using the extreme SSD storage type, submit an application by choosing [Service Tickets > Create Service Ticket](#) in the upper right corner of the management console.

5.3 DB Engine Versions

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.

Mapping Between GaussDB DB Engine and Kernel Engine Versions

Table 5-2 lists the mapping between GaussDB versions.

For details about how to query the DB engine version and kernel engine version of a GaussDB instance, see [Checking the DB Engine and Kernel Engine Version of a GaussDB Instance](#).

Table 5-2 Mapping between GaussDB DB engine and kernel engine versions

DB Engine Version	Kernel Engine Version
V2.0-8.210.0	505.2.1
V2.0-8.201.0	505.2.0
V2.0-3.226.0	503.1.0.SPC2300
V2.0-3.223.0	503.1.0.SPC2000
V2.0-3.222.0	503.1.0.SPC1700
V2.0-3.220.0	503.1.0.SPC1500
V2.0-3.208.0	503.1.0.SPC1200
V2.0-3.207.0	503.1.0.SPC1100
V2.0-3.201.0	503.1.0.SPC0200

Checking the DB Engine and Kernel Engine Version of a GaussDB Instance



1. [Log in to the management console](#).
2. Click  in the upper left corner and select a region and project.
3. Click  in the upper left corner of the page and choose **Databases > GaussDB**.
4. On the **Instances** page, click the name of the target instance to go to the **Basic Information** page.
5. On the **Basic Information** page of the target instance, view the instance version and kernel version in the **DB Engine Version** field and **Kernel Engine Version** field in the **Configuration** area.

Figure 5-8 Basic information



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 5-3 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.
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.
Frozen	The instance is frozen because of an unpaid account balance. Frozen instances are unfrozen only after your account is recharged and the overdue payments are cleared.
Storage full	The storage space of the instance is full. No more data can be written to the databases on this instance. You need to scale up the storage space to restore the instance to be available.
Deleted	The instance has been deleted. Deleted instances will not be displayed in the instance list.
Upgrading	The engine version is being upgraded.

Status	Description
Parameters change. Pending reboot	A modification to a database parameter is waiting for a DB instance reboot before it can take effect.
Changing replicas	The number of replicas for the DB instance is being reduced.

Backup Statuses

Table 5-4 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

GaussDB instances in different regions support x86 or Arm architecture. For details, see [Table 6-1](#). GaussDB instances also support different specifications. For details, see [Table 6-2](#).

- x86: Dedicated (with a vCPU-to-memory ratio of 1:4), dedicated (with a vCPU-to-memory ratio of 1:8), and general-purpose (with a vCPU-to-memory ratio of 1:4)
- Arm: Kunpeng dedicated (with a vCPU-to-memory ratio of 1:4) and Kunpeng dedicated (with a vCPU-to-memory ratio of 1:8)

Table 6-1 Supported regions

Edition Type	Instance Type	Specification Type	Billing Mode	Region
Enterprise edition	Distributed	<ul style="list-style-type: none">• Dedicated (1:4)• Dedicated (1:8)• Kunpeng dedicated (1:4)	Pay-per-use and yearly/monthly	ME-Riyadh TR-Istanbul AP-Singapore LA-Sao Paulo1 AF-Johannesburg AP-Jakarta AP-Bangkok CN-Hong Kong LA-Mexico City2 CN Southwest-Guiyang1 LA-Santiago CN North-Beijing4 RU-Moscow2 CN East 2

Edition Type	Instance Type	Specification Type	Billing Mode	Region
		Kunpeng dedicated (1:8)	Pay-per-use	TR-Istanbul AP-Singapore AF-Johannesburg LA-Sao Paulo1 AP-Jakarta AP-Bangkok CN-Hong Kong LA-Mexico City2 CN Southwest-Guiyang1 LA-Santiago CN North-Beijing4 RU-Moscow2 CN East 2
			Yearly/ Monthly	AP-Singapore LA-Sao Paulo1 AF-Johannesburg AP-Jakarta AP-Bangkok CN-Hong Kong LA-Mexico City2 CN Southwest-Guiyang1 LA-Santiago CN North-Beijing4 RU-Moscow2 CN East 2

Edition Type	Instance Type	Specification Type	Billing Mode	Region
	Centralized	<ul style="list-style-type: none"> • Dedicated (1:4) • Dedicated (1:8) • Kunpeng dedicated (1:4) • Kunpeng dedicated (1:8) 	Pay-per-use and yearly/monthly	ME-Riyadh TR-Istanbul AP-Singapore AF-Johannesburg LA-Sao Paulo1 AP-Bangkok AP-Jakarta CN-Hong Kong LA-Mexico City2 LA-Santiago CN Southwest-Guiyang1 CN North-Beijing4 RU-Moscow2 CN East 2
Basic edition	Distributed	<ul style="list-style-type: none"> • Dedicated (1:4) • Dedicated (1:8) • Kunpeng dedicated (1:4) 	Pay-per-use	ME-Riyadh AF-Johannesburg LA-Sao Paulo1 CN-Hong Kong CN South-Guangzhou LA-Santiago AP-Singapore TR-Istanbul CN Southwest-Guiyang1 AP-Bangkok LA-Mexico City2 CN East-Shanghai1 CN North-Beijing4 AP-Jakarta CN East 2
		Kunpeng dedicated (1:8)	Pay-per-use	CN South-Guangzhou CN East-Shanghai1 CN Southwest-Guiyang1 CN North-Beijing4 CN East 2

