

GeminiDB Cassandra

Issue 02
Date 2023-03-27



Copyright © Huawei Technologies Co., Ltd. 2023. 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 Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are trademarks 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 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.

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to "Vul. Response Process". For details about the policy, see the following website:<https://www.huawei.com/en/psirt/vul-response-process>
For enterprise customers who need to obtain vulnerability information, visit:<https://securitybulletin.huawei.com/enterprise/en/security-advisory>

Contents

1 Service Overview.....	1
1.1 What Is GeminiDB Cassandra API?.....	1
1.2 Compatible APIs and Versions.....	2
1.3 Instance Specifications.....	2
1.4 Instance Statuses.....	3
1.5 Database Constraints.....	4
1.5.1 Basic Design.....	5
1.5.2 DB Instance Specifications and Performance.....	7
1.5.3 Database Objects.....	7
1.5.4 Access and Connection Pools.....	9
1.5.5 Batches.....	9
1.5.6 Queries.....	9
2 Getting Started with GeminiDB Cassandra API.....	10
2.1 Service Overview.....	10
2.2 Buying an Instance.....	11
2.3 Instance Connections.....	18
2.3.1 Connection Methods.....	18
2.3.2 Connecting to an Instance over a Private Network.....	18
2.3.3 Connecting to an Instance over a Public Network.....	20
2.3.4 Connecting to an Instance Using Java.....	22
2.3.5 Connecting to an Instance Using Go.....	23
3 Working with GeminiDB Cassandra API.....	26
3.1 Permissions Management.....	26
3.1.1 Creating a User and Assigning Permissions.....	26
3.1.2 Creating a Custom Policy.....	27
3.2 Instance Lifecycle.....	29
3.2.1 Restarting an Instance.....	29
3.2.2 Deleting Instance.....	29
3.2.3 Recycling an Instance.....	30
3.3 Instance Modifications.....	31
3.3.1 Upgrading a Minor Version.....	31
3.3.2 Changing an Instance Name.....	33

3.3.3 Resetting the Administrator Password.....	34
3.3.4 Scaling Up Storage Space.....	35
3.3.5 Configuring Autoscaling.....	36
3.3.6 Changing vCPUs and Memory of an Instance.....	41
3.3.7 Adding Nodes.....	43
3.3.8 Deleting Nodes.....	45
3.3.9 Managing Tags.....	46
3.3.10 Updating the OS of an Instance.....	48
3.4 Connection Management.....	49
3.4.1 Configuring Security Group Rules.....	49
3.4.2 Binding and Unbinding an EIP.....	51
3.4.3 Viewing the IP Address and Port Number.....	52
3.5 Data Management.....	53
3.5.1 Importing and Exporting Data by Running COPY.....	53
3.6 Intra-region DR.....	72
3.6.1 Creating a DR Instance.....	73
3.6.2 Deleting the DR Relationship.....	77
3.7 Cross-region Dual-active DR.....	78
3.7.1 Overview.....	78
3.7.2 Creating a Dual-Active Relationship.....	78
3.7.3 Deleting a Dual-active Relationship.....	79
3.8 Data Backup.....	80
3.8.1 Overview.....	80
3.8.2 Managing Automated Backups.....	81
3.8.3 Setting a Cross-Region Backup Policy.....	86
3.8.4 Managing Manual Backups.....	89
3.9 Data Restoration.....	92
3.9.1 Restoration Methods.....	92
3.9.2 Restoring Data to a New Instance.....	92
3.9.3 Restoring a Backup to a Specific Point in Time.....	94
3.10 Parameter Template Management.....	95
3.10.1 Creating a Parameter Template.....	95
3.10.2 Modifying a Parameter Template.....	96
3.10.3 Viewing Parameter Change History.....	99
3.10.4 Exporting a Parameter Template.....	100
3.10.5 Comparing Parameter Templates.....	101
3.10.6 Replicating a Parameter Template.....	102
3.10.7 Resetting a Parameter Template.....	104
3.10.8 Applying a Parameter Template.....	104
3.10.9 Viewing Application Records of a Parameter Template.....	104
3.10.10 Modifying a Parameter Template Description.....	105
3.10.11 Deleting a Parameter Template.....	105

3.11 Audit on Instance Operations.....	106
3.11.1 Key Operations Supported by CTS.....	106
3.11.2 Querying Traces.....	107
3.12 Monitoring and Alarm Configuration.....	108
3.12.1 GeminiDB Cassandra Metrics.....	108
3.12.2 Configuring Alarm Rules.....	115
3.12.3 Viewing Metrics.....	120
3.13 Enterprise Project.....	121
3.13.1 Overview.....	121
3.13.2 Managing Quotas.....	121
3.14 Billing Management.....	123
3.14.1 Renewing Instances.....	123
3.14.2 Changing the Billing Mode from Pay-per-Use to Yearly/Monthly.....	124
3.14.3 Changing the Billing Mode from Yearly/Monthly to Pay-per-Use.....	126
3.14.4 Unsubscribing from a Yearly/Monthly Instance.....	127
4 FAQs.....	130
4.1 Product Consulting.....	130
4.1.1 What Should I Pay Attention to When Using GeminiDB Cassandra?.....	130
4.1.2 What Is GeminiDB Cassandra Instance Availability?.....	130
4.2 Billing.....	130
4.2.1 What Are the Differences Between Yearly/Monthly and Pay-per-use Billing Mode?.....	131
4.2.2 Can I Switch Between Yearly/Monthly and Pay-per-Use Payments?.....	131
4.3 Database Usage.....	131
4.3.1 Why Does the Overall Instance Performance Deteriorate When QPS Increases After the Batch Size Is Decreased?.....	131
4.3.2 What Can I Do if Error "field larger than field limit (131072)" Is Reported During Data Import?....	132
4.3.3 What Should I Pay Attention to When Creating a GeminiDB Cassandra Table?.....	133
4.3.4 How Do I Detect and Resolve BigKey and HotKey Issues?.....	137
4.3.5 How Do I Set Up a Materialized View?.....	141
4.3.6 How Do I Use a Secondary Index?.....	144
4.3.7 How Do I Set Paging Query?.....	145
4.4 Database Connection.....	146
4.4.1 What Can I Do If Spark Failed to Connect to Cassandra?.....	146
4.4.2 What Can I Do If an Error Occurs When I Use Java Driver and a Mapped IP Address to Connect to a Database?.....	147
4.4.3 How Can I Create and Connect to an ECS?.....	148
4.4.4 Can I Change the VPC of a GeminiDB Cassandra Instance?.....	148
4.5 Backup and Restoration.....	149
4.5.1 How Long Does GeminiDB Cassandra Store Backup Data?.....	149
4.6 Instance Freezing, Release, Deletion, and Unsubscription.....	149
A Change History.....	151

1 Service Overview

1.1 What Is GeminiDB Cassandra API?

GeminiDB Cassandra API is a cloud-native NoSQL database compatible with Cassandra. It supports Cassandra Query Language (CQL), which gives you SQL-like syntax. GeminiDB Cassandra API is secure, reliable, scalable, and easy to manage and provides outstanding read/write performance.

- High security and reliability
 - A multi-layer security system, including VPC, subnet, security group, SSL, and fine-grained permission control, ensures database security and user privacy.
 - You can deploy an instance across three AZs and quickly back up or restore data to improve data reliability.
 - The distributed architecture provides superlative fault tolerance (*N-1* reliability).
- Outstanding read/write performance

GeminiDB Cassandra API gives you 3 times the performance of the open source version. Data can be written to this high availability database 24/7, and with automated load balancing and elastic scaling, you always have all the performance you need.
- Flexible scaling

Decoupled compute and storage allows you to add compute nodes in minutes and scale up storage capacity in seconds without service interruptions.
- Friendly UI

On the instance management console, you can create or delete instances in a visual way. Backup and restoration, configuring alarms or scaling out or in compute nodes is just as easy.

Typical Application Scenarios

- Internet

GeminiDB Cassandra provides excellent read and write performance, flexibility, and fault tolerance, making it easy for those websites that provide

product catalogs, recommendations, personalization engines, and transaction records to handle high concurrency and ensure low latency.

Advantages

Large-scale clusters

Each cluster can include up to 100 nodes, helping write-intensive Internet applications process massive volumes of data.

High availability and scalability

The failure of one node does not affect the availability of the entire cluster. Compute resources and storage space can be quickly scaled out or up, with minimal service interruptions.

High-concurrency writes

Powerful write performance helps you handle a huge number of concurrent e-commerce transactions.

- Industrial data collection

GeminiDB Cassandra API is fully compatible with Cassandra, so it can help you collect, organize, and store data from different types of terminals, and aggregate and analyze the data in real-time.

Advantages

Large-scale clusters

The large-scale clusters are well suited to collect and store massive numbers of manufacturing metrics.

High availability and performance

Data can be written to this the database 24/7.

Fast backup and restoration

Snapshots allow for fast backup and recovery.

Scaling in minutes

Scaling operations complete in minutes, helping handle service or project peaks.

1.2 Compatible APIs and Versions

This section describes the compatible APIs and versions supported by GeminiDB Cassandra.

Table 1-1 Compatible APIs and versions

Compatible API	Instance Type	Version
Cassandra	Cluster	3.11

1.3 Instance Specifications

Instances of the same type can have different memory specifications. You can select instances of different specifications based on application scenarios.

This section describes the instance specifications supported by GeminiDB Cassandra. The instance specifications depend on the selected flavor.

Table 1-2 GeminiDB Cassandra cluster instance specifications

Flavor	vCPUs	Memory (GB)	Min. Storage Space (GB)	Max. Storage Space (GB)
geminidb.cassandra.large.4	2	8	100	12,000
geminidb.cassandra.xlarge.4	4	16	100	24,000
geminidb.cassandra.2xlarge.4	8	32	100	48,000
geminidb.cassandra.4xlarge.4	16	64	100	96,000
geminidb.cassandra.8xlarge.4	32	128	100	192,000
geminidb.cassandra.large.8	2	16	100	12,000
geminidb.cassandra.xlarge.8	4	32	100	24,000
geminidb.cassandra.2xlarge.8	8	64	100	48,000
geminidb.cassandra.4xlarge.8	16	128	100	96,000
geminidb.cassandra.6xlarge.8	24	192	100	144,000
geminidb.cassandra.8xlarge.8	32	256	100	192,000

1.4 Instance Statuses

The status of an instance indicates the health of the instance. You can view the status of an instance on the console.

Table 1-3 Instance statuses

Status	Description
Available	The DB instance is available.
Abnormal	The instance is abnormal.

Status	Description
Creating	The instance is being created.
Creation failed	DB instance creation fails.
Restarting	The instance is being restarted.
Resetting password	The administrator password is being reset.
Adding node	Nodes are being added to an instance.
Deleting node	Nodes are being deleted from an instance.
Scaling storage space	The storage space of an instance is being scaled up.
Changing specifications	The vCPUs and memory of an instance are being changed.
Uploading backup	The backup file is being uploaded.
Backing up	A database backup is being created.
Checking restoration	The backup of the instance is being restored to a new instance.
Changing to yearly/monthly	The billing mode is being changed from pay-per-use to yearly/monthly.
Changing to pay-per-use	The billing mode is being changed from yearly/monthly to pay-per-use.
Creating a DR cluster	A DR instance is being created.
Canceling DR relationship	A DR instance is being deleted.
Frozen	The instance is frozen because your balance drops to or below zero.
Unfreezing	Overdue payments are cleared, and the DB instance is being unfrozen.
Checking changes	The yearly/monthly instance is pending check when its billing mode is changed.

1.5 Database Constraints

1.5.1 Basic Design

Design Rules

Rule 1: Do not store big data such as images and files in databases.

Rule 2: The maximum size of the key and value in a single row cannot exceed 64 KB, and the average size of rows cannot exceed 10 KB.

Rule 3: The data deletion policy must be considered in the design of a table. Data in a table cannot increase infinitely without being deleted.

Rule 4: Partition keys can evenly distribute workloads to avoid data skew.

A partition key of a primary key determines a logical partition for storing table data. If partition keys are not evenly distributed, data and load between nodes are unbalanced, resulting in a data skew problem.

Rule 5: The design of partition keys can evenly distribute data access requests to avoid BigKey or HotKey issues.

- **BigKey issue:** The main cause of BigKey is that the primary key is improperly design. As a result, there are too many records or too much data in a single partition. Once a partition becomes extremely large, access to the partition increases load of a server where the partition is located, and even causes the Out of Memory (OOM) error.
- **HotKey issue:** This issue occurs when a key is frequently operated in a short period of time. For example, breaking news can cause a spike in traffic and large number of requests. As a result, the CPU usage and the load on the node on which the key is located increase, affecting other requests to the node and reducing the success rate of services. HotKey issues will also occur during promotion of popular products and Internet celebrity live streaming.

For details about how to handle BigKey and HotKey issues, see [How Do I Detect and Resolve BigKey and HotKey Issues?](#)

Rule 6: The number of rows of a single partition key cannot exceed 100,000, and the disk space of a single partition cannot exceed 100 MB.

- The number of rows of a single partition key cannot exceed 100,000.
- The size of records under a single partition key cannot exceed 100 MB.

Rule 7: Ensure strong consistency between data copies written to GeminiDB Cassandra, but do not support transactions.

Table 1-4 GeminiDB Cassandra consistency description

Consistency Model	Consistency Supported	Description
Concurrent write consistency	Yes	GeminiDB Cassandra does not support transactions, and data writing is strongly consistent.

Consistency Model	Consistency Supported	Description
Consistency between tables	Yes	GeminiDB Cassandra does not support transactions, and data writing is strongly consistent.
Data migration consistency	Eventual consistency	DRS migration provides the data sampling, comparison, and verification capabilities. After services are migrated, data verification is automatically performed.

Rule 8: For large-scale storage, database splitting must be considered.

Ensure that the number of nodes in the GeminiDB Cassandra cluster is less than 100. If the number of nodes exceeds 100, split the cluster vertically or horizontally.

- Vertical splitting: Data is split by functional module, for example, the order database, product database, and user database. In this mode, the table structures of multiple databases are different.
- Horizontal sharding: Data in the same table is divided into blocks and stored in different databases. The table structures in these databases are the same.

Rule 9: Avoid tombstones caused by large-scale deletion.

- Use TTL instead of Delete if possible.
- Do not delete a large amount of data. Delete data by primary key prefix.
- A maximum of 1,000 rows can be deleted at a time within a partition key.
- Avoid querying deleted data during range query.
- Do not frequently delete data of a large range in one partition.

Design Suggestion

Suggestion 1: Properly control the database scale and quantity.

- It is recommended that the number of data records in a single table be less than or equal to 100 billion.
- It is recommended that a single database contain no more than 100 tables.
- It is recommended that the maximum number of fields in a single table be 20 to 50.

Suggestion 2: Estimate how many resources that GeminiDB Cassandra servers can process.

- If it is estimated that N nodes need to be used, adding additional N/2 nodes is recommended for fault tolerance and performance consistency.
- In normal scenarios, the CPU usage of each node is limited to 50% to avoid fluctuation during peak hours.

Suggestion 3: To store large volumes of data, perform a test run based on service scenarios.

In service scenarios with a large number of requests and data volume, you need to test the performance in advance because the service read/write ratio, random access mode, and instance specifications vary greatly.

Suggestion 4: Split database cluster granularity properly.

- In distributed scenarios, microservices of a service can share a GeminiDB Cassandra cluster to reduce resource and maintenance costs.
- The service can be divided into different clusters based on the data importance, number of tables, and number of records in a single table.

Suggestion 5: Do not frequently update some fields in a single data record.

Suggestion 6: If there are too many nested elements such as List, Map, or Set, read and write performance will be affected. In this case, convert such elements into JSON data for storage.

1.5.2 DB Instance Specifications and Performance

Specifications and Performance

For details about GeminiDB Cassandra performance and specifications, see [Table 1-5](#).

Table 1-5 GeminiDB Cassandra performance and specifications

vCPUs	Max. Data Volume per Node (GB)	Max. Transactions per Second (TPS) on a Single Node
2	200	250
4	250	1,000
8	250	2,500
16	500	5,000
24	500	7,000
32	500	10,000

NOTE

- The TPS depends on the number of vCPUs of an instance.
- When the data volume or TPS of a single node goes beyond the above limits, problems like high latency, request failure, or OMM will occur. Scaling storage space is recommended. All issues caused by untimely scaling are not within the SLA commitment scope.

1.5.3 Database Objects

Naming Rules

Rule 1: The object name cannot be duplicated with any keyword of the database.

Rule 2: Object names (including database names, table names, field names, and index names) must be in lowercase and separated by underscores (_).

Rule 3: The length of an object name (including the database name, table name, field name, and index name) cannot exceed 30 characters.

Rule 4: The table alias must be short. Generally, aliases are in lowercase letters.

Table Design Rules

Rule 1: Compatibility must be considered during table design.

Columns can be added but cannot be deleted.

Rule 2: The table name and database name cannot exceed 48 bytes.

Rule 3: By default, tables are created based on the optimal performance specifications. If the high-performance table is not required, you can set performance parameter **ZOO_THROUGHPUT** to **big**, **medium**, or **small** when creating a table. By default, this parameter is not set to **big**. If you use RocksDB as the storage engine, memory needs to be allocated in advance and the number of tables created in an instance is limited. For details, see [What Should I Pay Attention to When Creating a GeminiDB Cassandra Table?](#).

If necessary, use denormalization and redundancy to improve the read performance.

Indexing Rules

Rule 1: Design all queries as primary-key based queries and do not rely too much on secondary indexes.

Rule 2: An index can be used for query only after it is configured.

Rule 3: Do not frequently update indexes.

Rule 4: Do not create an index column for a table that contains too many duplicate values. For example, if one table stores 100 million data records and one of its columns contains the same data or a few types data, creating an index column for this table is not recommended.

Rule 5: The **counter** column cannot be indexed.

Rule 6: Do not create an index for any column that is frequently updated or deleted.

Rule 7: Use indexes together with partition keys to minimize message forwarding between nodes and resource consumption and prevent out-of-memory or high CPU usage.

View Rules

- If a materialized view is used, ensure that the original table corresponds to no more than three views. The more views the original table corresponds to, the greater impacts on the synchronization of views.
- Do not use any frequently-updated field in the original table as the primary key of a view.

Flow Table Rules

One flow table stores 24 hours of data by default. If there is a large amount of data to be queried, return results on multiple pages. No more than 100 query results are returned each time and a retry is allowed if a query request times out.

1.5.4 Access and Connection Pools

Rule 1: A connection pool must be used to access the database to improve reliability.

Rule 2: GeminiDB Cassandra clusters use RoundRobinPolicy for load balancing.

1.5.5 Batches

Rule 1: Logged batches are not supported. Only unlogged batches are supported.

Rule 2: A maximum of 25 rows of data can be operated in a batch.

Rule 3: In a batch, a request size cannot exceed 5 KB.

Rule 4: In a batch, no more than 10 partitions are involved, and only one table is operated.

1.5.6 Queries

Using a Sort Key for Range Query

It is recommended that the sequence of the sort keys for range query be the same as that used during table creation. Otherwise, the performance deteriorates.

NOTE

If no sort key sequence is specified, the default sort key sequence is ASC during query and table creation.

COUNT Query

If a database contains a very large amount of data, do not run the following statement to query the database. Otherwise, the query may fail.

```
select count(*) from "test" where sds_uid='100000000000000006250004';
```

The following statement is recommended:

```
select sum(row_count) From system_distributed.size_estimates WHERE keyspace_name="" and table_name="";
```

NOTE

This query is an asynchronous task in the background, so the results are not accurate and for reference only.

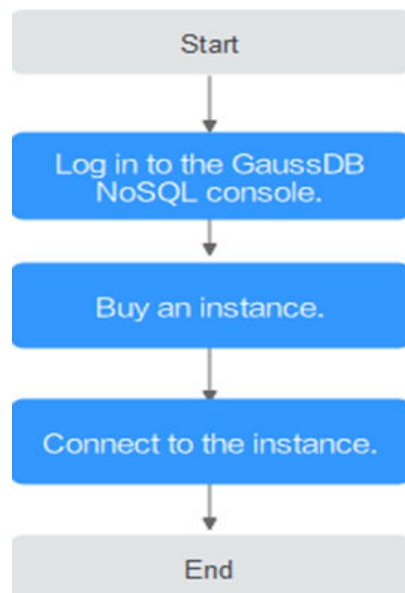
2 Getting Started with GeminiDB Cassandra API

2.1 Service Overview

This section describes how to buy an instance and then connect to and manage it.

Process

Figure 2-1 Flowchart



Operation Guide

The process of buying and using an instance involves the following steps:

Step 1: Log in to the GeminiDB Cassandra API console.

Step 2: **Buy an instance.**

Step 3: [Connect to the instance](#).

2.2 Buying an Instance

This section describes how to buy a GeminiDB Cassandra instance that is compatible with Cassandra APIs.

Each tenant can have up to 50 GeminiDB Cassandra instances by default. To request a higher quota, contact customer service.

Prerequisites

- You have registered a Huawei Cloud account.

Procedure

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases** > **GeminiDB**.

Step 3 On the **Instances** page, click **Buy DB Instance**.

Step 4 On the displayed page, specify instance specifications and click **Next**.

Figure 2-2 Billing mode and basic information

The screenshot displays the configuration interface for a GeminiDB instance. It is divided into two main sections by a horizontal line. The top section contains configuration options for Billing Mode, Region, and Project. The bottom section contains configuration options for DB Instance Name, Compatible API, DB Instance Type, DB Engine Version, and AZ.

Billing Mode	<input checked="" type="radio"/> Yearly/Monthly <input type="radio"/> Pay-per-use
Region	<input type="text" value="EU-Dublin"/> ▼
<small>Regions are geographic areas isolated from each other. Resources are region-specific and cannot t</small>	
Project	<input type="text" value="EU-Dublin"/> ▼
<hr/>	
DB Instance Name	<input type="text" value="nosql-dc71"/> ?
Compatible API	<input checked="" type="radio"/> Cassandra <input type="radio"/> InfluxDB <input type="radio"/> Redis
DB Instance Type	<input type="text" value="Cluster"/>
DB Engine Version	<input type="text" value="3.11"/>
AZ	<input type="text" value="eu-west-101a"/>

Table 2-1 Billing parameters

Parameter	Description
Billing Mode	<p>Method that the instance is billed in. The value can be Yearly/Monthly or Pay-per-use.</p> <ul style="list-style-type: none"> Yearly/Monthly <ul style="list-style-type: none"> In this mode, specify Required Duration at the bottom of the page. The system deducts the fees incurred from your account based on the service price. If you do not need such an instance any longer after it expires, change the billing mode to pay-per-use. For details, see Changing the Billing Mode from Yearly/Monthly to Pay-per-Use. <p>NOTE Yearly/Monthly instances cannot be deleted directly. If such an instance is no longer required, unsubscribe from it. For details, see Unsubscribing from a Yearly/Monthly Instance.</p> <ul style="list-style-type: none"> Pay-per-use <ul style="list-style-type: none"> If you select this billing mode, you are billed based on how much time the instance is in use. If you expect to use an instance for a long period of time, change its billing mode to yearly/monthly to optimize costs. For details, see Changing the Billing Mode from Pay-per-Use to Yearly/Monthly.

Table 2-2 Basic information

Parameter	Description
Region	<p>The region where the instance is deployed.</p> <p>NOTICE Select the region nearest where you will access the instance from so that latency can be kept to a minimum and response is faster. Instances deployed in different regions cannot communicate with each other over a private network. After you buy an instance, you cannot change its region.</p>
DB Instance Name	<p>The instance name:</p> <ul style="list-style-type: none"> Can be the same as an existing instance name. Can include 4 to 64 bytes and must start with a letter. It is case-sensitive and allows only letters, digits, hyphens (-), and underscores (_). <p>After an instance is created, you can change its name. For details, see Changing an Instance Name.</p>
Compatible API	Cassandra
DB Instance Type	Cluster

Parameter	Description
DB Engine Version	3.11
AZ	Availability zone where the instance is created. An AZ is a part of a region with its own independent power supplies and networks. AZs are physically isolated but can communicate through an internal network.

Figure 2-3 Specifications and storage



Table 2-3 Specifications and storage

Parameter	Description
Instance Specifications	Decoupled storage and compute and software-hardware synergy deliver twice or more the performance of an on-premises database with the same specifications. When you create an instance, select higher specification and specify as few nodes as possible. For example, if you need 8 vCPUs, 32 GB, and 6 nodes for an on-premises deployment, then for a GeminiDB Cassandra instance with 8 vCPUs and 32 GB of memory, you only need 3 nodes. Select appropriate specifications based on the CPU-memory ratio. After an instance is created, you can change its specifications. For details, see Changing vCPUs and Memory of an Instance .
Nodes	Number of nodes that the instance is deployed on. After an instance is created, you can add nodes. For details, see Adding Nodes .

Parameter	Description
Storage Space	<p>Instance storage space. The range depends on the instance specifications. For details, see Instance Specifications.</p> <p>Select at least 1 GB each time you scale up the storage, and ensure that the increment is an integer.</p> <p>Enable autoscaling to ensure that the instance has sufficient storage and keeps available. To enable this function, just switch on button Configure Autoscaling and set the following parameters:</p> <ul style="list-style-type: none"> • If available storage drops to or below: The storage threshold for triggering autoscaling. When the percentage of available storage drops to or below the threshold you set or 10 GB, the system automatically scales up your instance storage. • Increase by: The percentage that your instance storage will be scaled up at. If the increased storage is not a multiple of 10 GB, the system will round it up to the nearest multiple of 10 GB. At least 100 GB is added each time. • Autoscaling Limit: Maximum amount that the system can automatically scale up an instance's storage space to. The value must be no less than the total storage of the instance and cannot exceed its maximum storage. <p>After an instance is created, you can scale up its storage if necessary. For details, see Scaling Up Storage Space.</p> <p>NOTE</p> <ul style="list-style-type: none"> • Once autoscaling is enabled, an agency will be created and fees will be automatically deducted from your account. • Autoscaling is available only when you have the required permission. To enable this function, contact customer service. • You can enable autoscaling after an instance is created. For details, see Configuring Autoscaling.

Figure 2-4 Network configuration

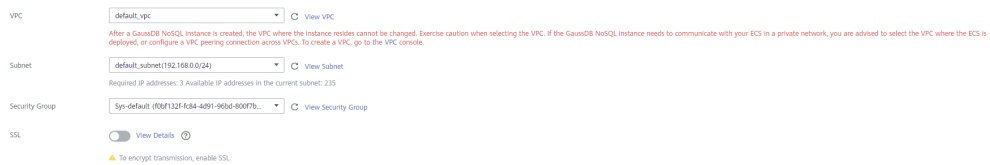


Table 2-4 Network configuration

Parameter	Description
VPC	<p>The virtual network where the instance is created. A VPC isolates networks for different services. You can select an existing VPC or create one.</p> <p>If there are no VPCs available, the system automatically allocates a VPC to you.</p> <p>For details, see "Creating a VPC" in the <i>Virtual Private Cloud User Guide</i>.</p> <p>NOTE</p> <ul style="list-style-type: none"> After a GeminiDB Cassandra instance is created, its VPC cannot be changed. If you want to connect to a GeminiDB Cassandra instance through an ECS over a private network, the instance and the ECS must be in the same VPC. If they are not in the same VPC, you can create a VPC peering connection to enable access.
Subnet	<p>A subnet where your instance is created. The subnet provides dedicated and isolated networks, improving network security.</p> <p>NOTE</p> <p>Create an IPv4 subnet or select an existing one. IPv6 subnets are not supported.</p>
Security Group	<p>A security group controls access between your instance and other services. Ensure that the security group you select allows the client to access the instance.</p> <p>If there are no security groups available, the system allocates one to you by default.</p>

Figure 2-5 Database configuration

Administrator rwuser

Administrator Password Keep your password secure. The system cannot retrieve your password.

Confirm Password

Parameter Template Default-Cassandra-3.11 [View Parameter Template](#)

Tags It is recommended that you use TMS's predefined tag function to add the same tags to different cloud resources. [View predefined tags](#)


You can add more tags.

Table 2-5 Database configuration

Parameter	Description
Administrator	Username of the administrator account. The default value is rwuser .
Administrator Password	<p>Password of the administrator account. The password:</p> <ul style="list-style-type: none"> • Can include 8 to 32 characters. • Can include uppercase letters, lowercase letters, digits, and any of the following special characters: ~!@#%^*-_ =+? • For security reasons, set a strong password. The system will verify the password strength. <p>Keep your password secure. The system cannot retrieve it if it is lost.</p>
Confirm Password	This password must be consistent with administrator password.
Parameter Template	<p>A template of parameters for creating an instance. The template contains API configuration values that are applied to one or more instances.</p> <p>After an instance is created, you can modify its parameters for optimal performance. For details, see Modifying a Parameter Template.</p>
Enterprise Project	<p>This parameter is provided for enterprise users.</p> <p>An enterprise project groups cloud resources, so you can manage resources and members by project. The default project is default.</p> <p>Select an enterprise project from the drop-down list. For more information about enterprise projects, see Enterprise Management User Guide.</p>

Figure 2-6 Tag configuration

Tags

It is recommended that you use TMS's predefined tag function to add the same tags to different cloud resources. [View predefined tags](#) 

Tag key Tag value

You can add 20 more tags.

Table 2-6 Tags

Parameter	Description
Tags	<p>This setting is optional. Adding tags helps you better identify and manage your GeminiDB Cassandra instances.</p> <p>Each instance can have up to 20 tags by default.</p> <p>A tag consists of a tag key and a tag value.</p> <ul style="list-style-type: none"> • Tag key: Mandatory if the instance is going to be tagged. Each tag key is unique for each instance. It can include up to 36 characters, including digits, letters, underscores (_), and hyphens (-). • Tag value: Optional if the instance is going to be tagged. The value can contain up to 43 characters, including digits, letters, underscores (_), periods (.), and hyphens (-). <p>After an instance is created, you can view its tags on the Tags page and can also add tags to, modify, and delete tags of your instance. For details, see Managing Tags.</p>

Figure 2-7 Required duration configuration




Table 2-7 Required duration


Parameter	Description
Required Duration	The length of your subscription if you select Yearly/Monthly billing. Subscription lengths range from one month to three years.
Auto-renew	<ul style="list-style-type: none"> • This option is not selected by default. • If you select this option, the auto-renew cycle is determined by the selected required duration.

Step 5 On the displayed page, confirm the instance details.

- For yearly/monthly instances
 - If you need to modify the settings, click **Previous** to modify parameters.
 - If no modification is required, read and agree to the service agreement and click **Pay Now**.
- For pay-per-use instances
 - If you need to modify the settings, click **Previous** to modify parameters.
 - If no modification is required, read and agree to the service agreement and click **Submit**.

Step 6 To view and manage your instance, go to the **Instances** page.

- It takes about 5 to 9 minutes to create an instance. During the process, the instance status is **Creating**.
- After the creation is complete, the status changes to **Available**.

You can click  in the upper right corner to refresh the instance status.

- Automated backup is enabled by default during instance creation. After the instance is created, a full backup is created.

----End

2.3 Instance Connections

2.3.1 Connection Methods

Table 2-8 Connection methods

Method	Scenario	Description
Private network	Private IP addresses are provided by default. Your applications are deployed on an ECS that is in the same region and VPC as your instances.	High security and performance
Public network	If you cannot access a DB instance through a private IP address, bind an EIP to the DB instance first and connect the ECS to the DB instance through the EIP.	<ul style="list-style-type: none"> • Low security • For faster transmission and improved security, migrate your applications to an ECS that is in the same subnet as your instance and use a private IP address to access the instance.
Java	An example of connecting to a GeminiDB Cassandra instance using Java is provided.	-
Go	An example of connecting to a GeminiDB Cassandra instance using Go is provided.	-

2.3.2 Connecting to an Instance over a Private Network

You can install the Cassandra client on the ECS and access the instance through a private IP address.

Precautions

- The DB instances must be in the same VPC and subnet as the ECS.
- The ECS must be in a security group that has access to the instances. For details, see [Configuring Security Group Rules](#).

Prerequisites

1. A GeminiDB Cassandra instance has been created and is running properly.
2. An ECS has been created. The following uses a Linux ECS as an example. For details, see [Purchasing an ECS](#) in *Getting Started with Elastic Cloud Server*.
3. Download and install the Cassandra client that matches the CPU type of the ECS.
 - If the CPU type is x86, download the [Cassandra client installation package](#).

Procedure

Step 1 Log in to ECS.

For details, see [Logging In to an ECS](#) in *Getting Started with Elastic Cloud Server*.

Step 2 Upload the Cassandra client installation package to the ECS.

Step 3 Run the following command to decompress the client installation package. The x86 client is used as an example.

```
unzip Cassandra_cqlsh_x86_64.zip
```

Step 4 Run the following command to grant the execute permission on all files:

```
chmod +x *
```

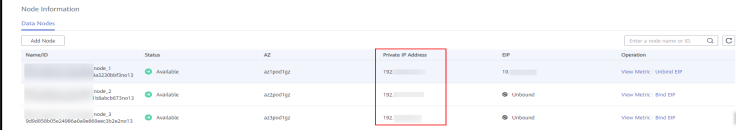

Step 5 Connect to the DB instance in the directory where the cqlsh tool is located.

```
./cqlsh <DB_HOST> <DB_PORT> -u <DB_USER>
```

Example:

```
./cqlsh 192.xx.xx.xx 8635 -u rwuser
```

Table 2-9 Parameter description

Parameter	Description
<DB_HOST>	<p>The private IP address of the instance to be connected.</p> <p>To obtain this IP address, go to the Instance Management page and click the target instance name. The IP address can be found in the Private IP Address field under Node Information on the Basic Information page.</p> <p>If the instance you purchased has multiple nodes, select the private IP address of any node.</p> <p>Figure 2-8 Viewing the private IP address</p> 
<DB_PORT>	<p>Port number of the instance to be connected. The default port number is 8635. Replace it with the actual port number.</p> <p>Click the instance name to go to the Basic Information page and obtain the port number in the Network Information area.</p> <p>Figure 2-9 Viewing the port number</p> 
<DB_USER>	<p>Database account. The default value is rwuser.</p>

Step 6 Check the results. If the following information is displayed, the connection is successful.

```
rwuser@cqlsh>
```

----End

2.3.3 Connecting to an Instance over a Public Network

You can use an ECS or local device to connect to a GeminiDB Cassandra instance over a public network.

This section describes how to use a Linux ECS to connect to a GeminiDB Cassandra instance over a public network.

Prerequisites

1. Bind an EIP to the GeminiDB Cassandra instance node and set security group rules. For details, see [Binding and Unbinding an EIP](#) and [Configuring Security Group Rules](#).
2. An ECS has been created. The following uses a Linux ECS as an example. For details, see [Purchasing an ECS](#) in *Getting Started with Elastic Cloud Server*.
3. Download and install the Cassandra client that matches the CPU type of the ECS.
 - If the CPU type is x86, download the [Cassandra client installation package](#).

Procedure

Step 1 Log in to the ECS. For details, see [Logging In to an ECS](#) in *Getting Started with Elastic Cloud Server*.

Step 2 Upload the Cassandra client installation package to the ECS.

Step 3 Run the following command to decompress the client installation package. The x86 client is used as an example.

```
unzip Cassandra_cqlsh_x86_64.zip
```

Step 4 Run the following command to grant the execute permission on all files:

```
chmod +x *
```

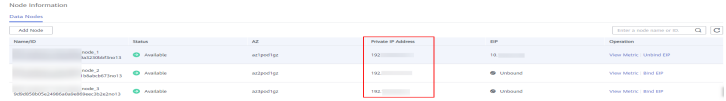

Step 5 Connect to the DB instance in the directory where the cqlsh tool is located.

```
./cqlsh <DB_HOST> <DB_PORT> -u <DB_USER>
```

Example:

```
./cqlsh 192.xx.xx.xx 8635 -u rwuser
```

Table 2-10 Parameter description

Parameter	Description
<DB_HOST>	<p>Specifies the EIP bound to the instance to be connected.</p> <p>To obtain the EIP, go to the Instance Management page and click the target instance name. The EIP can be found in the EIP column in the Node Information area on the Basic Information page.</p> <p>If the instance you has multiple nodes, you can bind the EIP to any node to connect to the GeminiDB Cassandra instance.</p> <p>Figure 2-10 Viewing EIPs</p>  <p>If there are no EIPs bound to the instance, bind an EIP to the instance by referring to Binding and Unbinding an EIP and then connect to the instance.</p>
<DB_PORT>	<p>Port number of the instance to be connected. The default port number is 8635. Replace it with the actual port number.</p> <p>Click the instance name to go to the Basic Information page and obtain the port number in the Network Information area.</p> <p>Figure 2-11 Viewing the port number</p> 
<DB_USER>	<p>Database account. The default value is rwuser.</p>

Step 6 Check the results. If the following information is displayed, the connection is successful.

```
rwuser@cqlsh>
```

----End

2.3.4 Connecting to an Instance Using Java

This section describes how to use the Java to connect to a GeminiDB Cassandra instance.

Prerequisites

- A GeminiDB Cassandra instance has been created and is running properly. For details about how to create a GeminiDB Cassandra instance, see [Buying an Instance](#).
- For details about how to create an ECS, see Getting Started > [Creating an ECS](#) in the *Elastic Cloud Server User Guide*.
- JDK has been installed on the ECS.

Procedure

Step 1 Obtain the private IP address and port number of the GeminiDB Cassandra instance.

For details about how to obtain the private IP address and port number, see [Viewing the IP Address and Port Number](#).

Step 2 Log in to the ECS. For details, see [Logging In to an ECS](#) in *Getting Started with Elastic Cloud Server*.

Step 3 Edit the code for connecting to the GeminiDB Cassandra instance.

```
import com.datastax.driver.core.*;

Cluster cluster = null;
try {
    cluster = Cluster.builder()
        .addContactPoint("127.0.0.1")//Private IP address of the GeminiDB Cassandra instance
        obtained in step 1
        .withPort(8635) //Port number of the GeminiDB Cassandra instance obtained in
        step 1
        .build();
    Session session = cluster.connect();

    ResultSet rs = session.execute("select release_version from system.local");
    Row row = rs.one();
    System.out.println(row.getString("release_version"));
} finally {
    if (cluster != null) cluster.close();
}
```

Step 4 Run the sample code to check whether the result is normal.

----End

2.3.5 Connecting to an Instance Using Go

This section describes how to connect to a GeminiDB Cassandra instance using Go.

Prerequisites

- A GeminiDB Cassandra instance has been created and is running normally. For details about how to create a GeminiDB Cassandra instance, see [Buying an Instance](#).
- For details about how to create an ECS, see Getting Started > [Creating an ECS](#) in the *Elastic Cloud Server User Guide*.
- You have installed the Go environment on the ECS. If not, download the [Go installation package](#).

Procedure

Step 1 Obtain the private IP address and port number of the GeminiDB Cassandra instance.

For details about how to obtain the private IP address and port number, see [Viewing the IP Address and Port Number](#).

Step 2 Log in to the ECS. For details, see [Logging In to an ECS](#) in *Getting Started with Elastic Cloud Server*.

Step 3 Edit the code for connecting to the GeminiDB Cassandra instance.

```
// Default LoadBalancingPolicy RoundRobinHostPolicy
cluster := gocql.NewCluster("127.0.0.1,127.0.0.2,127.0.0.3")
cluster.Authenticator = gocql.PasswordAuthenticator{
    Username: "user",
    Password: "password"
}
cluster.Keyspace = "ks1"
// connect to the cluster
session, err := cluster.CreateSession()
if err != nil {
    log.Fatal(err)
}
defer session.Close()
```

Step 4 Run sample code to check whether the result is normal.

----End

Executing Write and Read Operations

Create a session query. Query parameters cannot be used in other statements and cannot be modified after the query starts.

Use `Query.Exec` if you need to read the query results after a query is executed:

```
err := session.Query(`INSERT INTO tweet (timeline, id, text) VALUES (?, ?, ?)`,
    "me", gocql.TimeUUID(), "hello world").WithContext(ctx).Exec()
```

Use `Query.Scan` if you want to read one line of data:

```
err := session.Query(`SELECT id, text FROM tweet WHERE timeline = ? LIMIT 1`,
    "me").WithContext(ctx).Consistency(gocql.One).Scan(&id, &text)
```

Use `Iter.Scanner` if you want to read multiple lines of data:

```
scanner := session.Query(`SELECT id, text FROM tweet WHERE timeline = ?`,
    "me").WithContext(ctx).Iter().Scanner()
for scanner.Next() {
    var (
        id gocql.UUID
        text string
    )
    err = scanner.Scan(&id, &text)
    if err != nil {
        log.Fatal(err)
    }
    fmt.Println("Tweet:", id, text)
}
// scanner.Err() closes the iterator, so scanner nor iter should be used afterwards.
if err := scanner.Err(); err != nil {
    log.Fatal(err)
}
```

Executing Multiple Queries Concurrently

It is safe to share a session in multiple goroutines. If necessary, you execute multiple queries using multiple goroutines.

```
results := make(chan error, 2)
go func() {
    results <- session.Query(`INSERT INTO tweet (timeline, id, text) VALUES (?, ?, ?)`,
        "me", gocql.TimeUUID(), "hello world 1").Exec()
}()
go func() {
    results <- session.Query(`INSERT INTO tweet (timeline, id, text) VALUES (?, ?, ?)`,
        "me", gocql.TimeUUID(), "hello world 2").Exec()
}()
```

3 Working with GeminiDB Cassandra API

3.1 Permissions Management

3.1.1 Creating a User and Assigning Permissions

This section describes how to use [IAM](#) to control fine-grained permissions for your GeminiDB resources. With IAM, you can:

- Create IAM users for employees based on your enterprise's organizational structure. Each IAM user will have their own security credentials for accessing GeminiDB resources.
- Grant only the permissions required for users to perform a specific task.
- Entrust a Huawei Cloud account or cloud service to perform efficient O&M on your GeminiDB resources.

If your Huawei Cloud account does not require individual IAM users, skip this section.

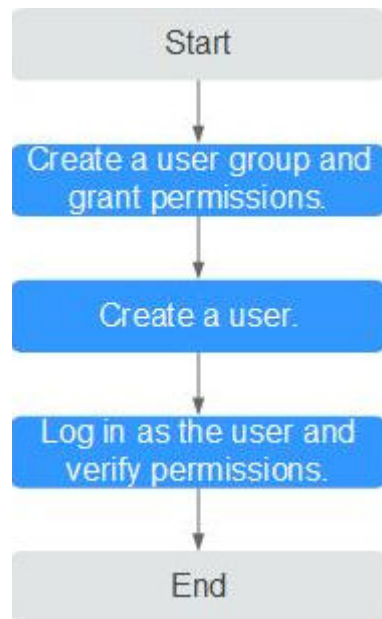
The following describes the procedure for granting permissions (see [Figure 3-1](#)).

Prerequisites

Learn about the permissions (see) supported by GeminiDB and choose policies or roles according to your requirements. For system policies of other services, see [Permissions Policies](#).

Process Flow

Figure 3-1 Process of granting GeminiDB permissions



1. **Create a user group and assign permissions** to it.

Create a user group on the IAM console and attach the **GaussDB NoSQL FullAccess** policy to the group.

NOTE

To use some interconnected services, you also need to configure permissions of such services.

For example, when using DAS to connect to a DB instance, you need to configure the GaussDB NoSQL FullAccess and DAS FullAccess permissions.

2. **Create an IAM user** and add it to a user group.

Create a user on the IAM console and add the user to the group created in **1**.

3. **Log in** and verify permissions.

Log in to the management console using the created user, and verify that the user only has read permissions.

Choose **Service List > GeminiDB** and click **Buy DB Instance**. If you can buy an instance, the required permission policy has taken effect.

3.1.2 Creating a Custom Policy

Custom policies can be created to supplement the system-defined policies of GeminiDB. For the actions supported for custom policies, see .

You can create custom policies in either of the following ways:

- Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.
- JSON: Edit JSON policies from scratch or based on an existing policy.

For details, see [Creating a Custom Policy](#). The following describes examples of common GeminiDB custom policies.