Edition Type	Instance Type	Specification Type	Billing Mode	Region
	Centralized	<ul style="list-style-type: none"> General-purpose (1:4) Dedicated (1:4) Dedicated (1:8) Kunpeng dedicated (1:4) 	Pay-per-use	ME-Riyadh AP-Jakarta AF-Johannesburg CN South-Guangzhou LA-Sao Paulo1 LA-Mexico City2 CN-Hong Kong CN East-Shanghai1 CN North-Beijing4 CN Southwest-Guiyang1 TR-Istanbul AP-Singapore AP-Bangkok LA-Santiago CN East 2
		Kunpeng dedicated (1:8)	Pay-per-use	CN East-Shanghai1 CN South-Guangzhou CN Southwest-Guiyang1 CN North-Beijing4 CN East 2

Table 6-2 Instance specifications

Edition Type	DB Instance Type	Specification Type	vCPUs	Memory (GB)	Default Maximum Connections (Per CN)
Enterprise edition	Distributed	Dedicated (1:4) NOTE Dedicated (1:4) is based on the x86 architecture.	8	32	100
			16	64	1,000
			24	96	3,000
			32	128	3,000
			48	192	3,000
			64	256	4,000
			96	384	6,000

Edition Type	DB Instance Type	Specification Type	vCPUs	Memory (GB)	Default Maximum Connections (Per CN)	
		Dedicated (1:8) NOTE Dedicated (1:8) is based on the x86 architecture.	8	64	1,000	
			16	128	2,000	
			24	192	3,000	
			32	256	4,000	
			48	384	6,000	
			64	512	8,000	
		Kunpeng dedicated (1:4) NOTE Kunpeng dedicated (1:4) is based on the Arm architecture.	8	32	100	
			16	64	1,000	
			24	96	3,000	
			32	128	2,000	
			48	192	3,000	
			64	256	4,000	
		Kunpeng dedicated (1:8) NOTE Kunpeng dedicated (1:8) is based on the Arm architecture.	96	384	3,000	
			8	64	1,000	
			16	128	2,000	
			24	192	3,000	
			32	256	4,000	
		Centralized	Dedicated (1:4) NOTE Dedicated (1:4) is based on the x86 architecture.	48	384	6,000
				60	480	8,000
				4	16	300
				8	32	600
	16			64	2,048	
	24			96	2,048	
	32			128	5,000	
	48	192	3,000			
			96	384	6,000	

Edition Type	DB Instance Type	Specification Type	vCPUs	Memory (GB)	Default Maximum Connections (Per CN)
		Dedicated (1:8) NOTE Dedicated (1:8) is based on the x86 architecture.	4	32 NOTE This specification is not available for production environments, but you can submit a service ticket to request it at Service Tickets > Create Service Ticket in the upper right corner of the management console.	600
			8	64 NOTE Only centralized instances of version V2.0-2.6 or later can be used in the production environment.	2,048
			16	128	5,000
			24	192	3,000
			32	256	11,000
			48	384	6,000
			64	512	25,000
			Kunpeng dedicated (1:4) NOTE Kunpeng dedicated (1:4) is based on the Arm architecture.	4	16
		8		32	600
		16		64	2,048
		24		96	2,048
		32		128	5,000
		48		192	3,000

Edition Type	DB Instance Type	Specification Type	vCPUs	Memory (GB)	Default Maximum Connections (Per CN)		
			96	384	6,000		
		Kunpeng dedicated (1:8) NOTE Kunpeng dedicated (1:8) is based on the Arm architecture.	4	32 NOTE This specification is not available for production environments, but you can submit a service ticket to request it at Service Tickets > Create Service Ticket in the upper right corner of the management console.	600		
			8	64 NOTE Only centralized instances of version V2.0-2.6 or later can be used in the production environment.	2,048		
			16	128	5,000		
			24	192	3,000		
			32	256	11,000		
			48	384	6,000		
			60	480	24,000		
		Basic edition	Distributed	Dedicated (1:4) NOTE Dedicated (1:4) is based on the x86 architecture.	16	64	200
					32	128	350
64	256				900		