Example Custom Policy

- Example 1: Allowing users to create GeminiDB instances

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "nosql:instance:create"
      ]
    }
  ]
}
```

- Example 2: Deny users the permission to delete GeminiDB instances.

A policy with only "Deny" permissions must be used in conjunction with other policies to take effect. If the policies assigned to a user contain both Allow and Deny actions, the Deny actions take precedence over the Allow actions.

The following method can be used if you need to assign permissions of the **GaussDB NoSQLFullAccess** policy to a user but you want to prevent the user from deleting GeminiDB instances. Create a custom policy for denying instance deletion, and attach both policies to the group to which the user belongs. Then, the user can perform all operations on GeminiDB instances except deleting GeminiDB instances. The following is an example of the deny policy:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Deny"
      "Action": [
        "nosql:instance:delete"
      ],
    }
  ]
}
```

- Example 3: Defining permissions for multiple services in a policy

A custom policy can contain the actions of multiple services that are of the global or project-level type. The following is an example policy containing actions of multiple services:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Action": [
        "nosql:instance:create",
        "nosql:instance:rename",
        "nosql:instance:delete",
        "vpc:publiclps:list",
        "vpc:publiclps:update"
      ],
      "Effect": "Allow"
    }
  ]
}
```

```
} ]
```

3.2 Instance Lifecycle

3.2.1 Restarting an Instance

You may need to restart an instance for routine maintenance.

Precautions

- Only instances in states **Available**, **Abnormal**, or **Checking restoration** can be restarted.
- Restarting an instance will interrupt services. Exercise caution when performing this operation. Wait until off-peak hours and ensure that your application can re-connect.
- After you restart an instance, all nodes in the instance are also restarted.
- If you enable operation protection to improve the security of your account and cloud products, two-factor authentication is required for sensitive operations. For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

Procedure

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases > GeminiDB**.

Step 3 On the **Instances** page, locate the instance you want to restart and choose **More > Restart** in the **Operation** column.

Alternatively, click the name of the instance you want to restart, and on the displayed **Basic Information** page, click **Restart** in the upper right corner of the page.

Step 4 If you have enabled operation protection, click **Start Verification** in the **Restart DB Instance** dialog box. On the displayed page, click **Send Code**, enter the verification code, and click **Verify**. The page is closed automatically.

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

----End

3.2.2 Deleting Instance

You can choose to delete a pay-per-use instance on the **Instances** page based on service requirements. To delete a yearly/monthly instance, unsubscribe from it. For details, see [Unsubscribing from a Yearly/Monthly Instance](#).

Precautions

- Instances where operations are being performed cannot be deleted. They can be deleted only after the operations are complete.

- If a pay-per-use instance is deleted, its automated backups will also be deleted and you will no longer be billed for them. Manual backups, however, will be retained and generate additional costs.
- After an instance is deleted, all its data and all automated backups are automatically deleted as well and cannot be recovered. Back up it before you delete an instance. For details, see [Creating a Manual Backup](#).
- After you delete an instance, all of its nodes are deleted.
- Deleted instances will be retained in the recycle bin for a period of time after being released, so you can rebuild the instance from it.

Procedure

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases > GeminiDB**.

Step 3 On the **Instances** page, locate the instance you want to delete and choose **More > Delete** in the **Operation** column.

Step 4 If you have enabled operation protection, click **Start Verification** in the **Delete DB Instance** dialog box. On the displayed page, click **Send Code**, enter the verification code, and click **Verify**. The page is closed automatically.

NOTE

If you enable operation protection, two-factor authentication is required for sensitive operations to secure your account and cloud products. For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

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

Deleted instances are not displayed in the instance list any longer.

----End

3.2.3 Recycling an Instance

You can restore unsubscribed yearly/monthly instances or deleted pay-per-use instances from the recycle bin.

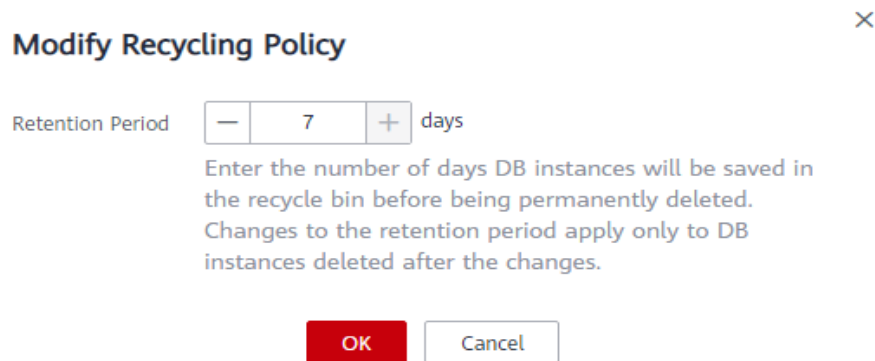
Precautions

- The recycling bin is enabled by default and cannot be disabled. Instances in the recycle bin are retained for 7 days by default, and this will not incur any charges.
- You can put up to 100 instances into the recycle bin. If the maximum number of instances is reached, you cannot put instances into the recycle bin any more.
- If you delete an instance of full storage, the deleted instance will not be moved to the recycle bin.
- You can modify the retention period, and the changes only apply to the DB instances deleted after the changes, so exercise caution when performing this operation.

Modifying the Recycling Policy

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases** > **GeminiDB**.
- Step 3** On the **Recycling Management** page, click **Modify Recycling Policy**. In the displayed dialog box, set the retention period for the deleted DB instances from 1 day to 7 days. Then, click **OK**.

Figure 3-2 Modify Recycling Policy



----End

Rebuilding an Instance

You can rebuild DB instances from the recycle bin within the retention period to restore data.

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases** > **GeminiDB**.
- Step 3** On the **Recycling Bin** page, locate the instance to be rebuilt and in the **Operation** column, click **Rebuild**.

Figure 3-3 Rebuilding an instance

Modify Recycling Policy							
DB Instance Name/ID	DB Instance Type	DB Engine Version	Billing Mode	Created	Deleted	Enterprise Project	Operation
ed55 84bd107d3798445ea05cb08071271a00e02			Pay-per-use	May 20, 2020 09:38:55 GMT...	May 20, 2020 09:42:00 GMT...	default	Rebuild

- Step 4** On the displayed page, set required parameters and submit the rebuilding task.

----End


3.3 Instance Modifications

3.3.1 Upgrading a Minor Version

GeminiDB Cassandra can be upgraded by installing patches to improve performance, release new features, or fix bugs.

If a new patch is released, you can upgrade your instance by clicking the upgrade button in the **Compatible API** column on the **Instances** page.

Figure 3-4 Upgrade button



Precautions

- Upgrade your instance once there are new patches released.
- The instance will be restarted and services may be interrupted during the upgrade. The interruption duration depends on services, quantity of nodes, and the amount of service data. Upgrade your instance during off-peak hours.

- When you upgrade a cluster, services may be interrupted a number of times equal to the number of nodes in the cluster plus one. Each interruption will last for no more than a minute and will only affect the services on that node. The upgrade duration is as follows:

$$600 + (N \times 60) \leq \text{Total upgrade duration (s)} \leq 600 + (N \times 120)$$

For example, if there are 9 nodes in a cluster instance, the upgrade duration is 19 to 28 minutes.


The upgrade duration of most instances is close to 600+ (N x 60). If there are too many tokens on a single node, the upgrade duration may be increased.

- Before you upgrade a DR instance, upgrade the corresponding standby instance first and then the primary instance afterwards.

Procedure

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, locate the instance you want to upgrade and click **Upgrade Minor Version** in the **Compatible API** column.

Figure 3-5 Patch installation



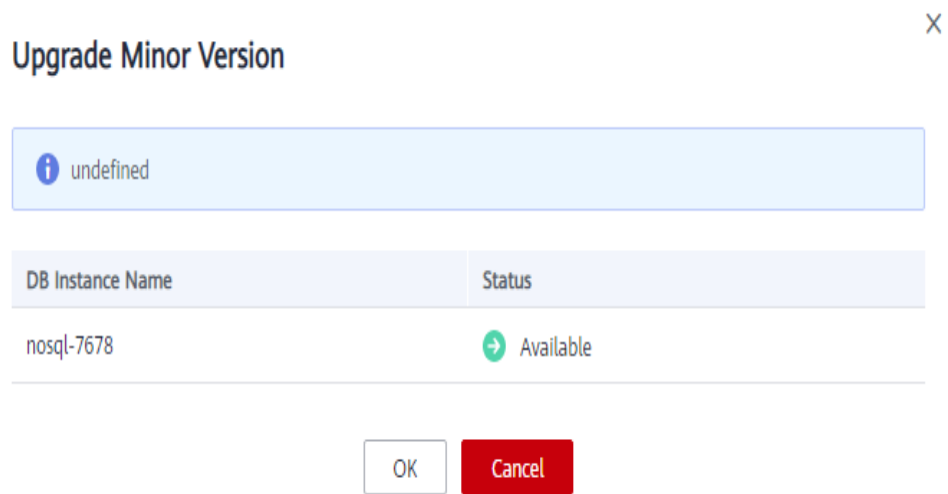
Alternatively, click the instance name to go to the **Basic Information** page. In the **DB Information** area, click **Upgrade Minor Version** in the **Compatible API** field.

Figure 3-6 Patch installation



Step 4 In the displayed dialog box, click **OK**.

Figure 3-7 Confirming dialog box



Step 5 View the upgrade result on the **Instances** page.

- When the upgrade is ongoing, the instance status is **Upgrading minor version**.
- After the upgrade is complete, the instance status changes **Available**.

----End


3.3.2 Changing an Instance Name

This section describes how to change the name of a GeminiDB Cassandra instance to identify different instances.

Method 1

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases > GeminiDB**.

Step 3 On the **Instances** page, click  to the right of the instance whose name you want to modify.

- To submit the change, click **OK**.
- To cancel the change, click **Cancel**.

NOTE

The instance name:

- Can be the same as an existing instance name.
- Can include 4 to 64 bytes and must start with a letter. It is case-sensitive and allows only letters, digits, hyphens (-), and underscores (_).

Step 4 View the results on the **Instances** page.


----End



Method 2

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases > GeminiDB**.

Step 3 On the **Instances** page, locate the instance whose name you want to modify.

Step 4 In the **Instance Information** area on the **Basic Information** page, click  in the **DB Instance Name** field to change the name.

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

NOTE

The instance name:

- Can be the same as an existing instance name.
- Can include 4 to 64 bytes and must start with a letter. It is case-sensitive and allows only letters, digits, hyphens (-), and underscores (_).

Step 5 View the results on the **Instance Management** page.

----End

3.3.3 Resetting the Administrator Password

For security reasons, change administrator passwords periodically.

Precautions

- You can reset the administrator password only when your instance is states **Available**, **Backing up**, **Checking restoration**, or **Scaling up**. You can also choose to reset the password if an instance node becomes abnormal.
- The administrator password takes effect immediately after being reset.
- If you enable operation protection to improve the security of your account and cloud products, two-factor authentication is required for sensitive operations. For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

Method 1

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases > GeminiDB**.

Step 3 On the **Instances** page, locate the instance whose password you want to reset and choose **More > Reset Password** in the **Operation** column.

Step 4 Enter and confirm the new administrator password and click **OK**.

The password must be 8 to 32 characters in length and contain uppercase letters, lowercase letters, digits, and any of the following special characters: ~!@#%^*-_ = +?

Step 5 If you have enabled operation protection, click **Start Verification** in the displayed dialog box. On the displayed page, click **Send Code**, enter the verification code, and click **Verify**. The page is closed automatically.

----End

Method 2

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases > GeminiDB**.

Step 3 On the **Instances** page, click the instance whose password you want to reset to go to the **Basic Information** page.

Step 4 In the **DB Information** area, click **Reset Password** in the **Administrator** field.

Step 5 Enter and confirm the new administrator password and click **OK**.

The password must be 8 to 32 characters in length and contain uppercase letters, lowercase letters, digits, and any of the following special characters: ~!@#%^*-_ = +?

Step 6 If you have enabled operation protection, click **Start Verification** in the displayed dialog box. On the displayed page, click **Send Code**, enter the verification code, and click **Verify**. The page is closed automatically.

----End

3.3.4 Scaling Up Storage Space

This section describes how to scale up storage space of an instance to suit your service requirements.

Precautions

- Storage space can only be scaled up.
- **Storage scaling does not interrupt your services. After storage scaling is complete, you do not need to restart your instance.**

Procedure

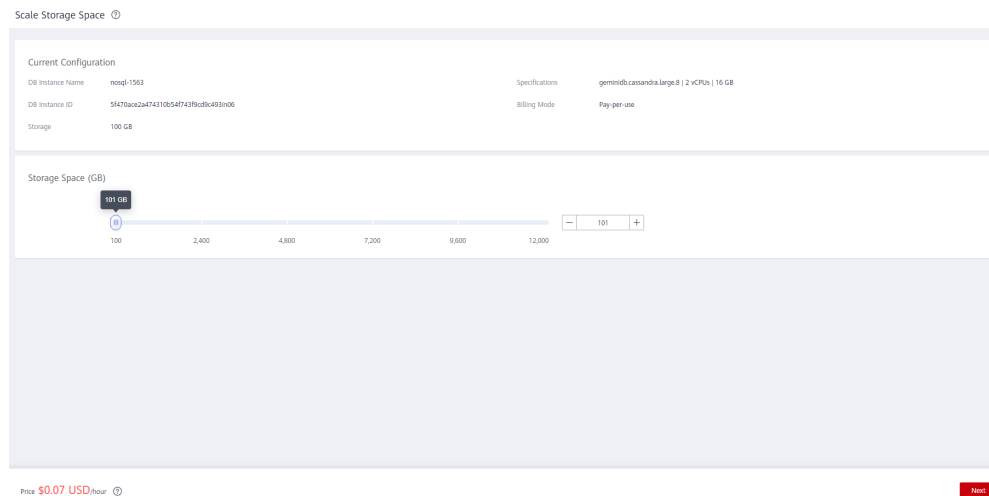
Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases > GeminiDB**.

Step 3 On the **Instances** page, locate the instance whose storage space you want to scale and choose **MoreScale Storage Space** in the **Operation** column.

Click the instance name. In the **Storage Space** area on the **Basic Information** page, click **Scale**.

Step 4 On the displayed page, specify a new storage capacity and click **Next**.

Figure 3-8 Scaling up storage space of an instance

You must select at least 1 GB each time you scale, and only an integer is allowed.

Step 5 On the displayed page, confirm the storage space.

- For yearly/monthly instances
 - If you need to modify your settings, click **Previous**.
 - If you do not need to modify your settings, click **Next** and complete the payment.
- For pay-per-use instances
 - If you need to modify your settings, click **Previous**.
 - If you do not need to modify your settings, click **Submit**.

Step 6 Check the scaling result.

- When the scale-up task is ongoing, the instance status is **Scaling up**.
- After the scale-up task is complete, the instance status becomes **Available**.
- In the **Storage Space** area on the **Basic Information** page, check whether the scale-up was successful.

----End

3.3.5 Configuring Autoscaling

You can enable storage autoscaling for GeminiDB Cassandra instances. When the storage space usage reaches the upper limit, autoscaling is triggered.

You can enable storage autoscaling:

1. When you create an instance. For details, see [Buying an Instance](#).
2. After you create an instance

This section describes how to configure storage autoscaling after an instance is created.

Configuring the Required Permissions

If you are an IAM user, perform the following operations before you enable storage autoscaling:

1. Configure fine-grained permissions for IAM and minimum permissions for GeminiDB.

For details about how to configure IAM permissions, see [Creating a Custom Policy](#).

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:permissions:grantRoleToAgencyOnDomain",
        "iam:permissions:grantRoleToAgencyOnProject",
        "iam:agencies:listAgencies",
        "iam:roles:listRoles",
        "iam:agencies:getAgency",
        "iam:agencies:createAgency",
        "iam:permissions:checkRoleForAgencyOnProject",
        "iam:permissions:listRolesForAgency",
        "iam:quotas:listQuotas",
        "iam:permissions:checkRoleForAgencyOnDomain",
        "iam:permissions:revokeRoleFromAgencyOnProject",
        "iam:permissions:grantRoleToAgency",
        "iam:permissions:listRolesForAgencyOnProject",
        "iam:agencies:updateAgency",
        "iam:permissions:revokeRoleFromAgency",
        "iam:permissions:checkRoleForAgency",
        "iam:permissions:revokeRoleFromAgencyOnDomain",
        "iam:agencies:deleteAgency",
        "iam:permissions:listRolesForAgencyOnDomain",
        "nosql:instance:list",
        "nosql:instance:modifyStorageSize"
      ]
    }
  ]
}
```

2. [Create a user group and assign permissions](#).

You can create a user group on the IAM console and grant it the permissions created in [1](#).

3. [Create an IAM user](#) and add it to a user group.

Log in to the IAM console using the major account or an account with the IAM permissions, locate the IAM user that the target instance belongs to, and add it to the user group created in [2](#). The IAM user will inherit permissions of the user group.


Precautions

- Autoscaling is available when your account balance is sufficient.
- Autoscaling is currently in open beta testing. If you need to use this function, contact customer service to apply for it.
- The instance is in the **Available** status.
- Once autoscaling is enabled, an agency will be created and fees will be automatically deducted.

Autoscaling of a Single Instance

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, click the instance. The **Basic Information** page is displayed.
- Step 4** In the **Storage Space** area, click **Configure Autoscaling**.

Figure 3-9 Configure Autoscaling



- Step 5** Enable the **Configure Autoscaling** toggle and specify the trigger condition, increment, and autoscaling limit.

Figure 3-10 Configuring autoscaling parameters

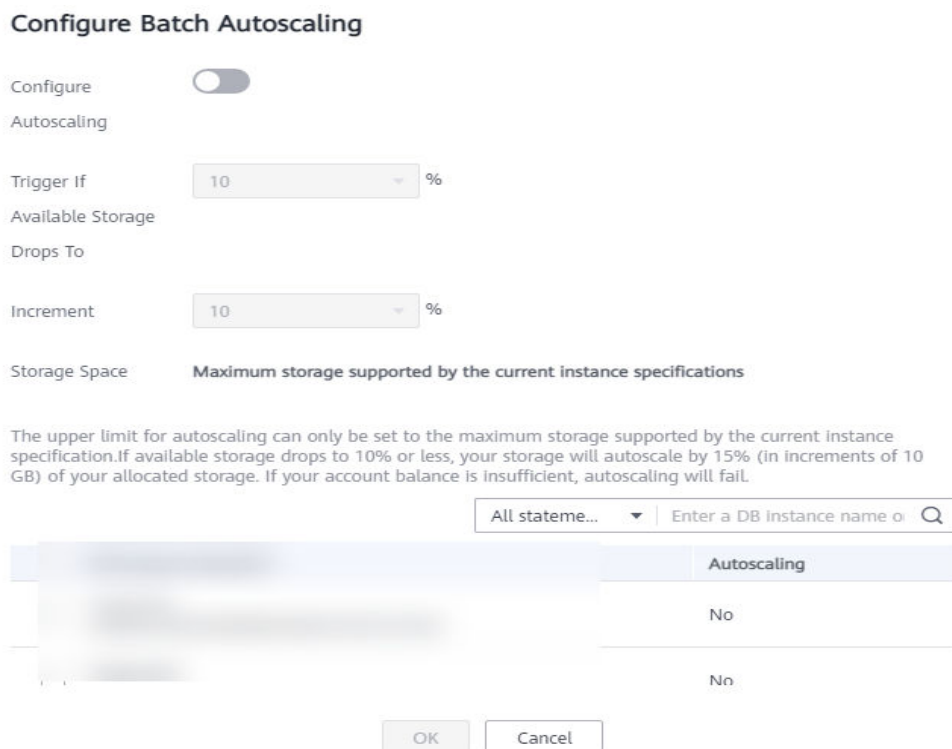


Table 3-1 Parameter description

Parameter	Description
Configure Autoscaling	The autoscaling toggle button.
Trigger If Available Storage Drops To	When the available storage drops to a specified threshold or 10 GB, autoscaling is triggered.
Increment	The percentage of the current storage to be automatically scaled. The value can be 10, 15, or 20. You can select a proper increment as required. At least 100 GB is added each time.
Storage Space	Upper limit of the storage space in GB that can be automatically scaled to The limit must be no less than the storage of your instance and cannot exceed the maximum storage space defined by your instance specifications.

Step 6 Click **OK**.

----End

Autoscaling Storage Space of Instances in Batches

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases > GeminiDB**.

Step 3 Select instances and click **Auto Scale**.

Figure 3-11 Storage autoscaling

The screenshot shows a management console interface. At the top left, there is a button labeled 'Auto Scale'. Below it is a table with the following columns: Name/ID, DB Instance Type, Compatible API, Status, Billing Mode, and Operation. A single instance is listed with Name/ID 'nosql-7689', DB Instance Type 'Cluster', Compatible API 'Cassandra 3.11', Status 'Available', Billing Mode 'Pay-per-use', and Operation 'Created on Nov. 30, 2021...'. There are also search and filter options at the top right of the table.

Name/ID	DB Instance Type	Compatible API	Status	Billing Mode	Operation
nosql-7689	Cluster	Cassandra 3.11	Available	Pay-per-use	Created on Nov. 30, 2021...

Step 4 Select the instance for which you want to configure autoscaling, enable autoscaling, and specify the trigger conditions and autoscaling limit.

Figure 3-12 Configuring autoscaling parameters

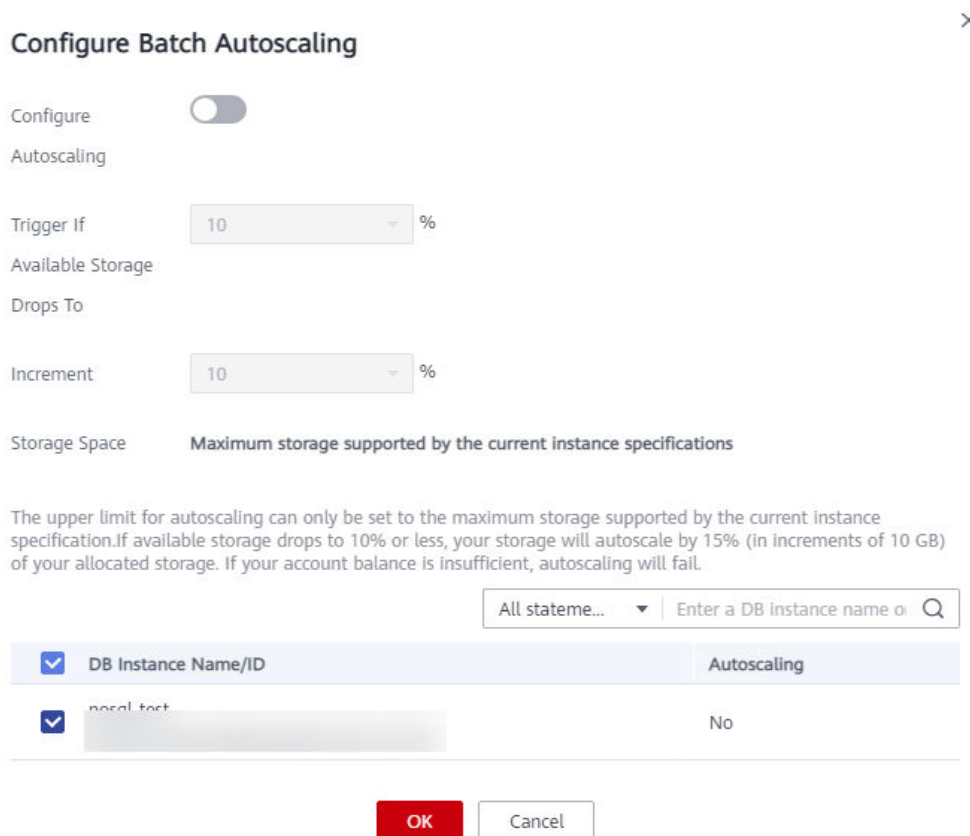


Table 3-2 Parameter description

Parameter	Description
Configure Autoscaling	The autoscaling toggle button.
Trigger If Available Storage Drops To	When the available storage drops to a specified threshold or 10 GB, autoscaling is triggered.
Increment	Preset percentage of currently allocated storage. You can choose to set it to 10 , 15 , or 20 based on service requirements. At least 100 GB is added each time.
Storage Space	Batch autoscaling does not allow you to specify an upper storage limit. The upper limit is the maximum storage defined by your instance specifications by default.

Step 5 Click **OK**.

----End

3.3.6 Changing vCPUs and Memory of an Instance

This section describes how to change your instance vCPUs and memory to suit your service requirements.

Precautions

- Instances can be scaled up or down by changing their specifications.
- If one instance has multiple nodes, the change will be performed on the nodes one by one. It takes about 5 to 10 minutes for each node, and the total time required depends on the number of the nodes.
- For a node whose specifications are being changed, its computing tasks are handed over to other nodes. Change specifications of nodes during off-peak hours to prevent the instance from overload.
- Do not perform DDL operations when you change the instance specifications.

NOTE

A data definition language (DDL) is a language for defining data structures and database objects. Common examples of DDL statements are CREATE, ALTER, and DROP. Data Definition Language (DDL) is used to create, modify, and delete database objects, such as tables, indexes, views, functions, stored procedures, and triggers.

Method 1

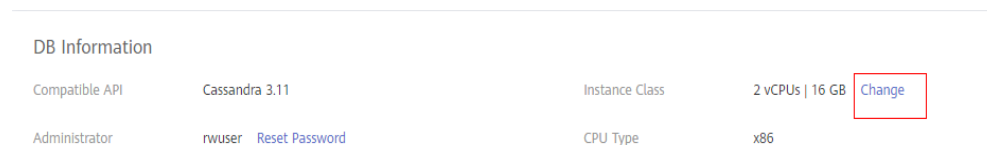
Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases > GeminiDB**.

Step 3 On the **Instances** page, locate the instance whose vCPUs and memory you want to change and click its name.

Step 4 In the **DB Information** area on the **Basic Information** page, click **Change** next to the **Instance Class** field.

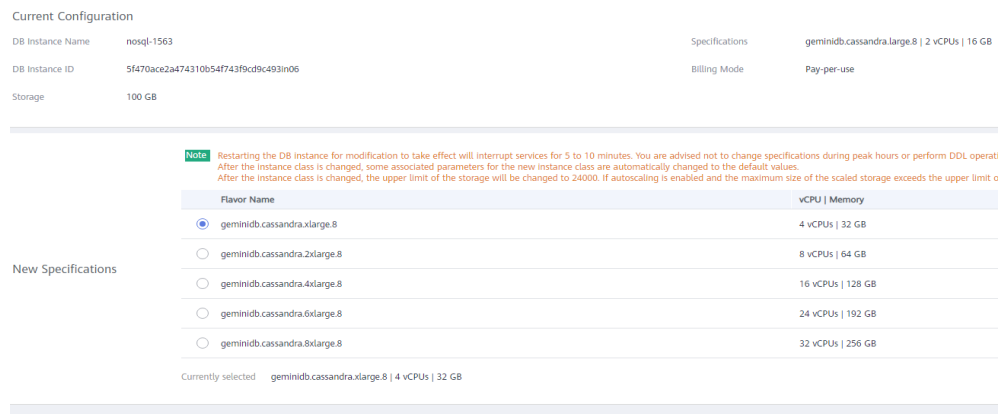
Figure 3-13 Changing specifications



DB Information			
Compatible API	Cassandra 3.11	Instance Class	2 vCPUs 16 GB Change
Administrator	rwuser Reset Password	CPU Type	x86

Step 5 On the displayed page, select the required specifications and click **Next**.

Figure 3-14 Changing specifications



Step 6 On the displayed page, confirm the instance specifications.

- For yearly/monthly instances
 - If you need to modify your settings, click **Previous**.
 - If you do not need to modify your settings, click **Submit**. If you are scaling up the instance specifications, go to the payment page, select a payment method, and complete the payment.
- For pay-per-use instances
 - If you need to modify your settings, click **Previous**.
 - If you do not need to modify your settings, click **Submit**.

Step 7 View the change result.

Go to the **Basic Information** page and in the **DB Information** area you can see the new instance specifications.

----End


Method 2

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases > GeminiDB**.

Step 3 On the **Instances** page, locate the instance whose specifications you want to change and choose **More > Change Specifications** in the **Operation** column.

Figure 3-15 Changing specifications



Step 4 On the displayed page, select the required specifications and click **Next**.

Figure 3-16 Changing specifications

Current Configuration		Specifications	
DB Instance Name	nosql-1563	Specifications	geminidb.cassandra.large.8 2 vCPUs 16 GB
DB Instance ID	5f470ace2a474310b54743f9cd9c493in06	Billing Mode	Pay-per-use
Storage	100 GB		

Note Restarting the DB instance for modification to take effect will interrupt services for 5 to 10 minutes. You are advised not to change specifications during peak hours or perform DDL operations. After the instance class is changed, some associated parameters for the new instance class are automatically changed to the default values. After the instance class is changed, the upper limit of the storage will be changed to 24000. If autoscaling is enabled and the maximum size of the scaled storage exceeds the upper limit of the storage, the storage will be automatically scaled to the upper limit.

Flavor Name	vCPU Memory
<input checked="" type="radio"/> geminidb.cassandra.xlarge.8	4 vCPUs 32 GB
<input type="radio"/> geminidb.cassandra.2xlarge.8	8 vCPUs 64 GB
<input type="radio"/> geminidb.cassandra.4xlarge.8	16 vCPUs 128 GB
<input type="radio"/> geminidb.cassandra.6xlarge.8	24 vCPUs 192 GB
<input type="radio"/> geminidb.cassandra.8xlarge.8	32 vCPUs 256 GB

New Specifications

Currently selected: geminidb.cassandra.xlarge.8 | 4 vCPUs | 32 GB

Step 5 On the displayed page, confirm the instance specifications.

- For yearly/monthly instances
 - If you need to modify your settings, click **Previous**.
 - If you do not need to modify your settings, click **Submit**. If you are scaling up the instance specifications, go to the payment page, select a payment method, and complete the payment.
- For pay-per-use instances
 - If you need to modify your settings, click **Previous**.
 - If you do not need to modify your settings, click **Submit**.

Step 6 View the change result.

Go to the **Basic Information** page and in the **DB Information** area you can see the new instance specifications.

----End

3.3.7 Adding Nodes

This section describes how to add nodes to an instance to suit your service requirements.

Precautions

- Adding nodes may lead to the decrease of operations per second (OPS). Perform this operation during off-peak hours.
- You can only add nodes when the instance status is **Available** or **Checking restoration**.
- Instances that one or more nodes are added to cannot be deleted.
- You can also delete nodes as required. For details, see [Deleting Nodes](#).

Method 1

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases** > **GeminiDB**.

Step 3 On the **Instances** page, locate the instance that you want to add nodes to and click its name.

Step 4 In the **Node Information** area on the **Basic Information** page, click **Add Node**.

Figure 3-17 Basic information

Node Information

Name/ID	Status	AZ	Private IP Address	EIP	Operation
	Creating	az2		Unbound	View Metric Bind EIP Delete
	Creating	az1		Unbound	View Metric Bind EIP Delete
	Creating	az3		Unbound	View Metric Bind EIP Delete

Step 5 Specify **Add Nodes** and click **Next**.

Figure 3-18 Adding nodes

<input type="checkbox"/>	Name/ID	DB Instance Type	Compatible API	Status	Enterprise Project	Billing Mode	Operation
<input type="checkbox"/>	node-1563 5f470ac2a474310b54f743f9c9ca493in06	Cluster	Cassandra 3.11	Available	default	Pay-per-Use Created on Feb 16, 202...	Change to Yearly/Monthly Change Specifications More
<input type="checkbox"/>	node-6882 6d3d3331f17f7d682b0778b4976a3172in12	Proxy-based general pu...	Redis 5.0	Available	default	Pay-per-Use Created on Feb 16, 202...	Change to Yearly/Monthly Cha...

Create Backup

Scale Storage Space

Add Node

Restart

Reset Password

Delete

NOTE

- New nodes are of the same specifications as existing nodes. Once a new node is added, its specifications cannot be changed.
- New nodes and the instance can be in different subnets of the same VPC.

Step 6 On the displayed page, confirm the node configurations.

- For yearly/monthly instances
 - If you need to modify your settings, click **Previous**.
 - If you do not need to modify your settings, click **Next** and complete the payment.
- For pay-per-use instances
 - If you need to modify your settings, click **Previous**.
 - If you do not need to modify your settings, click **Submit**.

Step 7 View the results.

- When new nodes are being added, the instance status is **Adding node**.
- After the nodes are added, the instance status becomes **Available**.
- Click the instance name. In the **Node Information** area on the **Basic Information** page, view information about the new nodes.

----End

Method 2

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases** > **GeminiDB**.

Step 3 On the **Instance Management** page, locate the instance you want to add nodes for and choose **More > Add Node** in the **Operation** column.

Figure 3-19 Adding nodes

Name/ID	DB Instance Type	Compatible API	Status	Enterprise Project	Billing Mode	Operation
rncj-1563 9470ba2a47431055474395c9c483106	Cluster	Cassandra 3.11	Available	default	Pay-per-Use Created on Feb 16, 202...	Change to Yearly/Monthly Change Specifications More
rncj-6582 6d3d3331f174682b07778b4974a31721n12	Proxy-based general pu...	Redis 5.0	Available	default	Pay-per-Use Created on Feb 16, 202...	Change to Yearly/Monthly Cha

Step 4 Specify **Add Nodes** and click **Next**.

Figure 3-20 Adding nodes

Name/ID	DB Instance Type	Compatible API	Status	Enterprise Project	Billing Mode	Operation
rncj-1563 9470ba2a47431055474395c9c483106	Cluster	Cassandra 3.11	Available	default	Pay-per-Use Created on Feb 16, 202...	Change to Yearly/Monthly Change Specifications More
rncj-6582 6d3d3331f174682b07778b4974a31721n12	Proxy-based general pu...	Redis 5.0	Available	default	Pay-per-Use Created on Feb 16, 202...	Change to Yearly/Monthly Cha

NOTE

- New nodes are of the same specifications as existing nodes. Once a new node is added, its specifications cannot be changed.
- New nodes and the instance can be in different subnets of the same VPC.

Step 5 On the displayed page, confirm the node configurations.

- For yearly/monthly instances
 - If you need to modify your settings, click **Previous**.
 - If you do not need to modify your settings, click **Next** and complete the payment.
- For pay-per-use instances
 - If you need to modify your settings, click **Previous**.
 - If you do not need to modify your settings, click **Submit**.

Step 6 View the results.

- When new nodes are being added, the instance status is **Adding node**.
- After the nodes are added, the instance status becomes **Available**.
- Click the instance name. In the **Node Information** area on the **Basic Information** page, view information about the new nodes.

----End

3.3.8 Deleting Nodes

You can delete nodes that are no longer used to release resources.

Precautions

- Deleted nodes cannot be recovered. Exercise caution when performing this operation.

- Only pay-per-use instances can be deleted.
- If you enable operation protection to improve the security of your account and cloud products, two-factor authentication is required for sensitive operations. For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

Procedure

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, locate the instance that you want to delete nodes from and click its name.
- Step 4** In the **Node Information** area on the **Basic Information** page, locate the node you want to delete and click **Delete** in the **Operation** column.
- Step 5** If you have enabled operation protection, click **Start Verification** in the **Delete Node** dialog box. On the displayed page, click **Send Code**, enter the verification code, and click **Verify**. The page is closed automatically.
- Step 6** In the displayed dialog box, click **Yes**.
- When the node is being deleted, the instance status is **Deleting node**.
 - After the node is deleted, the instance status becomes **Available**.

----End

3.3.9 Managing Tags

Tag Management Service (TMS) enables you to manage resources using tags on the management console. TMS works with other cloud services to manage global tags, and other cloud services manage their own tags.

Adding tags to GeminiDB Cassandra instances helps you better identify and manage them. An instance can be tagged when or after it is created.

After an instance is tagged, you can search for the tag key or value to quickly query the instance details.

Precautions

- You are advised to set predefined tags on the TMS console.
- A tag consists of a key and value. You can add only one value for each key. For details about the naming rules of tag keys and tag values, see [Table 3-3](#).
- Each instance can have up to 20 tags by default.
- The tag name must comply with the naming rules described in [Table 3-3](#).

Table 3-3 Naming rules

Parameter	Requirement	Example Value
Tag key	<ul style="list-style-type: none"> • Cannot be left blank. • Must be unique for each instance. • Contains a maximum of 36 characters. • Can only consist of digits, letters, underscores (_), and hyphens (-). 	Organization
Tag value	<ul style="list-style-type: none"> • Can be left blank. • Contains a maximum of 43 characters. • Can only consist of digits, letters, underscores (_), periods (.), and hyphens (-). 	nosql_01

Adding a Tag

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, locate the instance that you want to add tags to and click its name. The **Basic Information** page is displayed.
- Step 4** In the navigation pane on the left, click **Tags**.
- Step 5** On the **Tags** page, click **Add Tag**. In the displayed dialog box, enter a tag key and value, and click **OK**.
- Step 6** View and manage tags on the **Tags** page.
- End

Editing a Tag

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, locate the instance whose tags you want to edit and click its name. The **Basic Information** page is displayed.
- Step 4** In the navigation pane on the left, click **Tags**.
- Step 5** On the **Tags** page, locate the tag that you want to edit and click **Edit** in the **Operation** column. In the displayed dialog box, change the tag value and click **OK**.
- Only the tag value can be edited.
- Step 6** View and manage tags on the **Tags** page.
- End

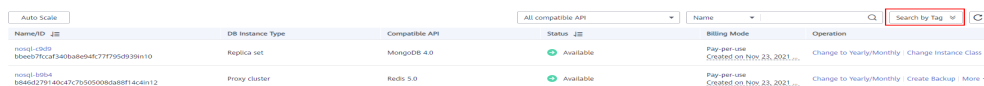
Deleting a Tag

- Step 1** Log in to the management console.
 - Step 2** In the service list, choose **Databases > GeminiDB**.
 - Step 3** On the **Instances** page, locate the instance whose tags you want to delete and click its name. The **Basic Information** page is displayed.
 - Step 4** In the navigation pane on the left, click **Tags**.
 - Step 5** On the **Tags** page, locate the tag that you want to delete and click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes**.
 - Step 6** Check whether the deleted tag is displayed on the **Tags** page.
- End

Searching by tag

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, click **Search by Tag** in the upper right corner of the instance list.

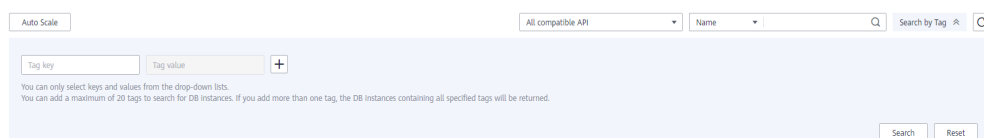
Figure 3-21 Search by Tag



Name/ID	DB Instance Type	Compatible API	Status	Billing Mode	Operation
msuj-csdr bb0eb7fcca340ba0e540c77f95d9391n10	Replica set	MongoDB 4.0	Available	Pay-per-use Created on Nov. 23, 2021...	Change to Yearly/Monthly · Change Instance Class
msuj-lsdr b046d279140c47c79505008d588f14c4n12	Proxy cluster	Redis 5.0	Available	Pay-per-use Created on Nov. 23, 2021...	Change to Yearly/Monthly · Create Backup · More

- Step 4** Enter a tag key or value and click **Search** to query the instance associated with the tag.

Figure 3-22 Searching by tag key



Auto Scale All compatible API

Tag key Tag value

You can only select keys and values from the drop-down lists.
You can add a maximum of 20 tags to search for DB instances. If you add more than one tag, the DB instances containing all specified tags will be returned.

----End

3.3.10 Updating the OS of an Instance

To improve database performance and security, the OS of a GeminiDB Cassandra instance needs to be updated timely.

Every time you upgrade the kernel version of your instance, GeminiDB Cassandra determines whether to update the OS and selects the right cold patch to upgrade the OS if necessary.

Updating the OS does not change the DB instance version or other information.

In addition, GeminiDB Cassandra installs hot patches as required to fix major OS vulnerabilities within the maintenance window you specified.

3.4 Connection Management

3.4.1 Configuring Security Group Rules

A security group is a collection of access control rules for ECSs and GeminiDB Cassandra instances that have the same security protection requirements and are mutually trusted in a VPC.

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

This section describes how to configure security group rules when you connect to a GeminiDB Cassandra instance over private and public networks.

Precautions

- Each account can create up to 500 security group rules by default.
- Too many security group rules will increase the first packet latency, so a maximum of 50 rules for each security group is recommended.
- One security group can be associated with only one GeminiDB Cassandra instance.
- For details about security group rules, see [Table 3-4](#).

Table 3-4 Parameter description

Scenario	Description
Connecting to an instance over a private network	<p>Check whether the ECS and GeminiDB Cassandra instance are in the same security group:</p> <ul style="list-style-type: none"> • If yes, no security group rules need to be configured. • If no, configure security group rules for them, respectively. <ul style="list-style-type: none"> – GeminiDB Cassandra instance: Configure inbound rules for its security group. For details, see Procedure. – ECS: The default security group rule allows all outbound data packets, so you do not need to configure a security rule for the ECS. If not all outbound traffic is allowed in the security group, configure an outbound rule for the ECS.
Connecting to an instance over a public network	<p>Add inbound rules for the security group associated with the GeminiDB Cassandra instance. For details, see Procedure.</p>


Procedure

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, locate the instance that you want to configure security group rules for and click its name.
- Step 4** Configure security group rules.

Method 1

In the **Network Information** area on the **Basic Information** page, click the name of security group.

Figure 3-23 Security group



- Step 5** Add an inbound rule.

1. Click the **Inbound Rules** tab.

Figure 3-24 Inbound rules

The screenshot shows the 'Inbound Rules' tab in a management console. It displays a table with the following columns: 'Protocol & Port', 'Type', 'Source', and 'Operation'. There are three rows of rules listed:

Protocol & Port	Type	Source	Operation
All	IPv4	0.0.0.0/0	Modify, Replicate, Delete
TCP: 22	IPv4	0.0.0.0/0	Modify, Replicate, Delete
TCP: 3389	IPv4	0.0.0.0/0	Modify, Replicate, Delete

2. Click **Add Rule**. The **Add Inbound Rule** dialog box is displayed.

Figure 3-25 Adding a rule

The screenshot shows the 'Add Inbound Rule' dialog box. At the top, it says 'Add Inbound Rule' with a link to 'Learn more about security group configuration.' Below this is a blue information box: 'Inbound rules allow incoming traffic to instances associated with the security group.' The 'Security Group' is set to 'dds-st-test-security-group'. A note says 'You can import multiple rules in a batch.' The main form has four columns: 'Protocol & Port', 'Type', 'Source', and 'Operation'. The 'Protocol & Port' dropdown is set to 'TCP' with an example '22 or 22-30'. The 'Type' dropdown is set to 'IPv4'. The 'Source' dropdown is set to 'IP address' with a text input field containing '0.0.0.0/0'. The 'Operation' dropdown is set to 'Operation'. At the bottom, there is an 'Add Rule' button, an 'OK' button, and a 'Cancel' button.

3. In the displayed **Add Rule** dialog box, set required parameters.

Table 3-5 Inbound rule settings

Parameter	Description	Example Value
Protocol & Port	<ul style="list-style-type: none"> - Network protocol. Available options are All, TCP, UDP, ICMP, or GRE - Port: The port or port range that allows the access to the ECS. Range: 1 to 65535 	TCP
Type	IP address type. This parameter is available only after the IPv6 function is enabled. <ul style="list-style-type: none"> - IPv4 - IPv6 	IPv4
Source	Source address. It can be a single IP address, an IP address group, or a security group to allow access from the IP address or instances in the security group. Example: <ul style="list-style-type: none"> - Single IP address: xxx.xxx.xxx.xxx/32 (IPv4) - Subnet: xxx.xxx.xxx.0/24 - All IP addresses: 0.0.0.0/0 - sg-abc (security group) 	0.0.0.0/0
Description	(Optional) Provides supplementary information about the security group rule. The description can contain up to 255 characters and cannot contain angle brackets (<>).	-

Step 6 Click **OK**.

----End

3.4.2 Binding and Unbinding an EIP

The Elastic IP service provides independent public IP addresses and bandwidth for public access. After you create a GeminiDB Cassandra instance, you can bind an EIP to it to allow external access. If later you want to prohibit external access, you can also unbind the EIP from the instance.

Precautions

- This function is in the open beta test (OBT) phase. To use this function, contact customer service.
- To change the EIP that has been bound to a node, unbind it from the node first.

Binding an EIP

Step 1 Log in to the management console.

- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, locate the GeminiDB Cassandra instance that you want to bind an EIP to and click its name.
- Step 4** On the **Basic Information** page, in the **Node Information** area, locate the target node and click **Bind EIP** in the **Operation** column.

Figure 3-26 Binding an EIP

Name/ID	Status	AZ	Subnet	Private IP Address	EIP	Operation
...	Available	az4	subnet-44d7(192.168.0.0/24)	...	Unbound	View Metric: Bind EIP
...	Available	az4	subnet-44d7(192.168.0.0/24)	...	Unbound	View Metric: Bind EIP

- Step 5** In the displayed dialog box, select the required EIP and click **Yes**. If no available EIPs are displayed, click **View EIP** and create an EIP on the VPC console.
- Step 6** In the **EIP** column, view the EIP that is successfully bound.
- To unbind the EIP from the instance, see [Unbinding an EIP](#).

----End

Unbinding an EIP

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, locate the GeminiDB Cassandra instance that you want to unbind an EIP from and click its name.
- Step 4** On the **Basic Information** page, in the **Node Information** area, locate the target node and click **Unbind EIP** in the **Operation** column.

Figure 3-27 Unbinding an EIP

Name/ID	Status	AZ	Subnet	Private IP Address	EIP	Operation
...	Available	az4	subnet-44d7(192.168.0.0/24)	...	10.90.9.80	View Metric: Unbind EIP
...	Available	az4	subnet-44d7(192.168.0.0/24)	...	Unbound	View Metric: Bind EIP

- Step 5** In the displayed dialog box, click **Yes**.
- To bind an EIP to the instance again, see [Binding an EIP](#).

----End

3.4.3 Viewing the IP Address and Port Number

This section describes how to query the IP address and port number of a GeminiDB Cassandra instance on the management console.

Procedure

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, locate the instance whose IP address and port you want to view and click its name.

Method 1

In the **Node Information** area on the **Basic Information** page, view the private IP address or EIP of each node in the instance.

Figure 3-28 Viewing IP addresses

Name/ID	Status	AZ	Subnet	Private IP Address	EIP	Operation
nosql-90cb_cpy_priam_node_1 dcd18e398bc34fe501b2019eb0e513...	Available	a24		0.0.0.0		View Metric Unbind EIP
nosql-90cb_cpy_priam_node_2 e84364429815455aa6b643414921...	Available	a24			Unbound	View Metric Bind EIP

In the **Network Information** area, view the port number of the instance. The default port is 8635.

Figure 3-29 Viewing the port number

VPC	vpc-demo	Security Group	Sys-default
Subnet		Database Port	8635
Address	192.168.0.116:8635, 192.168.0.160:8635		

Method 2

-----End

3.5 Data Management

3.5.1 Importing and Exporting Data by Running COPY

COPY is one of cqlsh commands. It includes **COPY TO** and **COPY FROM**. They are used to copy data to and from Cassandra.

COPY TO can export data from a table to a CSV, Parquet, or ORC file.

- If the exported file is in CSV format, it needs to be written into the target file by row, and fields are separated by delimiters.
- If no field name is specified, all fields are exported.
- To skip some fields, specify a field list.

COPY FROM allows you to import data from a CSV file to an existing table.

- The source file is imported by row.
- All rows in the dataset must contain the same number of fields, and the PRIMARY KEY field must have a value. During the import, the PRIMARY KEY field will be verified and the existing records are updated.
- If HEADER is set to **False** and no field name is specified, fields are imported in a specified order. After field names are specified, the fields are imported in sequence. The missing and empty fields are set to null.
- The source file can only have fewer fields than the target table.
- When only COPY FROM is used to import data, the number of rows in a dataset cannot exceed 2 million.

Precautions

- Import and export data during off-peak hours to minimize the impacts on your services.
- Obtain the latest binary package by referring to [Connecting to an Instance over a Private Network](#).

COPY Syntax

- **COPY TO**
COPY *table_name* [(*column_list*)] TO '*file_name*' [, '*file2_name*', ...] | STDOUT [WITH option = '*value*' [ADN ...]]
- **COPY FROM**
COPY *table_name* [(*column_list*)] FROM '*file_name*' [, '*file2_name*', ...] | STDIN [WITH option = '*value*' [ADN ...]]

NOTE

COPY supports one or more comma-separated file names or a list of Python glob expressions.

For some common syntax symbols in the COPY command, see [Table 3-6](#).

Table 3-6 Symbol conventions

Symbol	Description
Uppercase letters	Text keyword.
Lowercase letters	A variable, which needs to be replaced with a user-defined value.
Italic	(Optional) Enclose optional command parameters in square brackets ([]). Do not enter only square brackets.
()	Group. Parentheses (()) indicate the group to be selected. Do not input only brackets.
	Or. Use vertical bars () to separate elements. You can input any element. Do not enter only vertical bars.

Symbol	Description
...	Repeatable. The ellipsis (...) indicates that you can repeat syntax elements multiple times as required.
' <i>Literal string</i> '	The single quotation marks (') must contain the character string in the CQL statement. Use single quotation marks to keep uppercase letters.
{ key : value }	The map set. Include a map set or key-value pair in braces ({}). Separate keys and values with colons.
< <i>datatype1,datatype2</i> >	Set, list, map, or tuple of an ordered list. Angle brackets (< >) contain data types in collections, lists, maps, or tuples. Data types are separated by commas (,).
<i>cql_statement</i> ;	End a CQL statement. Semicolons (;) end all CQL statements.
[--]	Use two hyphens (--) to separate command line options from command arguments. This syntax is useful when parameters may be mistaken for command arguments.
' < <i>schema</i> > ... </ <i>schema</i> > '	Search CQL only; single quotation marks (') enclose the entire XML schema declaration.
@ <i>xml_entity</i> =' <i>xml_entity_type</i> '	Search CQL only; identify entities and literal values to overwrite XML elements in schemas and solrConfig files.