Edition Type	DB Instance Type	Specification Type	vCPUs	Memory (GB)	Default Maximum Connections (Per CN)
		Dedicated (1:8) NOTE Dedicated (1:8) is based on the x86 architecture.	16	128	350
			32	256	900
			64	512	2,000
		Kunpeng dedicated (1:4) NOTE Kunpeng dedicated (1:4) is based on the Arm architecture.	16	64	200
			32	128	350
		Kunpeng dedicated (1:8) NOTE Kunpeng dedicated (1:8) is based on the Arm architecture.	16	128	350
			32	256	900
			60	480	1,800
		Centralized	General-purpose (1:4) NOTE General-purpose (1:4) is based on the x86 architecture.	4	16
			8	32	600
		Dedicated (1:4) NOTE Dedicated (1:4) is based on the x86 architecture.	4	16	300
			8	32	600
			16	64	2,048
			32	128	5,000
		64	256	11,000	
		96	384	11000	
		128	512	25,000	

Edition Type	DB Instance Type	Specification Type	vCPUs	Memory (GB)	Default Maximum Connections (Per CN)
		Kunpeng dedicated (1:4) NOTE Kunpeng dedicated (1:4) is based on the Arm architecture.	4	16	300
			8	32	600
			16	64	2,048
			32	128	5000
			96	384	11000
			128	512	25,000
		Dedicated (1:8) NOTE Dedicated (1:8) is based on the x86 architecture.	8	64	2,048
			16	128	5,000
			32	256	11,000
			64	512	25,000
			96	768	40000
			128	1,024	55,000
		Kunpeng dedicated (1:8) NOTE Kunpeng dedicated (1:8) is based on the Arm architecture.	8	64	2,048
			16	128	5,000
			32	256	11,000
			96	768	40000
			128	1,024	55,000

7 Edition Types

GaussDB offers two editions: basic edition and enterprise edition. The enterprise edition distinguishes itself with a suite of advanced features, including advanced compression, dynamic data masking, and anti-tampering, while the basic edition includes only advanced compression. The basic edition of GaussDB is positioned as a cost-effective option but does not include certain advanced features. The basic edition delivers the same level of performance as the enterprise edition at a more affordable price. This edition is ideal for users who want to balance affordability with reliable performance and do not require the extra features.

For more details about the differences between the two editions, see [Table 7-1](#).

Table 7-1 Differences between the basic and enterprise editions of GaussDB instances

Edition Type	Basic edition	Enterprise edition
Supported Instance Type	<ul style="list-style-type: none"> • Centralized • Distributed 	<ul style="list-style-type: none"> • Centralized • Distributed
Supported Deployment Model	<ul style="list-style-type: none"> • Centralized <ul style="list-style-type: none"> - 1 primary + 2 standby - 1 primary + 1 standby + 1 log - Single • Distributed • Combined 	<ul style="list-style-type: none"> • Centralized <ul style="list-style-type: none"> - 1 primary + 2 standby - 1 primary + 1 standby + 1 log - Single • Distributed • Independent
Available Instance Specifications	Supported specification types: <ul style="list-style-type: none"> • General-purpose • Dedicated • Kunpeng dedicated 	Supported specification types: <ul style="list-style-type: none"> • Dedicated • Kunpeng dedicated

	<p>Minimum specifications:</p> <ul style="list-style-type: none"> Centralized 4 vCPUs 16 GB (General-purpose, Dedicated, and Kunpeng dedicated) Distributed Combined deployment: 16 vCPUs 64 GB (Dedicated and Kunpeng dedicated) <p>NOTE For a combined deployment, all components must be deployed on a single ECS, which demands higher CPU resources to maintain optimal performance. As a result, a minimum of 16 vCPUs is necessary for a combined deployment.</p>	<p>Minimum specifications:</p> <ul style="list-style-type: none"> Centralized 4 vCPUs 16 GB (Dedicated and Kunpeng dedicated) Distributed Independent deployment: 8 vCPUs 64 GB (Dedicated and Kunpeng dedicated)
	<p>Maximum specifications:</p> <ul style="list-style-type: none"> x86: 128 vCPUs and 1024 GB memory Arm: 128 vCPUs and 1024 GB memory 	<p>Maximum specifications:</p> <ul style="list-style-type: none"> x86: 64 vCPUs 512 GB Arm: 60 vCPUs 480 GB
Supported Advanced Feature	Advanced compression	Advanced compression, dynamic data masking, and anti-tampering
Performance (Benchmark : Enterprise Edition)	<ul style="list-style-type: none"> Centralized instance: same as the enterprise edition Distributed instance in a combined deployment: 50% of the performance of the enterprise edition, when the same specifications are configured 	-

Price (Benchmark : Enterprise Edition)	<ul style="list-style-type: none">• When the specifications are the same, the cost of a centralized instance in the basic edition is about 30% of the cost of the enterprise edition.• When the specifications are the same, the cost of a distributed instance in a combined deployment of the basic edition is about 7.5% of the cost of the enterprise edition.	-
---	---	---