COPY Usage Suggestions

Table 3-7 Description

Command	Parameter	Description	Default Value	Applicability
TO/FROM	DELIMITER	A single character used to separate fields.	English comma,	-
TO/FROM	QUOTE	A single character that contains a field value.	"	-
TO/FROM	ESCAPE	Escapes a single character using the QUOTE character.	\	-

Command	Parameter	Description	Default Value	Applicability
TO/FROM	HEADER	Boolean value (true false), indicating the name of the column in the first row. True matches the field name with the imported column name and inserts the column name into the first row of the exported data.	FALSE	-
TO/FROM	NULL	Filled value of the field whose query result is empty. You can set this parameter as required.	Empty string ()	-
TO/FROM	DATETIMEFORMAT	Time format for reading or writing CSV time data. The timestamp is in the strftime format. If this parameter is not set, the default value is the value of time_format in the cqlshrc file. Default format: %Y-%m-%d %H: %M: %S %z.	%Y-%m-%d %H:%M:%S%z	-
TO/FROM	MAXATTEMPTS	Maximum number of retry times when an error occurs.	5	-
TO/FROM	REPORTFREQUENCY	Frequency of displaying the status, in seconds.	0.25	-
TO/FROM	DECIMALSEP	Delimiter character for decimal values.	English full stop.	-
TO/FROM	THOUSANDSSEP	Separator of a thousand array.	None	-

Command	Parameter	Description	Default Value	Applicability
TO/FROM	BOOLEAN	Boolean values indicate True and False. The value is case-insensitive. For example, the values yes and no have the same effect as values YES and NO .	True,False	-

Command	Parameter	Description	Default Value	Applicability
TO/FROM	NUMPROCESSES	Number of working processes.	16	<p>The default value of this parameter is the number of kernels on the computer minus one. There is no maximum value for this parameter.</p> <p>You can run the dstat and dstat -lvrn 10 commands to check the CPU idle time. If the CPU idle time exists, use the default number of working processes. You can increase the number of processes while observing the CPU usage of the instance. It is recommended that the CPU usage be less than or equal to 60%. If the CPU usage of the executor is idle and the CPU usage of the instance exceeds the recommended value, expand the capacity to further improve the performance.</p>

Command	Parameter	Description	Default Value	Applicability
TO/ FROM	CONFIGFILE	Specifies a cqlshrc configuration file to set the WITH option. NOTE Command line options always overwrite the cqlshrc file.	None, user-defined	-
TO/ FROM	RATEFILE	Prints the output statistics to this file.	None, user-defined	You are advised to add this parameter when exporting data to improve statistics efficiency.
TO/ FROM	ORIGIN	Check whether the database to be imported or exported is an open-source Cassandra database. <ul style="list-style-type: none"> If the open-source Cassandra is used, the value is True. If GeminiDB Cassandra is used, the value is False. 	False	-
FROM	CHUNKSIZE	The block size is passed to the worker process.	5000	This parameter specifies the number of rows sent from the Feeder process (reading data from files) to the worker process. Depending on the average row size of the dataset, it may be advantageous to increase the value of this parameter.

Command	Parameter	Description	Default Value	Applicability
FROM	INGESTRATE	Approximate import rate per second.	100000	INGESTRATE indicates the rate (in rows) at which the feeder process sends data to the worker process per second. Generally, you do not need to change the value unless the rate is too high and needs to be limited.

Command	Parameter	Description	Default Value	Applicability
FROM	MAXBATCHSIZE	Maximum size of a batch file to be imported.	20	<p>The value of this parameter can be as large as possible but cannot exceed the upper limit.</p> <ul style="list-style-type: none"> • MAXBATCHSIZE x The size of a single row < batch_size_file_threshold_in_kb. • If the batch size is too large, an alarm will be reported and rejected. • Set the following parameters in cassandra.yaml: batch_size_warn_threshold_in_kb (The current value is 5.) batch_size_file_threshold_in_kb (The current value is 50.)

Command	Parameter	Description	Default Value	Applicability
FROM	MINBATC HSIZE	Minimum size of a batch import file.	2	For each chunk, the worker process writes data in batches based on the minimum batch size. The value may need to be adjusted based on the block size, number of nodes in the cluster, and number of VNODEs on each node. If the chunk size is larger, increase the value accordingly.
FROM	MAXROW S	Maximum number of rows. The value -1 indicates that there is no upper limit.	-1	-
FROM	SKIPROW S	Number of rows to skip.	0	-
FROM	SKIPCOLS	A comma-separated list of column names to skip.	None, user-defined	-
FROM	MAXPARS EERRORS	Maximum number of global parsing errors. The value -1 indicates that there is no upper limit.	-1	-
FROM	MAXINSE RTERROR S	Maximum number of global insertion errors. The value -1 indicates that there is no upper limit.	-1	-

Command	Parameter	Description	Default Value	Applicability
FROM	ERRFILE	A file that stores all rows that are not imported. If no value is set, the information is stored in import_ks_table.err , where ks is the key space and table is the table name.	import_ks_table.err	-
FROM	TTL	The time to live is in seconds. By default, data does not expire.	3600	-
TO	ENCODING	Output character string type.	UTF-8	-
TO	PAGESIZE	Size of the page for obtaining results.	1000	Size of the result page. The value is an integer. The default value is 1000 . The larger the page size, the longer the value of pagetimeout. If the data volume in a single row is large, set this parameter to a smaller value. If the data volume in a single row is small, set this parameter to a larger value. The best effect of this value depends on the local batch write capability of the executor. If the local batch write capability is strong.

Command	Parameter	Description	Default Value	Applicability
TO	PAGETIMEOUT	The page times out to obtain the result.	10	<p>The value is an integer, indicating the timeout interval for obtaining each page. The unit is second. The default value is 10 seconds.</p> <ul style="list-style-type: none"> • For a large page size or a large partition, increase the value of this parameter. • If a timeout occurs, increase the value of this parameter. • If the server times out, an exponential backoff policy is automatically initiated to prevent the server from being further overloaded, so you may notice the delay. The driver also generates a timeout. In this case, the driver does not know whether the server discards the request or returns the

Command	Parameter	Description	Default Value	Applicability
				<p>result later. There is a low probability that data may be lost or duplicated. Increasing the value of this parameter is helpful in preventing driver build timeouts.</p>
TO	BEGINTOKEN	Minimum token for exporting data.	None, user-defined	<p>The value is a string, indicating the minimum token to be considered during data export.</p> <p>Records with smaller tokens will not be exported.</p> <p>The default value is empty, indicating that there is no minimum token.</p>

Command	Parameter	Description	Default Value	Applicability
TO	ENDTOKEN	Maximum token used to export data.	None, user-defined	<p>The value is a string, indicating the maximum number of tokens to be considered during data export.</p> <p>Records with larger tokens will not be exported.</p> <p>This parameter is left empty by default, indicating that there is no maximum token.</p>
TO	MAXREQUESTS	Maximum number of requests that can be processed concurrently by each worker.	6	<p>The value of this parameter is an integer, indicating the maximum number of running requests that can be processed by each working process.</p> <p>Total degree of parallelism during data export = Number of working processes x Value of this parameter.</p> <p>Default value: 6 Each request will export data for the entire token range.</p>

Command	Parameter	Description	Default Value	Applicability
TO	MAXOUT PUTSIZE	Maximum size of an output file, in lines. After this parameter is set, the output file is split into multiple segments when the size of the output file exceeds the value of this parameter. The value -1 indicates that there is no upper limit.	-1	The value of this parameter is an integer, indicating the maximum size of an output file in the unit of lines. If the value of this parameter is exceeded, the output file is split into multiple segments. The default value is -1, indicating that there is no limit on the maximum value. Therefore, the file is the only output file. This parameter can be used together with MAXFILESIZE.

Command	Parameter	Description	Default Value	Applicability
TO	MAXFILE SIZE	Maximum size of an output file, in KB. After this parameter is set, the output file is split into multiple segments when the size of the output file exceeds the value of this parameter.	None, user-defined	The value of this parameter is an integer, indicating the maximum size of an output file in bytes. The final file size is close to the value of this parameter. If the file size exceeds this value, the output file is split into multiple segments. The default value is -1, indicating that there is no limit on the maximum value. Therefore, the file is the only output file. This parameter can be used together with MAXOUTPUTSIZE.
TO	dataformats	Output file format. Currently, this parameter can only be set to json.	None, user-defined	-
TO	DATATYPE	The file format can be Parquet or ORC.	None, user-defined	-
TO	RESULTFILE	The exported file containing detailed results.	None, user-defined	You are advised to add this parameter when exporting data to improve statistics efficiency.
TO	wherecondition	Export condition specified during the export.	None, user-defined	-

Procedure

The following uses an example to describe how to preconfigure data, export data, and import data.

Step 1 Pre-configuring Data

1. Create a keyspace.

```
CREATE KEYSPACE cycling WITH replication = {'class': 'SimpleStrategy', 'replication_factor': 3};
```

2. Create a table.

```
CREATE TABLE cycling.cyclist_name (  
  id UUID PRIMARY KEY,  
  lastname text,  
  firstname text  
);
```

3. Insert a data record.

```
INSERT INTO cycling.cyclist_name (id, lastname, firstname) VALUES  
(5b6962dd-3f90-4c93-8f61-eabfa4a803e2, 'VOS','Marianne');  
INSERT INTO cycling.cyclist_name (id, lastname, firstname) VALUES (e7cd5752-bc0d-4157-  
a80f-7523add8dbcd, 'VAN DER BREGGEN','Anna');  
INSERT INTO cycling.cyclist_name (id, lastname, firstname) VALUES (e7ae5cf3-d358-4d99-  
b900-85902fda9bb0, 'FRAME','Alex');  
INSERT INTO cycling.cyclist_name (id, lastname, firstname) VALUES  
(220844bf-4860-49d6-9a4b-6b5d3a79cbfb, 'TIRALONGO','Paolo');  
INSERT INTO cycling.cyclist_name (id, lastname, firstname) VALUES (6ab09bec-e68e-48d9-  
a5f8-97e6fb4c9b47, 'KRUIKSWIJK','Steven');  
INSERT INTO cycling.cyclist_name (id, lastname, firstname) VALUES (fb372533-  
eb95-4bb4-8685-6ef61e994caa, 'MATTHEWS', 'Michael');
```

Step 2 Exports data from and imports data to the **cyclist_name** table.

1. Export the **id** and **lastname** columns from the **cyclist_name** table to a CSV file.

```
COPY cycling.cyclist_name (id,lastname) TO './cyclist_lastname.csv' WITH HEADER =  
TRUE;
```

Figure 3-30 Exported successfully

```

Using 15 child processes

Starting copy of cycling.cyclist_name with columns [id, lastname].
Processed: 6 rows; Rate:      41 rows/s; Avg. rate:      41 rows/s
6 rows exported in 0.201 seconds.
Processed: 6 rows; Rate:      20 rows/s; Avg. rate:      40 rows/s
Results
Results          : success
Total operation  : 6
Total operation time : 0.201 seconds
Operation rate    : 40.468307430069224 rows/s
Total ranges     : 25
Success ranges   : 25
Failed ranges    : 0
Num processes    : 15
Max attempts     : 5

Ranges Results:
ranges          result      exported rows
(-7, 683212743470724096) success      1
(3074457345618258593, 3757670089088982528) success      1
(2220441416279853312, 3074457345618258593) success      1
(-6148914691236517207, -5465701947765792768) success      1
(-854015929338405120, -7) success      1
(-2391244602147534336, -1537228672809129307) success      1

```

After the preceding command is executed successfully, the **cyclist_lastname.csv** file is created in the upper-level directory of the current directory. If the file already exists, it will be overwritten.

2. Export the **id** and **first name** columns from the **cyclist_name** table to another CSV file.

```
COPY cycling.cyclist_name (id,firstname) TO './cyclist_firstname.csv' WITH HEADER = TRUE;
```

Figure 3-31 Exported successfully

```
Using 15 child processes

Starting copy of cycling.cyclist_name with columns [id, firstname].
Processed: 6 rows; Rate:      67 rows/s; Avg. rate:      67 rows/s
6 rows exported in 0.134 seconds.
Processed: 6 rows; Rate:      33 rows/s; Avg. rate:      67 rows/s
Results                               : success
Total operation                        : 6
Total operation time                   : 0.134 seconds
Operation rate                         : 66.57325993275435 rows/s
Total ranges                           : 25
Success ranges                         : 25
Failed ranges                          : 0
Num processes                          : 15
Max attempts                           : 5

Ranges Results:
ranges                                result    exported rows
(-854015929338405120, -7)            success   1
(-7, 683212743470724096)            success   1
(3074457345618258593, 3757670089088982528) success   1
(-6148914691236517207, -5465701947765792768) success   1
(2220441416279853312, 3074457345618258593) success   1
(-2391244602147534336, -1537228672809129307) success   1
```

After the preceding command is executed successfully, the **cyclist_firstname.csv** file is created in the upper-level directory of the current directory. If the file already exists, it will be overwritten.

3. Delete data from the **cyclist_name** table. To ensure data security, the TRUNCATE command is not supported.
`DELETE FROM cycling.cyclist_name WHERE id = 'fb372533-eb95-4bb4-8685-6ef61e994caa';`
4. No data exists in the table.
`SELECT * FROM cycling.cyclist_name ;`

Figure 3-32 Querying data

```
cqlsh> SELECT * FROM cycling.cyclist_name ;

 id | firstname | lastname
----+-----+-----
```

5. Import the **cyclist_firstname.csv** file.
`COPY cycling.cyclist_name (id,firstname) FROM './cyclist_firstname.csv' WITH HEADER = TRUE;`

Figure 3-33 Import succeeded

```
cqlsh> COPY cycling.cyclist_name (id,firstname) FROM './cyclist_firstname.csv' WITH HEADER = TRUE ;
Using 15 child processes

Starting copy of cycling.cyclist_name with columns [id, firstname].
Processed: 6 rows; Rate:      11 rows/s; Avg. rate:      15 rows/s
6 rows imported from 1 files in 0.387 seconds (0 skipped).
```

- Verify the imported data.
`SELECT * FROM cycling.cyclist_name;`

Figure 3-34 Import succeeded

```
cqlsh> SELECT * FROM cycling.cyclist_name ;
```

id	firstname	lastname
e7ae5cf3-d358-4d99-b900-85902fda9bb0	Alex	null
fb372533-eb95-4bb4-8685-6ef61e994caa	Michael	null
5b6962dd-3f90-4c93-8f61-eabfa4a803e2	Marianne	null
220844bf-4860-49d6-9a4b-6b5d3a79cbfb	Paolo	null
6ab09bec-e68e-48d9-a5f8-97e6fb4c9b47	Steven	null
e7cd5752-bc0d-4157-a80f-7523add8dbcd	Anna	null

- Import the `cyclist_lastname.csv` file.
`COPY cycling.cyclist_name (id,lastname) FROM './cyclist_lastname.csv' WITH HEADER = TRUE;`

Figure 3-35 Importing data

```
Using 15 child processes

Starting copy of cycling.cyclist_name with columns [id, lastname].
Processed: 6 rows; Rate: 11 rows/s; Avg. rate: 16 rows/s
6 rows imported from 1 files in 0.378 seconds (0 skipped).
```

- Check whether the data is updated.
`SELECT * FROM cycling.cyclist_name;`

The query result is displayed,

Figure 3-36 Import succeeded

```
cqlsh> SELECT * FROM cycling.cyclist_name ;
```

id	firstname	lastname
e7ae5cf3-d358-4d99-b900-85902fda9bb0	Alex	FRAME
fb372533-eb95-4bb4-8685-6ef61e994caa	Michael	MATTHEWS
5b6962dd-3f90-4c93-8f61-eabfa4a803e2	Marianne	VOS
220844bf-4860-49d6-9a4b-6b5d3a79cbfb	Paolo	TIRALONGO
6ab09bec-e68e-48d9-a5f8-97e6fb4c9b47	Steven	KRUIKSWIJK
e7cd5752-bc0d-4157-a80f-7523add8dbcd	Anna	VAN DER BREGGEN

(6 rows)

----End

3.6 Intra-region DR

3.6.1 Creating a DR Instance

GeminiDB instances support HA. If an instance fails to be connected due to a natural disaster, you can switch services to its DR instance and change the DB connection address on the application side to quickly recover service access.

Precautions

- A primary instance can have only one DR instance.
- This function is in the open beta test (OBT) phase. To use this function, contact customer service.

Prerequisites

A primary instance has been created.

Creating a DR Instance

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, locate the primary instance you want to create a DR instance for and choose **More > Create DR Instance** in the **Operation** column.
- Step 4** On the displayed page, configure required parameters and click **Next**.

Table 3-8 Basic information

Parameter	Description
Billing Mode	<p>Select Yearly/Monthly or Pay-per-use.</p> <ul style="list-style-type: none"> • Yearly/Monthly <ul style="list-style-type: none"> – In this mode, specify Required Duration at the bottom of the page. The system deducts the fees incurred from your account based on the service price. – If you do not need such an instance any longer after it expires, change the billing mode to pay-per-use to optimize costs. For details, see Changing the Billing Mode from Yearly/Monthly to Pay-per-Use. <p>NOTE Yearly/Monthly instances cannot be deleted directly. If such an instance is no longer required, unsubscribe from it. For details, see Unsubscribing from a Yearly/Monthly Instance.</p> <ul style="list-style-type: none"> • Pay-per-use <ul style="list-style-type: none"> – If you select this billing mode, you are billed based on how much time the instance is in use. – If you expect to use an instance for a long period of time, change its billing mode to yearly/monthly to optimize costs. For details, see Changing the Billing Mode from Pay-per-Use to Yearly/Monthly.

Parameter	Description
Region	The region is the same as that of the primary instance.
DB Instance Name	The instance name: Can include 4 to 64 bytes and must start with a letter. It is case-sensitive and allows only letters, digits, hyphens (-), and underscores (_).
Compatible API	Cassandra
DB Instance Type	Cluster
DB Engine Version	The compatible API version is the same as that of the primary instance.
CPU Type	The CPU type is the same as that of the primary instance.
AZ	<p>Availability zone where the instance is created. An AZ is a part of a region with its own independent power supplies and networks. AZs are physically isolated but can communicate through an internal network.</p> <p>An instance can be deployed in one or three AZs.</p> <ul style="list-style-type: none"> • If you want to deploy an instance in a single AZ, select one AZ. • If you want to deploy your instance across AZs for disaster recovery, select three AZs. Nodes of the instance are evenly distributed across the three AZs.

Table 3-9 Specifications and storage

Parameter	Description
Instance Specifications	<p>vCPUs and memory of the instance.</p> <p>Different performance specifications support different connections and maximum IOPS. Select CPUs and memory based on your service requirements. For details, see DB Instance Specifications and Performance.</p> <p>After an instance is created, you can change its vCPUs and memory by referring to Changing vCPUs and Memory of an Instance.</p>
Nodes	<p>Specify the number of nodes based on service requirements.</p> <p>After an instance is created, you can add nodes by referring to Adding Nodes.</p>

Parameter	Description
Storage Space	<p>Storage space depends on the instance specifications. The minimum storage space is 100 GB, and the storage space you set must be an integer. You can increase a minimum of 1 GB at a time.</p> <p>Enable autoscaling to ensure that the instance has sufficient storage and keeps available. To enable this function, just switch on button Configure Autoscaling and set the following parameters:</p> <ul style="list-style-type: none"> • If available storage drops to or below: The storage threshold for triggering autoscaling. When the percentage of available storage drops to or below the threshold you set or 10 GB, the system automatically scales up your instance storage. • Increase by: The percentage that your instance storage will be scaled up at. If the increased storage is not a multiple of 10 GB, the system will round it up to the nearest multiple of 10 GB. At least 100 GB is added each time. • Autoscaling Limit: Maximum amount that the system can automatically scale up an instance's storage space to. The value must be no less than the total storage of the instance and cannot exceed its maximum storage. <p>After an instance is created, you can scale up its storage space by referring to Scaling Up Storage Space.</p> <p>NOTE</p> <ul style="list-style-type: none"> • Once autoscaling is enabled, an agency will be created and fees will be automatically deducted from your account. • Autoscaling is available only when you have the required permission. To enable this function, contact customer service. • You can enable autoscaling after an instance is created. For details, see Configuring Autoscaling.

Table 3-10 Network

Parameter	Description
VPC	The VPC of the DR instance remains unchanged by default.
Subnet	The subnet of the DR instance remains unchanged by default. If you select another subnet in the same VPC, ensure that the selected subnet can be connected to the subnet of the primary instance.
Security Group	The security group of the DR instance remains unchanged by default.

Table 3-11 Database configuration

Parameter	Description
Administrator	Username of the administrator account. The default value is rwuser .
Administrator Password	The password must be the same as that of the primary instance to ensure that a switchover is performed in the event of a failure.
Confirm Password	Enter the administrator password again.
Parameter Template	A parameter template contains API configuration values that can be applied to one or more instances. After an instance is created, you can modify its parameters to better meet your service requirements. For details, see Modifying a Parameter Template .

Table 3-12 Tags

Parameter	Description
Tags	<p>The setting is optional. Adding tags helps you better identify and manage your instances. Each instance supports up to 20 tags. To use more tags, contact customer service to apply for a quota of 20 tags.</p> <p>A tag consists of a tag key and a tag value.</p> <ul style="list-style-type: none"> • Tag key: Mandatory if the instance is going to be tagged. Each tag key is unique for each instance. It can include up to 36 characters, including digits, letters, underscores (_), and hyphens (-). • Tag value: Optional if the instance is going to be tagged. The value can contain up to 43 characters, including digits, letters, underscores (_), periods (.), and hyphens (-). <p>After an instance is created, you can view its tags on the Tags tab and can also add, modify, and delete tags of your instance. For details, see Managing Tags.</p>


Table 3-13 Required duration


Parameter	Description
Required duration	The length of your subscription if you select Yearly/Monthly billing. Subscription lengths range from one month to three years.

Parameter	Description
Auto-renew	<ul style="list-style-type: none"> This option is not selected by default. If you select this option, the auto-renew cycle is determined by the selected required duration.

Step 5 On the displayed page, confirm the instance details.

- For yearly/monthly instances
 - If you need to modify the settings, click **Previous** to modify parameters.
 - If you do not need to modify the settings, read and agree to the service agreement, click **Pay Now**, and complete the payment.
- For pay-per-use instances
 - If you need to modify the settings, click **Previous** to modify parameters.
 - If no modification is required, read and agree to the service agreement and click **Submit**.

Step 6 On the **Instances** page, click  in front of the primary instance to view and manage the DR instance.

- During DR instance creation, the status of the primary instance is **DR cluster being created**, and the status of the DR instance is **Creating**. This process takes about 5 to 9 minutes.
 - After the creation is complete, the status changes to **Available**.
- You can click  in the upper right corner of the page to refresh the instance status.
- During creation, an automated backup policy is enabled by default. A full backup is automatically triggered after an instance is created.

----End

3.6.2 Deleting the DR Relationship

You can delete the primary or DR instance to delete the DR relationship.

Precautions

- When you delete an instance, all the data in it and all its automated backups are automatically deleted as well and cannot be restored.
- After you delete an instance, all nodes in the instance are also deleted.
- To delete a yearly/monthly instance, you need to unsubscribe from the order. For details, see [Unsubscribing from a Yearly/Monthly Instance](#).
- If you enable operation protection to improve the security of your account and cloud products, two-factor authentication is required for sensitive operations. For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

Procedure

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, locate the primary or DR instance that you want to delete and choose **More > Delete** in the **Operation** column.
- Step 4** If you have enabled operation protection, click **Start Verification** in the **Delete DB Instance** dialog box. On the displayed page, click **Send Code**, enter the verification code, and click **Verify**. The page is closed automatically.
- Step 5** In the displayed dialog box, click **Yes**.

When the instance is being deleted, its status is **DR relationship being canceled**. After the instance is deleted, it is not displayed in the instance list.