8 Security

8.1 Shared Responsibilities

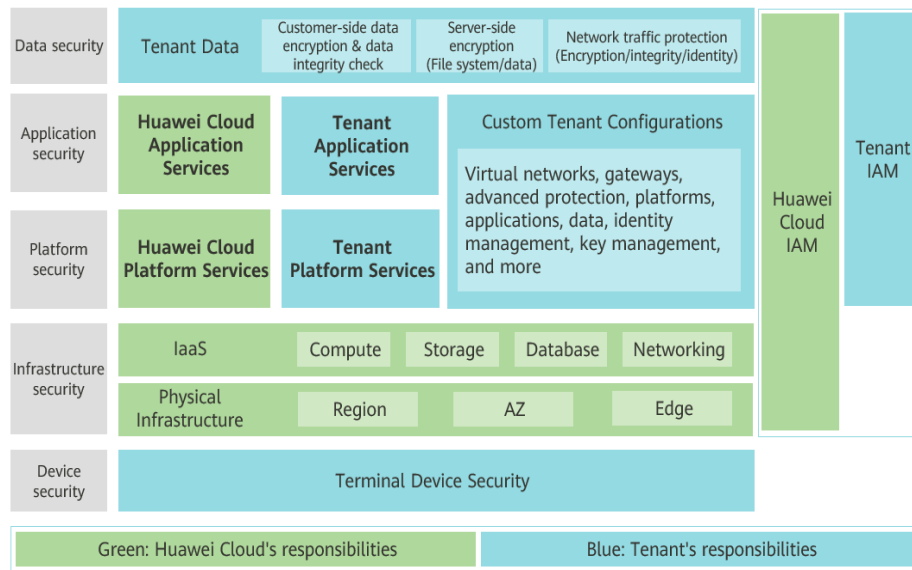
Huawei guarantees that its commitment to cyber security will never be outweighed by the consideration of commercial interests. To cope with emerging cloud security challenges and pervasive cloud security threats and attacks, Huawei Cloud builds a comprehensive cloud service security assurance system for different regions and industries based on Huawei's unique software and hardware advantages, laws, regulations, industry standards, and security ecosystem.

Figure 8-1 illustrates the responsibilities shared by Huawei Cloud and users.

- **Huawei Cloud:** Ensure the security of cloud services and provide secure clouds. Huawei Cloud's security responsibilities include ensuring the security of our IaaS, PaaS, and SaaS services, as well as the physical environments of the Huawei Cloud data centers where our IaaS, PaaS, and SaaS services operate. Huawei Cloud is responsible for not only the security functions and performance of our infrastructure, cloud services, and technologies, but also for the overall cloud O&M security and, in the broader sense, the security and compliance of our infrastructure and services.
- **Tenant:** Use the cloud securely. Tenants of Huawei Cloud are responsible for the secure and effective management of the tenant-customized configurations of cloud services including IaaS, PaaS, and SaaS. This includes but is not limited to virtual networks, the OS of virtual machine hosts and guests, virtual firewalls, API Gateway, advanced security services, all types of cloud services, tenant data, identity accounts, and key management.

Huawei Cloud Security White Paper elaborates on the ideas and measures for building Huawei Cloud security, including cloud security strategies, the shared responsibility model, compliance and privacy, security organizations and personnel, infrastructure security, tenant service and security, engineering security, O&M security, and ecosystem security.

Figure 8-1 Huawei Cloud shared security responsibility model



8.2 Identity Authentication and Access Control

8.2.1 Service Access Control

Identity Authentication

When you access GaussDB, the system authenticates your identity using password and IAM authentication.

- **Password verification**

To manage your instance, you need to use Data Admin Service (DAS) to log in to your instance. The login is successful only after your account and password are verified.

- **IAM verification**

You can use **Identity and Access Management (IAM)** to provide fine-grained control of GaussDB permissions. IAM provides identity authentication, permissions management, and access control, helping you to securely access your Huawei Cloud resources. IAM users can use GaussDB resources only after their accounts and passwords are verified. For details, see [Step 2: Create IAM Users and Log In](#).

Access Control

- **Permissions control**

If you need to assign different permissions to different employees in your enterprise to access your DB instance resources, IAM is a good choice. For details, see [Permissions Management](#).

- **VPCs and subnets**

A VPC is a logically isolated, configurable, and manageable virtual network. It helps improve the security of cloud resources and simplifies network

deployment. You can define security groups, virtual private networks (VPNs), IP address segments, and bandwidth for a VPC. This facilitates internal network configuration and management and allows you to change your network in a secure and convenient network manner.

A subnet provides dedicated network resources that are logically isolated from other networks for security.

For details, see [Creating a VPC](#).

- **Security groups**

A security group is a logical group that provides access control policies for the GaussDB instances that have the same security requirements and are mutually trusted in a VPC. To ensure database security and reliability, you need to configure security group rules to allow only specific IP addresses and ports to access your GaussDB instances.

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

8.3 Data Protection

GaussDB provides a series of methods and features to ensure data security and reliability.

Table 8-1 Methods for data security

Method	Description	Reference
Transmission encryption (HTTPS)	HTTP and HTTPS are both supported, but HTTPS is recommended for enhanced security.	Making an API Request
Data backup	You can back up and restore databases to ensure data reliability.	Working with Backups
Critical operation protection	With this function enabled, the system authenticates user's identity when they perform any risky operations like deleting an instance. This enhances the protection for your data and configuration.	Introduction to Sensitive Operation Protection
SSL	You can use SSL to encrypt the connection between GaussDB and the client. It ensures the privacy, authentication, and integrity of Internet communications.	SSL Connection
Deletion protection	GaussDB moves all deleted instances to the recycle bin. You can restore a DB instance that was deleted up to 7 days ago from the recycle bin.	Rebuilding a Deleted Instance

8.4 Audit and Logs

Audit

Cloud Trace Service (CTS) records operations on the cloud resources in your account. You can use the logs generated by CTS to perform security analysis, track resource changes, audit compliance, and locate faults.

For details about how to enable and configure CTS, see [Enabling CTS](#).

With CTS, you can record operations associated with GaussDB for future query, audit, and backtracking. For details, see [Key Operations Supported by CTS](#).

Logs

GaussDB provides a variety of log types and functions for database analysis or audit. You can download logs on the management console.

- Error logs
GaussDB allows you to download error logs. Error logs are generated when the database is running. These logs can help you analyze problems with the database. For details about error logs, see [Error Logs](#).
- Slow query logs
You can download slow query logs to locate slow SQL statement execution problems and tune the database accordingly. For details about slow query logs, see [Slow Query Logs](#).

8.5 Risk Monitoring

Cloud Eye is a comprehensive monitoring platform for resources like cloud databases and cloud servers. It enables you to monitor resources, configure alarm rules, identify resource exceptions, and quickly respond to resource changes.

Monitoring Metrics

You can monitor resources and operations, such as CPU usage and network throughput using Cloud Eye. For details about supported monitoring metrics and how to create alarm rules, see [Monitoring Metrics](#).

The monitoring interval can be 1 minute, or 10 seconds. The default monitoring interval is 1 minute. To improve the accuracy of monitoring metrics, you can enable **Monitoring by Seconds**.

Event Monitoring

Event monitoring provides reporting, query, and alarm functions for event data. You can create alarm rules for both system events and custom events. When specific events occur, Cloud Eye generates alarms for you.

8.6 Fault Recovery

Automated backups are created during the backup time window of your GaussDB instances. GaussDB saves automated backups based on the retention period (1 to 732 days) you specified. For details, see [Configuring an Automated Backup Policy](#).

Based on your service requirements, you can:

- [Restore data to the original instance, a new instance or an existing instance.](#)
- [Restoring a DB Instance to a Specific Point in Time.](#)

Multiple-AZ Deployment

An AZ is a physical region where resources have their own independent power supply and networks. AZs are physically isolated but interconnected through an internal network. GaussDB supports multiple-AZ deployment for cross-AZ DR.

Failover

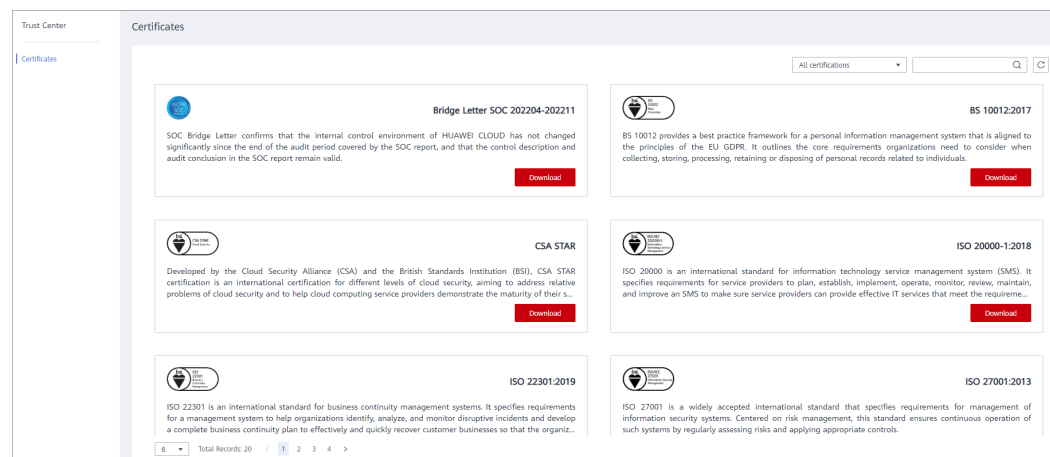
GaussDB is a multi-node instance. By default, a centralized instance has only one primary node; a distributed instance has only one primary DN and multiple standby DNs. If the primary DN becomes unavailable, GaussDB automatically fails over to a standby DN.