----End

3.7 Cross-region Dual-active DR

3.7.1 Overview

GeminiDB Cassandra supports cross-region dual-active DR and bidirectional synchronization between two instances at different sites. Once an instance becomes faulty, the other instance takes over read/write traffic to ensure service continuity.

Cross-region dual-active DR allows you to deploy two GeminiDB Cassandra instances in different data centers. Both of the two instances can handle service requests. If a data center becomes faulty, services in the faulty data center can be switched to the other data center to recover services without any interruption.

For how to configure cross-region dual-active DR, see [Creating a Dual-Active Relationship](#).

3.7.2 Creating a Dual-Active Relationship

If you already have an instance in one region, GeminiDB Cassandra allows you to create a dual-active DR instance in a different region, and synchronizes data between them.

This section describes how to create such a dual-active DR instance.

Precautions

- Before you create a dual-active relationship, create a destination instance in a specific region and ensure that the instance has the same or higher specifications than your existing instance. If you want to create an instance with smaller specifications than your existing instance, contact customer service and make sure that CPU and memory resources of the instance are abundant.
- Ensure that the standby instance has no additional tables before a dual-active DR relationship is created for it.

- The standby instance must have the same administrator password as the primary instance.
- If you want to create a dual-active relationship again after it is removed, you need to execute the DROP statement to clear tables in the standby instance. Otherwise, the dual-active relationship cannot be created.
- The existing instance functions as the primary instance and backs up data and transfers it to the destination instance.

Procedure

Step 1 In the service list, choose **Databases > GeminiDB**.


Step 2 On the **Instances** page, locate the instance that you want to create a dual-active relationship for and choose **More > Create Dual-Active Relationship** in the **Operation** column.

Step 3 On the **Create Dual-Active Relationship** dialog box, locate the destination instance as the dual-active DR instance.


NOTE

The destination instance must be in a different CIDR block from the source instance and has the same or higher specifications and no less nodes and storage space than the source, to synchronize data in real time between them and avoid subnet conflicts across regions.

Step 4 Click **OK**.

Step 5 On the **Instances** page, click  before the source instance and view and manage its DR instance.

- When the DR instance is being created, its status is **Creating dual-active relationship**.
- After the creation is complete, the status changes to **Available**.

You can click  in the upper right corner of the page to refresh the instance status.

----End

3.7.3 Deleting a Dual-active Relationship

This section describes how to delete a dual-active relationship on the GeminiDB Cassandra console.

Procedure

Step 1 In the service list, choose **Databases > GeminiDB**.

Step 2 On the **Instances** page, locate the instance that you want to delete a dual-active relationship from and choose **More > Delete Dual-Active Relationship** in the **Operation** column.

Step 3 In the displayed dialog box, click **Yes**.

When the instance is being deleted, its status is **Deleting dual-active relationship**. After the relationship is deleted, the instance status changes to **Available**.

----End

3.8 Data Backup

3.8.1 Overview

GeminiDB Cassandra API supports instance backups and restorations to ensure data reliability. After an instance is deleted, the manual backup data is retained. Automated backup data is released together with instances. Backup data cannot be downloaded or exported.

Backup Methods

Both automatic backup and manual backup are supported.

- Automated backup

You can click [Modify Backup Policy](#) on the GeminiDB console, and the system will automatically back up your instance data based on the time window and backup cycle you set in the backup policy and will store the data for the retention period you specified.

Automated backups cannot be manually deleted. You can adjust their retention period by referring to [Modifying an Automated Backup Policy](#), and backups that expire will be automatically deleted.

- Manual backup

A manual backup is a full backup of a DB instance and can be retained until you manually delete it. Manual backup can be triggered at any time to meet your service requirements.

Regularly backing up your database is recommended. If your database becomes faulty or data is corrupted, you can restore it from backups.

Table 3-14 Backup methods

Method	Scenario
Automated backup	After you set a backup policy, the system automatically backs up your database based on the policy. You can also modify the policy based on service requirements.
Manual backup	You can enable full backup for your instance based on service requirements.
Cross-region backup	GeminiDB Cassandra API allows you to store backups in the destination region. Then for disaster recovery, you can restore the backups to a new instance in another region.

Backup Storage

Backups are stored in OBS buckets, providing disaster recovery and saving space.

After you purchase an instance, GeminiDB Cassandra API will provide additional backup storage of the same size as you purchased. For example, if you purchase an instance of 100 GB, you will obtain additional backup storage of 100 GB free of charge. If the size of backup data does not exceed 100 GB, the backup data is stored on OBS free of charge. If the size of the backup data exceeds 100 GB, you will be charged based on the OBS billing rules.

3.8.2 Managing Automated Backups

Automated backups can be created to ensure data reliability. If a database or table is deleted, maliciously or accidentally, backups can help restore your data.

Precautions

- Backup files are saved as packages in OBS buckets. Upload of backup files and service reads both consume bandwidth, so the upload bandwidth of OBS is limited. The bandwidth of a single node ranges from 20 MB/s to 70 MB/s. For better performance, you need to specify appropriate nodes for an instance and take into account the bandwidth for uploading backups.
- The CPU usage may increase 5% to 15% because uploading backups consumes CPU resources.
- The memory usage may increase by about 300 MB during the upload of backups. The increase depends on the instance's data volume. The increased memory mainly caches data during backup upload and service read. After the backup upload is complete, the memory recovers.

Automated Backup Policy

Automated backups are generated according to a backup policy and saved as packages in OBS buckets to ensure data confidentiality and durability. You are advised to regularly back up your database, in case it becomes faulty or damaged. However, backing up data might affect the database read and write performance so it is recommended that you enable automated backups during off-peak hours.

When you create an instance, an automated backup policy is enabled by default.

Figure 3-37 Enabling the automated backup policy

×

Modify Backup Policy

Automated Backup

Retention Period days
Enter an integer from 1 to 35.

Time Zone GMT+08:00

Time Window

Backup Cycle

All

Monday Tuesday Wednesday

Thursday Friday Saturday

Sunday

A minimum of one day must be selected.

Incremental Backup Interval minutes

- **Incremental Backup Interval:** Incremental backups are generated every 15 minutes. To enable automated backup, contact technical support.
- **Retention Period:** Automated backup files are saved for seven days by default. The backup retention period can range from 1 to 35 days. Full backups are retained till the retention period expires. However, even if the retention period has expired, the most recent backup will be retained.
 - Extending the retention period improves data reliability. You can extend the retention period as needed.
 - If you shorten the retention period, the new backup policy takes effect for existing backups. Any automated backups (including full and incremental backups) that have expired will be automatically deleted. Manual backups will not be automatically deleted but you can delete them manually.

 NOTE

- If the retention period is less than seven days, the system automatically backs up data every day.
- The system checks existing automated backup files and deletes the files that exceed the backup retention period you set.
- **Time Window:** A one-hour period the backup will be scheduled within 24 hours, such as 04:00–05:00. The backup time is in GMT format. If the DST or standard time is switched, the backup time segment changes with the time zone.

If **Retention Period** is set to **2**, full and incremental backups that have been stored for more than two days will be automatically deleted. That is, the backup generated on Monday will be deleted on Wednesday. Similarly, the backup generated on Tuesday will be deleted on Thursday.

Policy for automatically deleting full backups:

To ensure data integrity, even after the retention period expires, the most recent backup will be retained, for example,

If **Backup Cycle** was set to **Monday** and **Tuesday** and the **Retention Period** was set to **2**:

- The full backup generated on Monday will be automatically deleted on Thursday. The reasons are as follows:

The full backup generated on Monday expires on Wednesday, but it is the last backup, so it will be retained until a new backup expires. The next backup will be generated on Tuesday and will expire on Thursday. So the full backup generated on Monday will not be automatically deleted until Thursday.

- The full backup generated on Tuesday will be automatically deleted on the following Wednesday. The reasons are as follows:

The backup generated on Tuesday will expire on Thursday, but as it is the last backup, so it will be retained until a new backup expires. The next backup will be generated on the following Monday and will expire on the following Wednesday. So the full backup generated on Tuesday will not be automatically deleted until the following Wednesday.

- **Backup Cycle:** By default, each day of the week is selected.
 - **All:** Each day of the week is selected. The system automatically backs up data every day.
 - **Select a cycle:** You can select one or more days in a week. The system automatically backs up data at the specified time.

 NOTE

A full backup starts within one hour of the time you specify. The amount of time required for the backup depends on the amount of data to be backed up. The more data has to be backed up, the longer it will take.

- After the DB instance is created, you can modify the automated backup policy as needed. You can change the time window after the DB instance is created. The system backs up data based on the automated backup policy you have set.
- After the automated backup policy is disabled, any automated backups in progress stop immediately.

Modifying an Automated Backup Policy

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, click the instance whose backup policy you want to modify and click its name.
- Step 4** Choose **Backups & Restorations** in the navigation pane on the left, and click **Modify Backup Policy**. In the displayed dialog box, set the backup policy and click **OK**.

For details, see [Automated Backup Policy](#).

Figure 3-38 Modifying the backup policy

Modify Backup Policy ×

Automated Backup

Retention Period days
Enter an integer from 1 to 35.

Time Zone GMT+08:00

Time Window

Backup Cycle

All

Monday Tuesday Wednesday

Thursday Friday Saturday

Sunday

A minimum of one day must be selected.

Incremental Backup Interval minutes

OK Cancel

- Step 5** Check or manage the generated backups on the **Backups** or **Backups & Restorations** page.

----End

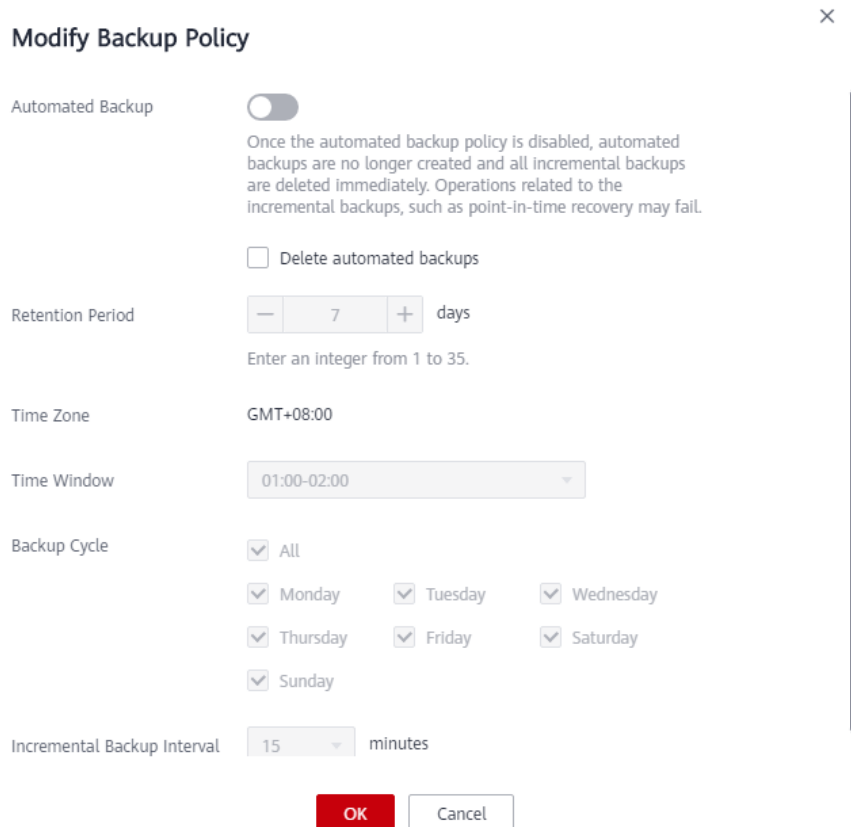
Disabling an Automated Backup Policy

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, locate the instance that you want to disable automated backup for and click its name.

Step 4 Choose **Backups & Restorations** in the navigation pane on the left, and click **Modify Backup Policy**.

Step 5 In the displayed dialog box, click  to disable automated backup and click **OK**.

Figure 3-39 Disabling automated backup



Modify Backup Policy ×

Automated Backup
 Once the automated backup policy is disabled, automated backups are no longer created and all incremental backups are deleted immediately. Operations related to the incremental backups, such as point-in-time recovery may fail.

Delete automated backups

Retention Period days
 Enter an integer from 1 to 35.

Time Zone GMT+08:00

Time Window

Backup Cycle All
 Monday Tuesday Wednesday
 Thursday Friday Saturday
 Sunday

Incremental Backup Interval minutes

When you disable automated backup, specify whether to delete the automated backups:

- If you select **Delete automated backups**, all backup files within the retention period will be deleted. There are no automated backups displayed until you enable automated backup again.
- If you do not select **Delete automated backups**, backup files within the retention period will be retained, but you can still manually delete them later if needed. For details, see [Deleting an Automated Backup](#).

If automated backup is disabled, any automated backups in progress stop immediately.

----End

Deleting an Automated Backup

If automated backup is disabled, you can delete stored automated backups to free up storage space.

If automated backup is enabled, the system will delete automated backups when they expire. You cannot delete them manually.

NOTICE

Deleted backups cannot be recovered. Exercise caution when performing this operation.

- **Method 1**
 - a. Log in to the management console.
 - b. In the service list, choose **Databases > GeminiDB**.
 - c. On the **Instances** page, click the instance whose automated backups you want to delete and click its name.
 - d. Choose **Backups & Restorations** in the navigation pane on the left, locate the backup you want to delete and click **Delete** in the **Operation** column.
 - e. In the displayed dialog box, confirm the backup details and click **Yes**.
- **Method 2**
 - a. Log in to the management console.
 - b. In the service list, choose **Databases > GeminiDB**.
 - c. On the **Backups** page, locate the backup that you want to delete and click **Delete** in the **Operation** column.
 - d. In the displayed dialog box, confirm the backup details and click **Yes**.

3.8.3 Setting a Cross-Region Backup Policy

GeminiDB Cassandra allows you to store backups in the destination region or OBS buckets. Then for disaster recovery, you can restore the backups to a new instance in another region.

After a cross-region backup policy is set for an instance, the system will synchronize backups of the instance to the destination region you specified. You can manage cross-region backup files on the **Backups** page.

Precautions

- To apply for the permissions to set cross-region backup policies, contact customer service.
- Before you configure a cross-region backup policy, make sure to enable automated backup first. Otherwise, the cross-region backup policy cannot take effect. For details, see [Modifying an Automated Backup Policy](#).

Setting or Modifying a Cross-Region Backup Policy

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases > GeminiDB**.

Step 3 On the **Instances** page, locate the instance that you want to connect to and click its name.

Step 4 In the navigation pane on the left, choose **Backups & Restorations**.

Step 5 On the displayed page, click **Set Cross-Region Backup Policy**.

Step 6 In the displayed dialog box, set required parameters.

Figure 3-40 Setting a cross-region backup policy

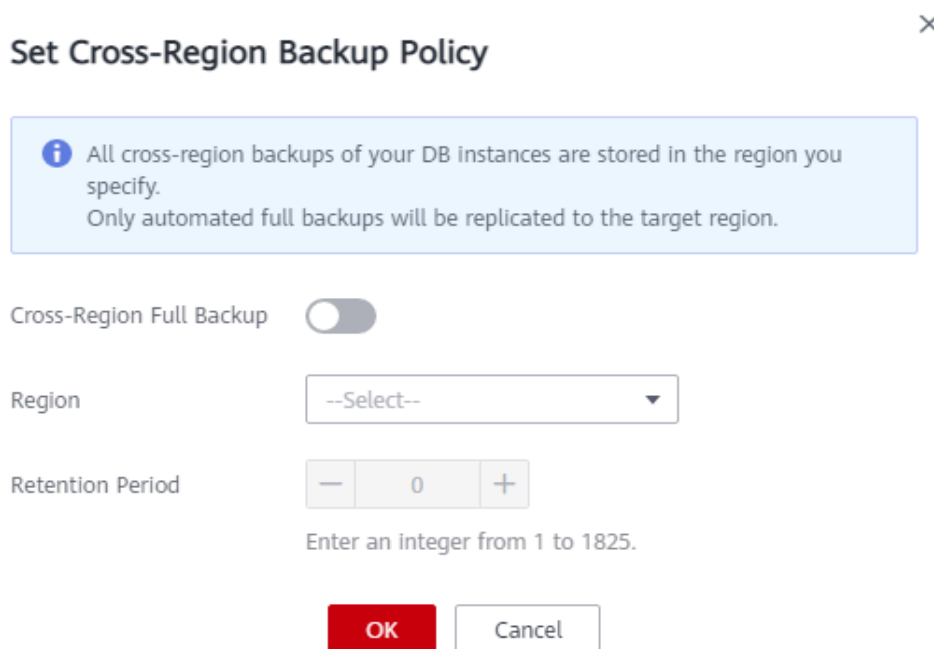


Table 3-15 Parameter description

Parameter	Description
Cross-Region Full Backup	If you enable Cross-Region Full Backup , automated full backup files of the instance will be stored in the region you specify.
Region	You can select the region for storing backups based on service requirements.
Retention Period	Number of days that cross-region backups are kept. The value ranges from 1 to 1825 . You can increase the retention period to improve data reliability.

NOTE

- Only new backups generated after you set a cross-region backup policy will be stored in the region you specify.
- All cross-region backups of your DB instances are stored in the same region you specify.
- Cross-region backups are synchronized to the destination region you specify only after your instance is backed up locally.
- Only automated full backups are replicated to the destination region.

Step 5 On the displayed page, click **Set Cross-Region Backup Policy**.

Step 6 In the displayed dialog box, disable **Cross-Region Full Backup**.

Figure 3-43 Disabling cross-region backup

NOTE

- After cross-region backup is disabled, the cross-region backup task is stopped and all cross-region backups are deleted immediately. As a result, operations using cross-region backups will fail.
- If an instance with cross-region backup enabled is deleted, its cross-region backups will be retained. The retention period depends on settings of the cross-region backup policy.

Step 7 Click **OK**.

----End

3.8.4 Managing Manual Backups

To ensure data reliability, GeminiDB Cassandra API allows you to manually back up instances whose status is **Available**. If a database or table is deleted, maliciously or accidentally, backups can help recover your data.

Precautions

- By default, you can create up to 50 backups.
- Manual backups are full backups.
- Backup files are saved as packages in OBS buckets. Upload of backup files and service reads both consume bandwidth, so the upload bandwidth of OBS is limited. The bandwidth of a single node ranges from 20 MB/s to 70 MB/s.

For better performance, you need to specify appropriate nodes for an instance and take into account the bandwidth for uploading backups.

- The CPU usage may increase 5% to 15% because uploading backups consumes CPU resources.
- The memory usage may increase by about 300 MB during the upload of backups. The increase depends on the instance's data volume. The increased memory mainly caches data during backup upload and service read. After the backup upload is complete, the memory recovers.

Creating a Manual Backup

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases** > **GeminiDB**.

Step 3 Create a manual backup.

Method 1

On the **Instances** page, locate the instance that you want to create a backup for and choose **More** > **Create Backup** in the **Operation** column.

Method 2

1. On the **Instances** page, click the instance that you want to create a backup for and click its name.
2. Choose **Backups & Restorations** in the navigation pane on the left, click **Create Backup**.

Method 3

In the navigation pane on the left, choose **Backups** and click **Create Backup**.

Step 4 In the displayed dialog box, enter a backup name and description and click **OK**.

Figure 3-44 Creating a manual backup

Create Backup ×

DB Instance Name nosql-6693

* Backup Name ?

Description ?

/256

OK Cancel

Table 3-16 Parameter description

Parameter	Description
DB Instance Name	Must be the name of the DB instance to be backed up and cannot be modified.
Backup Name	Must be 4 to 64 characters in length and start with a letter. It is case-insensitive and contains only letters, digits, hyphens (-), and underscores (_).
Description	Contains a maximum of 256 characters and cannot include line breaks or special characters > <"&'=

Step 5 View the backup status.

- When the backup is being created, query the backup status on the **Backups** or **Backups & Restorations** page. The backup status is **Backing up**.
- After the backup is created, the backup status is **Completed**.

----End

Deleting a Manual Backup

If you do not need a manual backup any longer, delete it on the **Backups** or **Backups & Restorations** page.

Deleted backups are not displayed in the backup list.

NOTICE

Deleted backups cannot be recovered. Exercise caution when performing this operation.

Method 1

1. Log in to the management console.
2. In the service list, choose **Databases > GeminiDB**.
3. On the **Instances** page, locate the instance whose backup you want to delete and click its name.
4. Choose **Backups & Restorations** in the navigation pane on the left, locate the backup you want to delete and click **Delete** in the **Operation** column.
5. In the displayed dialog box, confirm the backup details and click **Yes**.

Method 2

1. Log in to the management console.
2. In the service list, choose **Databases > GeminiDB**.
3. On the **Backups** page, locate the backup you want to delete and click **Delete** in the **Operation** column.
4. In the displayed dialog box, confirm the backup details and click **Yes**.

3.9 Data Restoration

3.9.1 Restoration Methods

GeminiDB Cassandra API supports multiple forms of data restoration. You can select one based on service requirements.

Table 3-17 Restoration methods

Method	Scenario
Restoring Data to a New Instance	You can restore an existing backup file to a new instance.
Restoring a Backup to a Specific Point in Time	You can use an automated backup to restore an instance to a specified point in time.

3.9.2 Restoring Data to a New Instance

GeminiDB Cassandra API allows you to use an existing backup to restore data to a new instance.

Precautions

- The new instances must have at least as many nodes as the original instance.
- The new instance must have at least as much storage as the original instance.
- Incremental backup and PITR are not supported.
- Restoration to the current instance is not supported.
- You can scale in the memory, but the memory decrease cannot become less than the actual memory used during the backup.
- The restored instance uses the same parameter group as the original instance.
- During the instance restoration, backups are downloaded from OBS buckets to the data directory of the restored instance. The download bandwidth of OBS is 40 MB/s.

Procedure

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** Restore an instance from backup.

Method 1

1. On the **Instances** page, locate the instance whose backup you want to restore and click its name.

2. Choose **Backups & Restorations** in the navigation pane on the left, locate the backup that you want to restore and click **Restore** in the **Operation** column.

Figure 3-45 Backups and restorations

Backup Name	DB Instance Name/ID	DB Engine Version	Backup Type	Backup Time	Status	Size	Description	Operation
cassandra-nosql-6693-202...	nosql-6693-63a20b6942c4f3d92d5176-449f...	cassandra 3.11	Automated	Jul 05, 2020 04:30:40 - Jul ...	Completed	2.42 MB	--	Restore
cassandra-nosql-6693-202...	nosql-6693-63a20b6942c4f3d92d5176-449f...	cassandra 3.11	Automated	Jul 05, 2020 04:30:41 - Jul ...	Completed	2.44 MB	--	Restore
cassandra-nosql-6693-202...	nosql-6693-63a20b6942c4f3d92d5176-449f...	cassandra 3.11	Automated	Jul 04, 2020 04:30:41 - Jul ...	Completed	2.42 MB	--	Restore
cassandra-nosql-6693-202...	nosql-6693-63a20b6942c4f3d92d5176-449f...	cassandra 3.11	Automated	Jul 03, 2020 04:30:40 - Jul ...	Completed	2.44 MB	--	Restore
cassandra-nosql-6693-202...	nosql-6693-63a20b6942c4f3d92d5176-449f...	cassandra 3.11	Automated	Jul 02, 2020 17:07:38 - Jul ...	Completed	2.44 MB	--	Restore

Method 2

On the **Backups** page, locate the backup that you want to restore and click **Restore** in the **Operation** column.

Figure 3-46 Backup management

Backup Name	DB Instance Name/ID	DB Engine Version	Backup Type	Backup Time	Status	Size	Description	Operation
cassandra-nosql-6693-202...	nosql-6693-63a20b6942c4f3d92d5176-449f...	cassandra 3.11	Automated	Jul 05, 2020 04:30:40 - Jul ...	Completed	2.42 MB	--	Restore
cassandra-nosql-6693-202...	nosql-6693-63a20b6942c4f3d92d5176-449f...	cassandra 3.11	Automated	Jul 05, 2020 04:30:41 - Jul ...	Completed	2.44 MB	--	Restore
cassandra-nosql-6693-202...	nosql-6693-63a20b6942c4f3d92d5176-449f...	cassandra 3.11	Automated	Jul 04, 2020 04:30:41 - Jul ...	Completed	2.42 MB	--	Restore
cassandra-nosql-6693-202...	nosql-6693-63a20b6942c4f3d92d5176-449f...	cassandra 3.11	Automated	Jul 03, 2020 04:30:40 - Jul ...	Completed	2.44 MB	--	Restore
cassandra-nosql-6693-202...	nosql-6693-63a20b6942c4f3d92d5176-449f...	cassandra 3.11	Automated	Jul 02, 2020 17:07:38 - Jul ...	Completed	2.44 MB	--	Restore
Cassandra_backup-159073...	autotest3-80864a20274645d80ac5f837c4...	cassandra 3.11	Manual	May 28, 2020 15:31:09 - ...	Completed	9.74 GB	strutt_testing	Restore Delete
backuptrp1587451431770...	autotest810c6f4e1c8844c2b8035c96c705d0...	cassandra 3.11	Manual	Apr 21, 2020 14:43:01 - A...	Completed	24.81 MB	--	Restore Delete
bak-nosql2restore	cas-restore2NewInst-fae8f2bc4205450985de283402b2c...	cassandra 3.11	Manual	Sep 09, 2019 14:34:49 - Se...	Completed	13.31 MB	--	Restore Delete

Step 4 In the displayed dialog box, confirm the current instance details and restoration method and click **OK**.

Figure 3-47 Restoring data to a new instance

Restore DB Instance ✕

DB Instance	Backup Name	DB Instance Name
	Cassandra_backup-1590737546429372500	autotest3

Restoration Method Create New Instance

OK
Cancel

- The default API type and DB engine version are the same as those of the original instance and cannot be changed.
- GeminiDB automatically calculates the minimum storage space required for restoration based on the size of the selected backup file. The storage capacity depends on the instance specifications, and must be an integer.
- The administrator password needs to be reset.
- To modify other parameters, see the description of buying instances of other DB engines in *Getting Started*.

Step 5 View the results.

A new instance is created using the backup data. The status of the new instance changes from **Creating** to **Available**.

After the restoration, the system will perform a full backup.

The new instance is independent from the original one.

----End

3.9.3 Restoring a Backup to a Specific Point in Time

You can restore an existing automated backup to a specific point in time.

The most recent full backup will be downloaded from OBS for restoration. After the restoration is complete, incremental backups will be replayed to the specified point in time. The time required depends on the amount of data to be restored.

Precautions

- GeminiDB Cassandra instances allow you to restore data to a new instance at a specific point in time.
- After automated backup is enabled, the system performs an incremental backup based on the preset incremental backup interval. The incremental backup is stored in OBS.
- Data can be restored to a specified time point only after the automated backup policy is enabled.
- During the instance restoration, backup files are downloaded from OBS buckets to the data directory of the restored instance. The download bandwidth of OBS is 40 MB/s.

Procedure

Step 1 Log in to the management console.

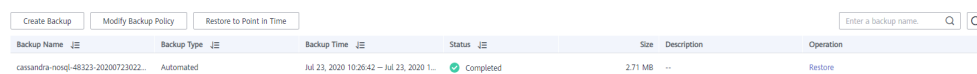
Step 2 In the service list, choose **Databases > GeminiDB**.

Step 3 On the **Instances** page, click the instance that you want to perform a PITR for.

Step 4 In the navigation pane on the left, choose **Backups & Restorations**.

Step 5 On the **Backups & Restorations** page, click **Restore to Point in Time**.

Figure 3-48 Restoring data to a point in time



Backup Name	Backup Type	Backup Time	Status	Size	Description	Operation
cassandra-moqj-48323-20200723022...	Automated	Jul 23, 2020 10:26:42 - Jul 23, 2020 1...	Completed	2.71 MB	--	Restore

Step 6 Select a restoration date and a time point to which data is restored, and then click **Yes**.

Figure 3-49 Restore to Point in Time

Step 7 On the **Create New Instance** page, create an instance of the same specifications as the instance to be restored. The new instance is independent from the original one.

- The new instance should be deployed in a different AZ to ensure that your applications will not be affected by SPOFs.
- The compatible API, instance type, instance version, and CPU type are the same as those of the original and cannot be changed.
- Other settings are the same as those of the original instance by default but can be modified based on service requirements. For details, see [Buying an Instance](#).

----End

3.10 Parameter Template Management

3.10.1 Creating a Parameter Template

You can use database parameter templates to manage DB API configurations. A database parameter template acts as a container for API configuration values that can be applied to one or more DB instances.

Each user can create up to 100 parameter templates. All types of instances in the same project can share the quota.

Procedure

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** In the navigation pane on the left, choose **Parameter Templates**.
- Step 4** On the **Parameter Templates** page, click **Create Parameter Template**.
- Step 5** Select a compatible API, specify a DB engine version and a parameter group description, and click **OK**.

Figure 3-50 Creating a parameter template

Create Parameter Template

* Compatible API

* DB Instance Type

* DB Engine Version

* Parameter Template Name

Description

0/256

You can create 98 more parameter templates. The parameter template quota is shared by all DB instances in a project.

- **Compatible API:** Select the API type and instance type that are compatible with your DB API parameter template.
- **DB Engine Version:** Select a DB engine version, for example, 3.11.
- **Parameter Template Name:** The template name can include 1 to 64 characters. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), underscores (_), and periods (.).
- **Description:** The description contains a maximum of 256 characters and cannot include line breaks or the following special characters >!<"&'=

Step 6 On the **Parameter Templates** page, view the created parameter template.

----End

3.10.2 Modifying a Parameter Template

You can modify parameters in a custom parameter template so that your instance can deliver spectacular performance.

Precautions

- Note that parameter values in default parameter templates cannot be changed.
- Though parameter values in a default template cannot be changed, you can view details about a default parameter template.
- If a custom parameter template is set incorrectly, the database startup may fail. You can re-configure the custom parameter template according to the configurations of the default parameter template.

Modifying Parameters in a Custom Parameter Template

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** In the navigation pane on the left, choose **Parameter Templates**.
- Step 4** Click the **Custom Templates** tab, locate the parameter template you want to modify, and click its name.
- Step 5** Change parameter values as required.

Figure 3-51 Modifying parameters in a parameter template

Parameter Name	Effective upon Restart	Value	Allowed Values	Description
slow_query_log_timeout_in_ms	Yes	300	1-60,000	Slow query time threshold, in milliseconds.
tombstone_failure_threshold	Yes	100000	10-10,000,000	If the number of tombstones in a query result exceed...
tombstone_warn_threshold	Yes	1000	10-100,000	When the number of tombstones in a query result ex...
unlogged_batch_across_partitions_warn_threshold	Yes	10	1-10,000	When the number of partitions for batch processing e...

- To save the modifications, click **Save**.
- To cancel the modifications, click **Cancel**.
- To preview the modifications, click **Preview**.

Figure 3-52 Preview Change

Preview Change ✕

Parameter Name	Current	New
slow_query_log_timeout...	500	300

Close

- Step 6** After parameters are modified, click **Change History** to view parameter modification details.

For details about how to view parameter modification details, see [Viewing Parameter Change History](#).

NOTICE

- The modifications take effect only after you apply the parameter template to DB instances. For details, see [Applying a Parameter Template](#).
- The change history page displays only the modifications of the last seven days.

----End

Modifying Parameters of a Specified Instance

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** In the navigation pane on the left, choose **Instances**. On the displayed page, locate the instance whose parameters you want to modify and click its name.
- Step 4** In the navigation pane on the left, choose **Parameters**. On the displayed page, modify parameters as required.

Figure 3-53 Parameters

Parameter Name	Effective upon Restart	Value	Allowed Values	Description
slow_query_log_timeout_in_ms	Yes	300	1-60,000	Slow query time threshold, in milliseconds.
tombstone_failure_threshold	Yes	100000	10-10,000,000	If the number of tombstones in a query result exce...
tombstone_warn_threshold	Yes	1000	10-100,000	When the number of tombstones in a query result ex...
unlogged_batch_across_partitions_warn_threshold	Yes	10	1-10,000	When the number of partitions for batch processing ...

- To save the modifications, click **Save**.
 - To cancel the modifications, click **Cancel**.
 - To preview the modifications, click **Preview**.
- Step 5** After parameters are modified, click **Change History** to view parameter modification details.

For details about how to view parameter modification details, see [Viewing Parameter Change History](#).

NOTICE

After you modify instance parameters, the modifications immediately take effect for the instance.

Check the value in the **Effective upon Restart** column.

- If the value is **Yes** and the DB instance status is **Pending restart**, restart the instance for the modifications to take effect.
- If the value is **No**, the modifications take effect immediately.

----End

3.10.3 Viewing Parameter Change History

Scenarios

You can view parameter change history of an instance or one of its custom parameter templates based on service requirements.

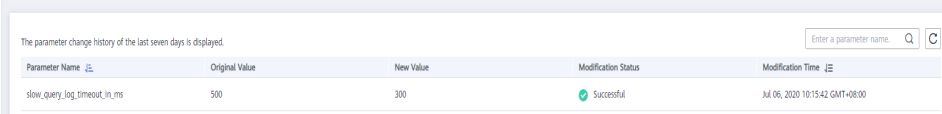
Precautions

In a newly exported or created parameter template, change history is left blank.

Viewing Change History of a Custom Parameter Template

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** In the navigation pane on the left, choose **Parameter Templates**. On the **Custom Templates** page, click the parameter template whose change history you want to view.
- Step 4** In the navigation pane on the left, choose **Change History**. Then, view the name, original value, new value, modification status, and modification time of the target parameter.

Figure 3-54 Viewing change history of a customer parameter template



Parameter Name	Original Value	New Value	Modification Status	Modification Time
slow_query_log_timeout_in_ms	500	300	Successful	Jul 06, 2020 10:15:42 GMT+08:00

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

-----End

Viewing Parameter Change History of an Instance

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, locate the instance whose parameter change history you want to view and click its name.
- Step 4** In the navigation pane on the left, choose **Parameters**. On the **Change History** page, view the name, original value, new value, modification status, and modification time of the target parameter.

Figure 3-55 Viewing parameter change history of an instance

Parameter Name	Original Value	New Value	Modification Status	Modification Time	Application Status	Application Time
slow_query_log_timeout_in_ms	500	300	Successful	Jul 06, 2020 10:17:06 GMT+08:00		--

----End

3.10.4 Exporting a Parameter Template

- You can export a parameter template of a DB instance for future use. To learn how to apply the exported parameter template to a DB instance, refer to section [Applying a Parameter Template](#).
- You can export the parameter template details (parameter names, values, and descriptions) of a DB instance to a CSV file for review and analysis.

Procedure

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases** > **GeminiDB**.
- Step 3** In the navigation pane on the left, choose **Instances**, locate the instance whose parameters you want to export, and click its name.
- Step 4** In the navigation pane on the left, choose **Parameters**. On the **Parameters** tab, above the parameter list, click **Export**.

Figure 3-56 Exporting a parameter template

- **Parameter Template:** You can export the parameters of the DB instance to a template for future use.

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

 **NOTE**

- **Parameter Template Name:** The template name can be 1 to 64 characters long. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), underscores (_), and periods (.).
- **Description:** The template description consists of a maximum of 256 characters and cannot include line breaks or the following special characters: >!<"&'=

After the parameter template is exported, a new template is generated in the list on the **Parameter Templates** page.

- **File:** You can export the parameter template details (parameter names, values, and descriptions) of a DB instance to a CSV file for review and analysis.

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

 **NOTE**

The file name must start with a letter and consist of 4 to 81 characters. It can contain only letters, digits, hyphens (-), and underscores (_).

----End

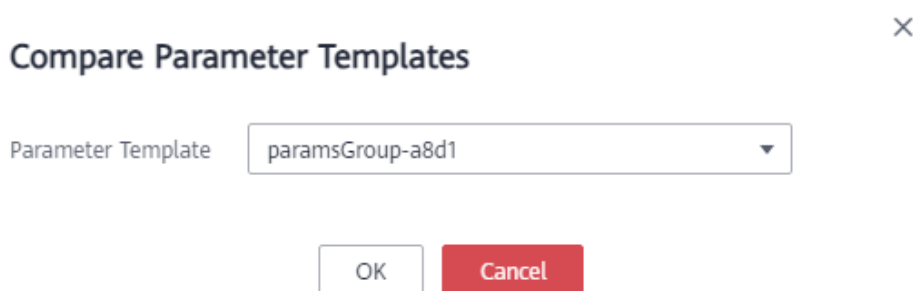
3.10.5 Comparing Parameter Templates

This section describes how to compare two parameter templates of the same instance type and compatible API to learn about their configurations.

Comparing Parameter Templates

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** In the navigation pane on the left, choose **Parameter Templates**.
- Step 4** In the parameter template list, locate the parameter template that you created and click **Compare** in the **Operation** column.
- Step 5** In the displayed dialog box, select a parameter template that is of the same instance type and compatible API as the selected template and click **OK**.

Figure 3-57 Comparing two parameter templates



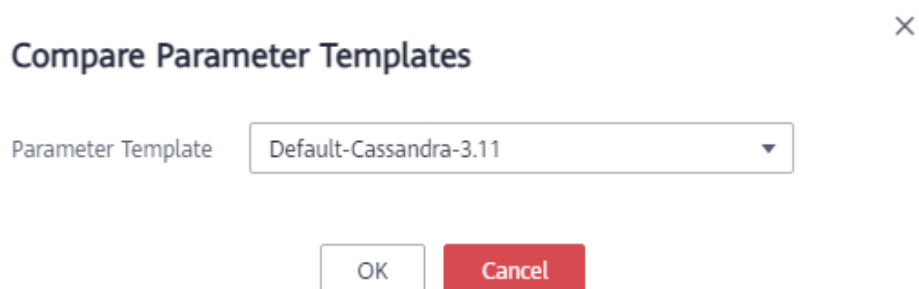
- If their parameters are different, the different parameter names and values are displayed.
- If their parameters are the same, no data is displayed.

----End

Comparing Parameter Templates of a Specified Instance

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** In the navigation pane on the left, choose **Instances**.
- Step 4** In the instance list, locate the instance whose parameter templates you want to compare and click its name.
- Step 5** In the navigation pane on the left, choose **Parameters** and then click **Compare** above the parameter list.
- Step 6** In the displayed dialog box, select a parameter template that is of the same instance type as the template of current instance and click **OK**.

Figure 3-58 Comparing the instance parameter template with another parameter template



- If their parameters are different, the different parameter names and values are displayed.
- If their parameters are the same, no data is displayed.

----End

3.10.6 Replicating a Parameter Template

You can replicate a parameter template you have created. When you have already created a parameter template and want to include most of the custom parameters and values from that template in a new parameter template, you can replicate that parameter template. You can also export a parameter template of a DB instance for future use.

Default parameter templates cannot be replicated, but you can create parameter templates based on the default templates provided.

Procedure

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** In the navigation pane on the left, choose **Parameter Templates**.
- Step 4** On the **Parameter Templates** page, click the **Custom Templates** tab. Locate the target parameter template and click **Replicate** in the **Operation** column.
- Alternatively, click the target DB instance on the **Instances** page. On the **Parameters** page, click **Export** to generate a new parameter template for future use.
- Step 5** In the displayed dialog box, enter the parameter template name and description and click **OK**.

Figure 3-59 Replicating a parameter template

Replicate Parameter Template ×

★ Source Parameter Template paramsGroup-22d1

★ New Parameter Template ?

Description ?

/256

You can replicate 97 more parameter templates. The parameter template quota is shared by all DB instances in a project.

- **New Parameter Template:** The template name can be up to 64 characters long. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), underscores (_), and periods (.).
- **Description:** The description contains a maximum of 256 characters and cannot include line breaks or the following special characters >!<"&'=

After the parameter template is replicated, a new template is generated in the list on the **Parameter Templates** page.

----End

3.10.7 Resetting a Parameter Template

You can reset all parameters in a custom parameter template to their default settings.

Procedure

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** In the navigation pane on the left, choose **Parameter Templates**.
- Step 4** On the **Parameter Templates** page, click the **Custom Templates** tab. Locate the target parameter template and choose **More > Reset** in the **Operation** column.
- Step 5** Click **Yes** to reset the parameter template.

----End

3.10.8 Applying a Parameter Template

GeminiDB Cassandra allows you to apply a parameter template. Modifications to parameters in a custom parameter template take effect only after you have applied the template to the target instance.

Procedure

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** In the navigation pane on the left, choose **Parameter Templates**.
- Step 4** On the **Parameter Templates** page, perform the following operations based on the template type:
 - To apply a default template, click **Default Templates**, locate the template, and in the **Operation** column, click **Apply**.
 - To apply a custom template, click **Custom Templates**, locate the template, and in the **Operation** column, choose **More > Apply**.

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

- Step 5** In the displayed dialog box, select one or more instances that the parameter template will be applied to and click **OK**.

After a parameter template is applied, you can [view its application records](#).

----End

3.10.9 Viewing Application Records of a Parameter Template

GeminiDB Cassandra allows you to view application records of a parameter template.

Procedure

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** In the navigation pane on the left, choose **Parameter Templates**.
- Step 4** On the **Parameter Templates** page, locate the parameter template whose application records you want to view and choose **More > View Application Record** in the **Operation** column.




You can view the name or ID of the instance that the parameter template applies to, as well as the application status, application time, and causes of any failures that have occurred.

----End

3.10.10 Modifying a Parameter Template Description

You can modify the description of a custom parameter template if needed.

Procedure

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** In the navigation pane on the left, choose **Parameter Templates**.
- Step 4** On the **Parameter Templates** page, click the **Custom Templates** tab. Locate the target parameter template and click  in the **Description** column.
- Step 5** Enter a new description. You can click  to submit or  to cancel the modification.
 - After you submit the modification, you can view the new description in the **Description** column on the **Parameter Templates** page.
 - The description can include up to 256 characters but cannot contain the following special characters: >!<"&'=

----End

3.10.11 Deleting a Parameter Template

You can delete a custom parameter template that is no longer in use.

Precautions

- Deleted templates cannot be recovered, so exercise caution when performing this operation.
- Default parameter templates cannot be deleted.

Procedure

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

- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** In the navigation pane on the left, choose **Parameter Templates**.
- Step 4** On the **Parameter Templates** page, click the **Custom Templates** tab. Locate the parameter template you want to delete and choose **More > Delete** in the **Operation** column.
- Step 5** Click **Yes** to delete the parameter template.
- End

3.11 Audit on Instance Operations

3.11.1 Key Operations Supported by CTS

With CTS, you can record GeminiDB Cassandra key operations for later query, audit, and backtracking.

Table 3-18 GeminiDB Cassandra key operations

Operation	Resource Type	Trace Name
Creating an instance	instance	NoSQLCreateInstance
Deleting an instance	instance	NoSQLDeleteInstance
Adding nodes	instance	NoSQLEnlargeInstance
Deleting nodes	instance	NoSQLReduceInstance
Restarting an instance	instance	NoSQLRestartInstance
Restoring data to a new instance	instance	NoSQLRestoreNewInstance
Scaling up storage space of an instance	instance	NoSQLExtendInstanceVolume
Resetting the password of an instance	instance	NoSQLResetPassword
Modifying the name of an instance	instance	NoSQLRenameInstance
Changing specifications	instance	NoSQLResizeInstance
Binding an EIP	instance	NoSQLBindEIP
Unbinding an EIP	instance	NoSQLUnBindEIP
Freezing an instance	instance	NoSQLFreezeInstance
Unfreezing an instance	instance	NoSQLUnfreezeInstance
Creating a backup	backup	NoSQLCreateBackup



Operation	Resource Type	Trace Name
Deleting a backup	backup	NoSQLDeleteBackup
Modifying the backup policy of an instance	backup	NoSQLSetBackupPolicy
Adding a tag for an instance	tag	NoSQLAddTags
Modifying an instance tag	tag	NoSQLModifyInstanceTag
Deleting an instance tag	tag	NoSQLDeleteInstanceTag
Creating a parameter template	parameterGroup	NoSQLCreateConfigurations
Modifying a parameter template	parameterGroup	NoSQLUpdateConfigurations
Modifying instance parameters	parameterGroup	NoSQLUpdateInstanceConfigurations
Replicating a parameter template	parameterGroup	NoSQLCopyConfigurations
Resetting a parameter template	parameterGroup	NoSQLResetConfigurations
Applying a parameter template	parameterGroup	NoSQLApplyConfigurations
Deleting a parameter template	parameterGroup	NoSQLDeleteConfigurations
Deleting the node that fails to be added	instance	NoSQLDeleteEnlargeFail-Node
Changing the security group of an instance	instance	NoSQLModifySecurityGroup
Configuring autoscaling	instance	NoSQLModifyAutoEnlarge-Policy
Exporting parameter template information for an instance	instance	NoSQLSaveConfigurations
Modifying the recycling policy	instance	NoSQLModifyRecyclePolicy

3.11.2 Querying Traces

After CTS is enabled, CTS starts recording operations on cloud resources. The CTS console stores the last seven days of operation records.

This section describes how to query the last seven days of operation records on the CTS console.

Procedure

- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Click **Service List**. Under **Management & Governance**, click **Cloud Trace Service**.
- Step 4** In the navigation pane on the left, click **Trace List**.
- Step 5** Specify filter criteria to search for the required traces. The following four filter criteria are available:
 - **Trace Source, Resource Type, and Search By**
Select filters from the drop-down list.
When you select **Trace name** for **Search By**, you need to select a specific trace name.
When you select **Resource ID** for **Search By**, you also need to select or enter a specific resource ID.
When you select **Resource name** for **Search By**, you also need to select or enter a specific resource name.
 - **Operator**: Select a specific operator (a user other than the tenant).
 - **Trace Status**: Select **All trace statuses**, **Normal**, **Warning**, or **Incident**.
 - **Start Date and End Date**: You can specify a time range to query traces.
- Step 6** Locate the required trace and click  on the left of the trace to view details.
- Step 7** Click **View Trace** in the **Operation** column. In the displayed dialog box, the trace structure details are displayed.

----End

3.12 Monitoring and Alarm Configuration

3.12.1 GeminiDB Cassandra Metrics

This section describes GeminiDB Cassandra metrics reported to Cloud Eye as well as their namespaces and dimensions. You can use APIs provided by Cloud Eye to query the metrics of the monitored object and alarms generated.

Namespace

SYS.NoSQL

Metrics

NOTE

You can view metrics on instance nodes by referring to [Viewing Metrics](#).