8.7 Certificates

Compliance Certificates

Huawei Cloud services and platforms have obtained various security and compliance certifications from authoritative organizations, such as International Organization for Standardization (ISO), system and organization controls (SOC), and Payment card industry (PCI) compliance standards. These certifications are available for [download](#).


Figure 8-2 Downloading compliance certificates



Resource Center

Huawei Cloud also provides the following resources to help users meet compliance requirements. For details, see [Resource Center](#).

Figure 8-3 Resource center



9 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.

IAM can be used for free. You pay only for the resources in your account. For more information about IAM, see [IAM Service Overview](#).

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 role-based 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 9-1 lists all the system-defined policies supported by GaussDB.

Table 9-1 System-defined permissions for GaussDB

Policy Name	Description	Category	Dependency Configuration
GaussDB FullAccess	Full permissions for GaussDB	System-defined policy	<p>To use storage autoscaling, create DR tasks, reset DR configurations, and rectify agency permissions, configure the following actions for IAM users:</p> <ul style="list-style-type: none"> • Actions required for creating a custom policy: <ul style="list-style-type: none"> - iam:agencies:listAgencies - iam:agencies:createAgency - iam:permissions:listRolesForAgencyOnProject - iam:permissions:grantRoleToGroupOnProject - iam:roles:listRoles - iam:permissions:listRolesForAgencyOnDomain - iam:permissions:revokeRoleFromAgencyOnProject - iam:permissions:revokeRoleFromAgencyOnDomain • Adding system role Security Administrator: <ol style="list-style-type: none"> 1. Select a user group to which the user belongs. 2. Click Authorize in the Operation column. 3. Add the Security Administrator role. <p>Actions required for creating a yearly/monthly instance using a RAM-based shared KMS key:</p>

Policy Name	Description	Category	Dependency Configuration
			<ul style="list-style-type: none"> iam:agencies:listAgencies iam:roles:listRoles iam:agencies:pass iam:agencies:createAgency iam:permissions:grantRoleToAgency <p>GaussDB FullAccess already contains the iam:agencies:listAgencies, iam:roles:listRoles, and iam:agencies:pass actions.</p> <p>GaussDB is a region-level service, and IAM is a global service. If you want to grant GaussDB FullAccess to a project, grant BSS ServiceAgencyReadPolicy (global service) to it as well. Granting GaussDB FullAccess to all projects eliminates the need for additional configuration when using IAM actions.</p> <p>BSS ServiceAgencyCreatePolicy contains the following actions: iam:agencies:createAgency and iam:permissions:grantRoleToAgency.</p>
GaussDB ReadOnlyAccess	Read-only permissions for GaussDB	System-defined policy	None

Table 9-2 lists the common operations supported by each system policy of GaussDB. Choose appropriate system policies based on this table.

Table 9-2 Common operations supported by the GaussDB system policies

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

Table 9-3 Common operations and supported actions

Operation	Action	Remarks
Creating a DB instance	gaussdb:instance:create gaussdb:param:list	<p>To select a VPC, subnet, and security group, configure the following actions:</p> <p>vpc:vpcs:list vpc:vpcs:get vpc:subnets:get vpc:securityGroups:get</p> <p>To create a yearly/monthly instance, configure the following actions:</p> <p>bss:order:update bss:order:view bss:balance:view</p> <p>To create an encrypted instance, configure the following actions for the project:</p> <p>kms:cmk:get kms:cmk:list</p> <p>To report event monitoring of a failed operation, configure the following actions:</p> <p>"ces:alarmsOnOff:put" "ces:alarms:create"</p>

Operation	Action	Remarks
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"
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 unsubscribe from a yearly/monthly instance, configure the following actions: "bss:unsubscribe:update" To report event monitoring of a failed operation, configure the following actions: "ces:alarmsOnOff:put" "ces:alarms:create"
Querying instances	gaussdb:instance:list	None

Operation	Action	Remarks
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:publiclips:get vpc:publiclips: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
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

Operation	Action	Remarks
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:modifyBackupPolicy	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"
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

Operation	Action	Remarks
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
Viewing metrics	-	ces:metric:listGroups ces:metric:listConfig

Operation	Action	Remarks
Enabling autoscaling	gaussdb:instance:autoEnlargePolicy	<p>To enable autoscaling, configure the following actions for the IAM users instead of your Huawei account:</p> <ul style="list-style-type: none"> iam:agencies:listAgencies iam:agencies:createAgency iam:permissions:listRolesForAgencyOnProject iam:permissions:grantRoleToGroupOnProject iam:roles:listRoles iam:permissions:listRolesForAgencyOnDomain iam:permissions:revokeRoleFromAgencyOnProject iam:permissions:revokeRoleFromAgencyOnDomain Adding system role Security Administrator: <ol style="list-style-type: none"> Select a user group to which the user belongs. Click Authorize in the Operation column. Add the Security Administrator role.
Querying advanced features	gaussdb:instance:listFeatures	None
Enabling advanced features	gaussdb:instance:updateFeatures	None

Table 9-4 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

Operation	Action	Remarks
Resetting the DR relationship	gaussdb:disasterRecovery:construct gaussdb:disasterRecovery:create	<p>To enable the DR operation, configure the following actions for the IAM users instead of your Huawei account:</p> <ul style="list-style-type: none"> • iam:agencies:listAgencies • iam:agencies:createAgency • iam:permissions:listRolesForAgencyOnProject • iam:permissions:grantRoleToGroupOnProject • iam:roles:listRoles • iam:permissions:listRolesForAgencyOnDomain • iam:permissions:revokeRoleFromAgencyOnProject • iam:permissions:revokeRoleFromAgencyOnDomain • Adding system role Security Administrator: <ol style="list-style-type: none"> 1. Select a user group to which the user belongs. 2. Click Authorize in the Operation column. 3. Add the Security Administrator role. <p>An agency RDSAccessProjectResource (including the policy DBS AgencyPolicy) will be automatically created in the current region.</p>

Operation	Action	Remarks
Establishing a DR relationship	gaussdb:disasterRecovery:construct gaussdb:disasterRecovery:create	<p>To enable the DR operation, configure the following actions for the IAM users instead of your Huawei account:</p> <ul style="list-style-type: none"> • iam:agencies:listAgencies • iam:agencies:createAgency • iam:permissions:listRolesForAgencyOnProject • iam:permissions:grantRoleToGroupOnProject • iam:roles:listRoles • iam:permissions:listRolesForAgencyOnDomain • iam:permissions:revokeRoleFromAgencyOnProject • iam:permissions:revokeRoleFromAgencyOnDomain • Adding system role Security Administrator: <ol style="list-style-type: none"> 1. Select a user group to which the user belongs. 2. Click Authorize in the Operation column. 3. Add the Security Administrator role. <p>An agency RDSAccessProjectResource (including the policy DBS AgencyPolicy) will be automatically created in the current region.</p>

Operation	Action	Remarks
Promoting the DR instance to primary	gaussdb:disasterRecovery:failover	The RDSAccessProjectResource agency is required.
Removing a DR relationship	gaussdb:disasterRecovery:release	The RDSAccessProjectResource agency is required.
Switching roles of primary and DR instances	gaussdb:disasterRecovery:switchover	The RDSAccessProjectResource agency is required.
Re-establishing a DR relationship	gaussdb:disasterRecovery:construct gaussdb:disasterRecovery:create	The RDSAccessProjectResource agency is required.
Performing a DR drill	gaussdb:disasterRecovery:simulation	The RDSAccessProjectResource agency is required.
Caching logs for DR	gaussdb:disasterRecovery:keeplog	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.

10 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 10-6](#).

Specifications

Table 10-1 Specification description

Resource Type	Specifications	Description
Storage space	<ul style="list-style-type: none">Centralized instances: 40 GB to 24,000 GBDistributed instances: 40 GB to 6,144,000 GB	Each shard has a maximum storage capacity of 24 TB. Adding a new shard increases the total storage capacity by an additional 24 TB. For a distributed instance in independent deployment, up to 256 shards can be added.
Connections	The minimum number is 10 and the theoretical maximum number is 262,143. The actual maximum value dynamically changes.	The maximum number of connections supported depends on memory specifications. For details, see Configuration of Maximum Connections .

Resource Type	Specifications	Description
IOPS	<ul style="list-style-type: none"> Ultra-high I/O: a maximum of 50,000 Extreme SSD: a maximum of 128,000 	The input/output operations per second (IOPS) supported depends on the I/O performance of Elastic Volume Service (EVS) disks. For details, see the description about ultra-high I/O and extreme SSDs in Disk Types and Performance of Elastic Volume Service Service Overview .

Quotas

Table 10-2 Quota limits

Resource Type	Constraint	Description
Tags	Each instance can have up to 20 tags.	For more information, see Tags .
Free backup space	GaussDB provides free backup space of the same size as your purchased storage space.	After you pay for the storage space of your instance, you will get a backup space of the same size for free. For more information, see How Is GaussDB Backup Data Billed?
Retention period of automated backups	The default value is 7 days. The value ranges from 1 to 732 days.	For more information, see Configuring an Automated Backup Policy for Instances .

Naming

Table 10-3 Naming constraints

Item	Description
Instance name	<ul style="list-style-type: none"> 4 to 64 characters long Must start with a letter. Only letters (case-sensitive), digits, hyphens (-), and underscores (_) are allowed.
Backup name	<ul style="list-style-type: none"> 4 to 64 characters long Must start with a letter. Only letters (case-sensitive), digits, hyphens (-), and underscores (_) are allowed.

Item	Description
Parameter template name	<ul style="list-style-type: none"> 1 to 64 characters long Only letters (case-sensitive), digits, hyphens (-), underscores (_), and periods (.) are allowed.

Security

Table 10-4 Security constraints

Item	Description
root permissions	<p>Only the root user permissions are available on the instance creation page.</p> <p>NOTE After August 30, 2022, the root user of newly created instances and upgraded existing instances has the sysadmin permission. To upgrade your existing instance, contact technical support.</p> <p>If your existing instances are not upgraded, the root user only has the following permissions: createrole, createdb, and monadmin. The root user has fewer permissions than the administrator. When some SQL syntax or functions, such as CREATE TABLESPACE, are executed, an error message is displayed, indicating that the permissions are insufficient.</p>
root password	<ul style="list-style-type: none"> 8 to 32 characters long Must contain at least three types of the following: uppercase letters, lowercase letters, digits, and special characters. Supported characters: ~!@#%^*_-=+?, <p>For more information, see Resetting the Administrator Password.</p>
Database port	<p>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.</p> <p>For more information, see Changing the Database Port.</p>
Disk encryption	<p>If you enable disk encryption during instance creation, the disk encryption status and the key cannot be changed later.</p>
VPC	<p>The VPC where a GaussDB instance is located cannot be changed after the instance is created.</p>

Item	Description
Security group	<ul style="list-style-type: none"> 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. For more information, see Configuring Security Group Rules.
System accounts	<p>To provide O&M services, the system automatically creates system accounts when you create GaussDB instances. These accounts cannot be used by users.</p> <ul style="list-style-type: none"> 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 10-5 Constraints on instance operations

Function	Constraints
Database access	<ul style="list-style-type: none"> To access a GaussDB instance that is not publicly accessible, the instance must be in the same VPC subnet as the ECSs. 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. The database port can be set when a DB instance is created and can be changed after the DB instance is created.
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.

Function	Constraints
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	<ul style="list-style-type: none"> • By default, the specifications cannot be reduced. If you need to reduce the specifications, contact Huawei Cloud technical support. • 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.
Recycle bin management	<ul style="list-style-type: none"> • 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 10-6 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
Restoring data to a new instance	No	Yes
Restoring data to an existing instance	No	Yes
Scaling up storage	Yes	Yes

Function	Supported by Versions Earlier Than V2.0-3.0	Supported by V2.0-3.0 and Later
Changing instance specifications	Yes	Yes
Hot patch installation	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
Deleting a schema and user	Yes	Yes
Performing database operations	Yes	Yes

11 Related Services

Table 11-1 shows the relationship between GaussDB and other services.

Table 11-1 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.

12 Compatibility

This section describes the compatibility of GaussDB in Oracle-compatible mode with the native Oracle database, and that of GaussDB in MySQL-compatible mode with the native MySQL database.

Compatibility with Oracle Database

GaussDB is basically compatible with Oracle Database in terms of basic functions (such as data types, SQL statements, and database objects) and PL/SQL. However, due to architecture design differences, there are still some incompatible items. For details about the differences between GaussDB in Oracle-compatible mode and Oracle Database 19c, see the following sections:

- [Basic SQL Elements](#)
- [Pseudocolumns](#)
- [Operators](#)
- [Expressions](#)
- [Conditions](#)
- [Common SQL DDL Clauses](#)
- [SQL Queries and Subqueries](#)
- [PL/SQL Language](#)
- [System Functions](#)
- [System Views](#)
- [Advanced Packages](#)

Compatibility with MySQL

GaussDB is compatible with MySQL in terms of data types, SQL functions, and database objects. Due to differences in the underlying framework implementations, there are still some distinctions between GaussDB and MySQL.

- For details about the differences between GaussDB in MySQL-compatible mode (with **sql_compatibility** set to 'MYSQL', **b_format_version** set to '5.7', and **b_format_dev_version** set to 's1') and MySQL 5.7, see the following sections:

- [Data Types](#)
- [System Functions](#)
- [Operators](#)
- [Character Sets](#)
- [Collation Rules](#)
- [SQL](#)
- [Drivers](#)
- For details about the differences between GaussDB in MySQL-compatible mode (with **sql_compatibility** set to 'M') and MySQL 5.7, see the following sections:
 - [Data Types](#)
 - [System Functions](#)
 - [Operators](#)
 - [Character Sets](#)
 - [Collation Rules](#)
 - [Transactions](#)
 - [SQL](#)