Table 3-19 GeminiDB Cassandra metrics

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
nosql005_disk_usage	Storage Space Usage	Storage space usage of the monitored object. Unit: Percent	0–100	GeminiDB Cassandra instances	1 minute
nosql006_disk_total_size	Total Storage Space	Total storage space of the monitored object. Unit: GB	≥ 0	GeminiDB Cassandra instances	1 minute
nosql007_disk_used_size	Used Storage Space	Used storage space of the monitored object. Unit: GB	≥ 0	GeminiDB Cassandra instances	1 minute
nosql009_dfv_write_delay	Storage Write Latency	Average delay of writing data to the storage layer in a specified period Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
nosql010_dfv_read_delay	Storage Read Latency	Average latency of reading data from the storage layer in a specified period Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra001_cpu_usage	CPU Usage	CPU usage of an instance Unit: Percent	0–100	GeminiDB Cassandra instance nodes	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cassandra002_mem_usage	Memory Usage	Memory usage of the instance Unit: Percent	0-100	GeminiDB Cassandra instance nodes	1 minute
cassandra003_bytes_out	Network Output Throughput	Outgoing traffic in bytes per second Unit: byte/s	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra004_bytes_in	Network Input Throughput	Incoming traffic in bytes per second Unit: byte/s	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra014_connections	Active Node Connections	Total number of connections attempting to connect to Cassandra instance nodes Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra015_read_latency	Average Read Latency	Average amount of time consumed by read requests Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra016_write_latency	Average Write Latency	Average amount of time consumed by write requests Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cassandra037_pending_write	Suspended Write Tasks	Number of write tasks in waiting status Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra038_pending_read	Suspended Read Tasks	Number of read tasks in waiting status Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra044_range_slice_latency	Scan Duration	Average amount of time consumed by scan operations Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra049_dropped_mutation	Dropped Writes	Average number of dropped writes Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra052_dropped_read	Dropped Reads	Average number of dropped reads Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra092_load_info	Data Volume on a Node	Data volume on a node Unit: byte	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra093_write_count_latency	Accumulated Write Requests	Total number of write requests initiated by a node Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cassandra094_write_1min_rate	Average Write Rate in the Last Minute	Average write rate in the last minute Unit: count/s	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra095_write_p75_latency	p75 Write Latency	p75 write latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra096_write_p95_latency	p95 Write Latency	p95 write latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra097_write_p99_latency	p99 Write Latency	p99 write latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra098_read_count_latency	Accumulated Read Requests	Total number of read requests initiated by a node Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra099_read_1min_rate	Average Read Rate in the Last Minute	Average read rate in the last minute Unit: count/s	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra100_read_p75_latency	p75 Read Latency	p75 read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cassandra101_read_p95_latency	p95 Read Latency	p95 read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra102_read_p99_latency	p99 Read Latency	p99 read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra103_range_slice_count_latency	Accumulated Range Read Requests	Accumulated range read requests Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra104_range_slice_1min_rate	Average Range Read Rate in the Last Minute	Average range read rate in the last minute Unit: count/s	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra105_range_slice_p75_latency	p75 Range Read Latency	p75 range read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra106_range_slice_p95_latency	p95 Range Read Latency	p95 range read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra107_range_slice_p99_latency	p99 Range Read Latency	p99 range read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cassandra163_write_p999_latency	p999 Write Latency	p999 write latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra164_read_p999_latency	p999 Read Latency	p999 read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra165_large_partition_num	Big Keys	Number of big keys on the current node Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra166_write_max_latency	Maximum Write Latency	Maximum write latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra167_read_max_latency	Maximum Read Latency	Maximum read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra168_imbalance_table_num	Tables with Uneven Data Distribution	Number of tables in which data is not evenly distributed. Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute

Dimensions

Key	Value
cassandra_cluster_id	Cluster ID of the GeminiDB Cassandra instance
cassandra_node_id	Node ID of the GeminiDB Cassandra instance

3.12.2 Configuring Alarm Rules

Setting alarm rules allows you to customize objects to be monitored and notification policies so that you can closely monitor your instances.

Alarm rules include the alarm rule name, instance, metric, threshold, monitoring interval, and whether to send notifications. This section describes how to set alarm rules.

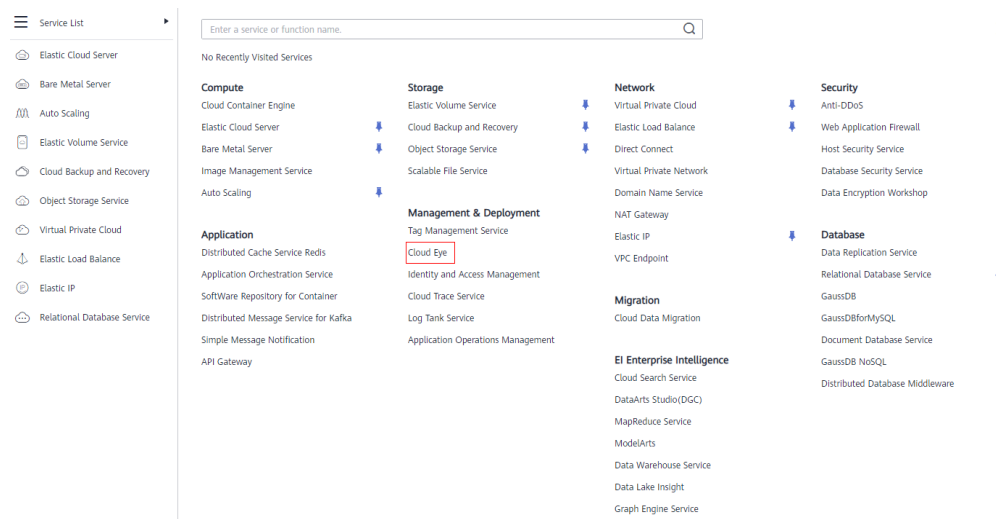
 **NOTE**

For more information about alarm rules, see *Cloud Eye User Guide*.

Procedure

- Step 1** Log in to the management console.
- Step 2** Click **Service List**. Under **Management & Deployment**, click **Cloud Eye**.

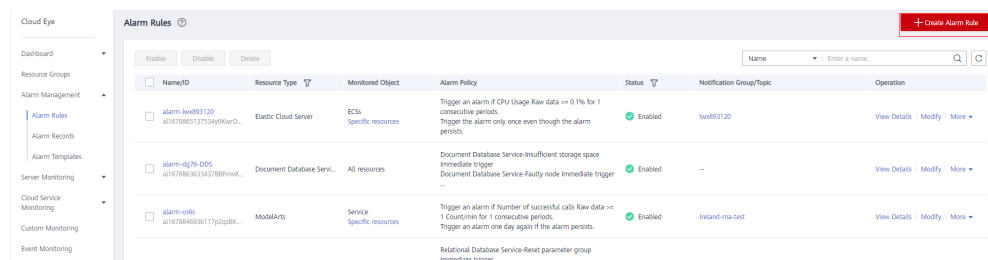
Figure 3-60 Selecting Cloud Eye



Step 3 In the navigation pane on the left, choose **Alarm Management > Alarm Rules**.

Step 4 On the **Alarm Rules** page, click **Create Alarm Rule**.

Figure 3-61 Creating an alarm rule



Step 5 Set alarm parameters.

1. Configure basic alarm information.

Figure 3-62 Configuring basic information for an alarm rule

* Name

Description

0/256

Table 3-20 Basic alarm rule information

Parameter	Description	Example Value
Name	Name of the rule. The system generates a random name and you can modify it.	alarm-cag2
Description	(Optional) Alarm rule description.	-

2. Select objects to be monitored and specify the monitoring scope.

Figure 3-63 Configuring objects to be monitored

* Alarm Type Metric Event

* Resource Type GaussDB PostgreSQL

* Dimension Cassandra - Cassandra Nodes

* Monitoring Scope Resource groups Specific resource

DB instance name

Name ID

nosql-test 686cc95e2624a42968359995f84e93dn06

Select All

nosql-test_nosql_prim_node_3 a4f8ae9fbac466ba94c833b228can06

nosql-test_nosql_prim_node_1 0e2ed575ac814100095a50426520no06

nosql-test_nosql_prim_node_2 3-f56b-a6518aa48ea907-f90701a10714-f906v06


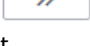
DB instance name

Name ID

No data available

Table 3-21 Parameter description

Parameter	Description	Example Value
Alarm Type	Alarm type that the alarm rule is created for. The value can be Metric or Event .	Metric
Resource Type	Type of the resource the alarm rule is created for. Select GeminiDB .	-
Dimension	Metric dimension of the alarm rule. Select Cassandra - Cassandra Nodes .	-

Parameter	Description	Example Value
Monitoring Scope	Monitoring scope the alarm rule applies to. NOTE <ul style="list-style-type: none"> - If you select Resource groups and any resource in the group meets the alarm policy, an alarm notification will be sent. - After you select Specific resources, select  to one or more resources and click  to add them to the box on the right. 	Specified resources
Group	This parameter is mandatory when Monitoring Scope is set to Resource groups .	-

3. Configure an alarm policy.

Figure 3-64 Configuring an alarm policy




Table 3-22 Parameter description


Parameter	Description	Example Value
Method	Select Associate template , Use existing template , or Configure manually . NOTE If you set Monitoring Scope to Specific resources , you can set Method to Use existing template .	Configure manually
Template	Select the template to be used. This parameter is available only when you set Method to Use existing template .	-


Parameter	Description	Example Value
Alarm Policy	<p>Policy for triggering an alarm. You can configure the threshold, consecutive periods, alarm interval, and alarm severity based on service requirements.</p> <ul style="list-style-type: none"> – Metric Name: specifies the the metric that the alarm rule is created for. The following metrics are recommended: <ul style="list-style-type: none"> Storage Space Usage, which is used to monitor the storage usage of GeminiDB Cassandra instances. If the storage usage is greater than 80%, scale up the storage in a timely manner by referring to Scaling Up Storage Space. CPU Usage and Memory Usage, which are used to monitor the compute resource usage of each GeminiDB Cassandra instance node. If the CPU usage or memory usage is greater than 80%, you can add nodes or upgrade node specifications in a timely manner. For more metrics, see GeminiDB Cassandra Metrics. – Alarm Severity: specifies the severity of the alarm. Valid values are Critical, Major, Minor, and Informational. <p>NOTE A maximum of 50 alarm policies can be added to an alarm rule. If any one of these alarm policies is met, an alarm is triggered.</p>	<p>Take the CPU usage as an example. The alarm policy configured in Figure 3-64 indicates that a major alarm notification will be sent to users every 10 minutes if the original CPU usage reaches 80% or above for three consecutive periods.</p>

4. Configure alarm notification information.

Figure 3-65 Configuring alarm notification information

Alarm Notification

* Notification Object 
Create an SMN topic and click refresh to make it available for selection.

* Notification Window Daily - 

* Trigger Condition Generated alarm Cleared alarm

Table 3-23 Parameter description

Parameter	Description	Example Value
Alarm Notification	<p>Whether to notify users when alarms are triggered. Notifications can be sent by email, text message, or HTTP/HTTPS message.</p> <p>Enabling alarm notification is recommended. When the metric data reaches the threshold set in the alarm rule, Cloud Eye immediately notifies you through SMN that an exception has occurred.</p>	Enabled Alarm Notification .
Notification Object	<p>Object that receives alarm notifications. You can select the account contact or a topic.</p> <ul style="list-style-type: none"> - Account contact is the mobile phone number and email address of the registered account. - Topic is used to publish messages and subscribe to notifications. If the required topic is unavailable, create one first and add subscriptions to it. 	-
Notification Window	<p>Cloud Eye sends notifications only within the notification window specified in the alarm rule.</p> <p>For example, if Notification Window is set to 00:00-8:00, Cloud Eye sends notifications only within 00:00-08:00.</p>	-
Trigger Condition	<p>Condition for triggering an alarm notification. You can select Generated alarm (when an alarm is generated), Cleared alarm (when an alarm is cleared), or both.</p>	-

5. Configure advanced settings.

Figure 3-66 Advanced settings**Table 3-24** Parameter description

Parameter	Description	Example Value
Enterprise Project	Enterprise project that the alarm rule belongs to. Only users with the enterprise project permissions can view and manage the alarm rule.	default

Step 6 After the configuration is complete, click **Create**.

When the metric data reaches the threshold set in the alarm rule, Cloud Eye immediately notifies you through SMN that an exception has occurred.

----End

3.12.3 Viewing Metrics

Cloud Eye monitors GeminiDB Cassandra statuses. You can obtain the GeminiDB Cassandra metrics on the management console.

Monitored data requires a period of time for transmission and display. The status of the monitored object displayed on the Cloud Eye page is the status obtained 5 to 10 minutes before. You can view the monitored data of a newly created DB instance 5 to 10 minutes later.

Precautions

- The DB instance is running properly.
Cloud Eye does not display the metrics of a faulty or deleted DB instance. You can view the monitoring information only after the instance is restarted or recovered.
- The DB instance has been properly running for at least 10 minutes.
The monitoring data and graphics are available for a new DB instance after the instance runs for at least 10 minutes.

Procedure

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases** > **GeminiDB**.

- Step 3** On the **Instance** page, click the instance whose metrics you want to view and click its name.
- Step 4** In the **Node Information** area on the **Basic Information** page, click **View Metric** in the **Operation** column.


Figure 3-67 Viewing metrics

Node Information

Name/ID	Status	AZ	Private IP Address	EIP	Operation
nsqj-6693_prim_node_1 68eb71550b534744b07044558a0d55dedno06	Available	az2pod1gz	192.168.139.180	Unbound	View Metric Bind EIP
nsqj-6693_prim_node_2 c31c0a639464f99a024538753bda7no06	Available	az2pod1gz	192.168.179.92	Unbound	View Metric Bind EIP
nsqj-6693_prim_node_3 95c13f81641f4008b4a095a3fe0672bdrno06	Available	az2pod1gz	192.168.128.224	Unbound	View Metric Bind EIP

- Step 5** In the monitoring area, select a time range to view monitoring data.

You can view the monitoring data in the last 1, 3, or 12 hours.

To view the monitoring curve in a longer time range, click  to enlarge the graph.

----End

3.13 Enterprise Project

3.13.1 Overview

An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is **default**.

You can also customize enterprise projects to meet your service requirements. For details, see [Enterprise Management User Guide](#).

3.13.2 Managing Quotas

GeminiDB Cassandra API provides a quota function that allows you to manage resources by controlling the number of resources in each enterprise project to ensure that resources can be used and managed properly.

This section describes how to query used resources in each enterprise project and its resource quotas.

This function is in the open beta test (OBT) phase. To use the function, contact customer service.

Viewing Resource Quotas in Each Enterprise Project

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases** > **GeminiDB**.
- Step 3** In the navigation pane on the left, choose **My Quotas** to view quota details of the current enterprise project.

Figure 3-68 Viewing quotas

The screenshot shows the 'Quota Management' interface. At the top right, there is a search bar with the placeholder text 'Enter an enterprise project.' and a search icon. Below the search bar is a table with the following columns: 'Enterprise Project', 'Used/Total DB Instances', 'Used/Total vCPUs', 'Used/Total Memory (GB)', and 'Operation'. The table contains 11 rows of data, each representing an enterprise project with its respective resource usage and a link to edit the quotas.

Enterprise Project	Used/Total DB Instances	Used/Total vCPUs	Used/Total Memory (GB)	Operation
default	1/1297	12/133583	48/11108044	Edit
QA	0/2212	0/10011	0/1799999	Edit
EPS0	0/1111	0/121	0/481	Edit
P8	0/100	0/1000	0/1000	Edit
P4	0/0	0/0	0/0	Edit
P3	0/0	0/0	0/0	Edit
P1	0/1001	0/11	0/11111	Edit
P6	0/0	0/0	0/0	Edit
P5	0/0	0/0	0/0	Edit
P7	0/0	0/0	0/0	Edit

At the bottom of the table, there is a pagination control showing 'Total Records: 11' and page navigation arrows.

Table 3-25 Parameter description

Parameter	Description
Enterprise Project	Enterprise project that an instance belongs to.
Used/Total DB Instances	Number of used instances in the current enterprise project
Used/Total vCPUs	vCPUs of all instances in the current enterprise project
Used/Total Memory (GB)	Memory of all instances in the current enterprise project

 **NOTE**

If there are no resources in an enterprise project, the default quota is 0. Before creating an instance, you need to set quotas first by referring to [Modifying Resource Quotas of an Enterprise Project](#).

----End

Modifying Resource Quotas of an Enterprise Project

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases** > **GeminiDB**.
- Step 3** In the navigation pane on the left, choose **My Quotas**. In the quota list, select the enterprise project you want to set quotas for and click **Modify** in the **Operation** column.

Figure 3-69 Modifying quotas

Table 3-26 Quota management

Parameter	Value Range
DB Instances	0–5,000
vCPUs	0–8,000,000
Memory (GB)	0–16,000,000

----End

3.14 Billing Management

3.14.1 Renewing Instances

This section describes how to renew your yearly/monthly GeminiDB Cassandra instances.

Precautions

Pay-per-use instances do not support this function.

Renewing a Yearly/Monthly Instance

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, locate the instance that you want to renew and click **Renew** in the **Operation** column.

Figure 3-70 Renewing an instance

Name/ID	DB Instance Type	Compatible API	Status	Enterprise Project	Billing Mode	Operation
nosql-sk3 ba7a27ce9f6d431d9b10c7598fa50243n12	Proxy-based general pu...	Redis 5.0	Available	default	Pay-per-Use Created on May 05, 2023 10...	Log In Change to Yearly/Monthly More ▾
nosql-sk1 83516d54bc10486bacc5162df881e44c1n12	Proxy-based general pu...	Redis 5.0	Available	default	Yearly/Monthly 31 days until expiration	Log In Renew More ▾

Alternatively, click the instance name to go to the **Basic Information** page. In the **Billing Information** area, click **Renew** next to the **Billing Mode** field.

Figure 3-71 Renewing an instance

Billing Information

Billing Mode: Yearly/Monthly **Renew**

Order: [blurred]

Created: Jun 17, 2022 11:41:11 GMT+08:00

Expiration Date: Jul 17, 2022 23:59:59 GMT+08:00

Step 4 On the displayed page, renew the instance.

----End

Renewing Instances

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases > GeminiDB**.

Step 3 On the **Instances** page, select the instances that you want to renew and click **Renew** above the instance list.

Figure 3-72 Renewing instances

Name/ID	DB Instance Type	Compatible API	Status	Billing Mode	Operation
<input checked="" type="checkbox"/>	[blurred]	Proxy-based general purpose	Available	Yearly/Monthly 30 days until expiration	Renew Change to Pay-per-Use More ▾
<input checked="" type="checkbox"/>	[blurred]	Proxy-based general purpose	Available	Yearly/Monthly 30 days until expiration	Renew Change to Pay-per-Use More ▾

Step 4 In the displayed dialog box, click **Yes**.

----End

3.14.2 Changing the Billing Mode from Pay-per-Use to Yearly/Monthly

This section describes how to change the billing mode of a GeminiDB Cassandra instance from pay-per-use to yearly/monthly. If you use the service for a long

time, you can change the billing mode of a DB instance from pay-per-use to yearly/monthly for lower costs.

Precautions

Only when the status of a pay-per-use instance is **Available**, its billing mode can be changed to yearly/monthly.


Changing the Billing Mode of a Single Instance

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, locate the instance whose billing mode you want to change and choose **Change to Yearly/Monthly** in the **Operation** column.

Figure 3-73 Changing the billing mode from pay-per-use to yearly/monthly

Name/ID	DB Instance Type	Compatible API	Status	Billing Mode	Operation
nooqi-43d2 a9d5f6c768714809d969217d4e81b55c0m12	Cluster	Redis 5.0	Available	Pay-per-use	Create Backup Restart Delete
nooqi-ed4d f18290384084b0390c4a35454599e2b3m06	Cluster	Cassandra 3.11	Available	Pay-per-use	Log In Change to Yearly/Monthly More

- Step 4** On the displayed page, select the renewal duration in month. The minimum duration is one month.
If you do not need to modify your settings, click **Pay Now**.
- Step 5** Select a payment method and click **Pay**.
- Step 6** View the results on the **Instances** page.

In the upper right corner of the DB instance list, click  to refresh the list. The instance status will become **Available** after the change is successful. The billing mode becomes to **Yearly/Monthly**.

----End

Changing the Billing Mode of Multiple Instance in Batches

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, select the instances whose billing mode you want to change and click **Change to Yearly/Monthly** above the DB instance list. In displayed dialog box, click **Yes**.

Figure 3-74 Changing the billing mode of multiple instances


Name/ID	DB Instance Type	Compatible API	Status	Billing Mode	Operation
X11451d5372m12			Available	Yearly/Monthly 30 days until expiration	Renew Change to Pay-per-Use More
9e618b39ccdm12			Available	Yearly/Monthly 30 days until expiration	Renew Change to Pay-per-Use More
2f3ea83efdein06			Available	Pay-per-use Created on Jun 17, 202...	Log In Change to Yearly/Monthly More
14706d40ccdm12			Available	Pay-per-use Created on Jun 17, 202...	Change to Yearly/Monthly Change Specificatio

Step 4 On the displayed page, select the renewal duration in month. The minimum duration is one month.

If you do not need to modify your settings, click **Pay Now**.

Step 5 Select a payment method and click **Pay**.

Step 6 View the results on the **Instances** page.

In the upper right corner of the instance list, click  to refresh the list. The instance status will become **Available** after the change is successful. The billing mode changes to **Yearly/Monthly**.

----End

3.14.3 Changing the Billing Mode from Yearly/Monthly to Pay-per-Use

You can change yearly/monthly GeminiDB Cassandra instances to pay-per-use so you only pay for the actual usage of your resources.

Precautions

- The billing mode of a yearly/monthly instance can only be changed to pay-per-use when the instance is in the **Available** status.
- Auto renewal will be disabled after the billing mode of your instances is changed to pay-per-use. Exercise caution when performing this operation.

Changing the Billing Mode of a Single Instance to Pay-per-Use

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases > GeminiDB**.

Step 3 On the **Instances** page, locate the instance whose billing mode you want to change and click **Change to Pay-per-Use** in the **Operation** column.

Figure 3-75 Changing yearly/monthly to pay-per-Use

Name/ID	DB Instance Type	Compatible API	Status	Billing Mode	Operation
nosql-83d2 a56dfec768714809a96217de81b550m12	Cluster	Redis 5.0	Available	Pay-per-use	Create Backup Restart Delete
nosql-edc4 f1e290384d84b0390c4a3545599e2b3m06	Cluster	Cassandra 3.11	Available	Pay-per-use	Log In Change to Yearly/Monthly More
nosql-f03 bf0fad8208647b69a4407796b021957m06	Cluster	Cassandra 3.11	Changing instance class	Yearly/Monthly	Log In Renew More
nosql-7896_c 06adff8d9a4c420929a51e2a24951cm06	Cluster	Cassandra 3.11	Available Pending restart	Yearly/Monthly	Log In Renew More
nosql-79a8 8a51ab7a519443e9a22bd3c7a901e1afm06	Cluster	Cassandra 3.11	Available	Pay-per-use	Change to Pay-per-Use Create Backup Restart Reset Password Unsubscribe
nosql-3516 91f5994ce38c4519a9c7f4948ba43580m06	Cluster	Cassandra 3.11	Available	Yearly/Monthly	Log In Renew More

Step 4 On the displayed page, confirm the instance information and click **Change to Pay-per-Use** to submit the change. The billing mode will change to pay-per-use after the DB instance expires.

Step 5 After you submit the change, a message is displayed in the **Billing Mode** column of the target DB instance, indicating that the billing mode will be changed to pay-per-use after the DB instance expires.

Step 6 To cancel the change, choose **Billing > Renewal** to enter the Billing Center. On the **Renewals** page, locate the target DB instance and click **More > Cancel Change to Pay-per-Use**.

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

----End

Changing the billing mode of multiple instances to pay-per-use

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases > GeminiDB**.

Step 3 On the **Instances** page, select the instances whose billing mode you want to change and click **Change to Pay-per-Use** above the instance list.

Figure 3-76 Changing the billing mode of multiple instances to pay-per-use

Name/ID	DB Instance Type	Compatible API	Status	Billing Mode	Operation
1d5372im12			Available	Yearly/Monthly 30 days until expiration	Renew Change to Pay-per-Use More
b70685178c			Available	Yearly/Monthly 30 days until expiration	Renew Change to Pay-per-Use More

Step 4 In the displayed dialog box, click **Yes**.

Step 5 On the displayed page, confirm the instance information and click **Change to Pay-per-Use**. The billing mode will change to pay-per-use after the instance expires.

NOTICE

Auto renewal will be disabled after the billing mode of your instances is changed to pay-per-use. Exercise caution when performing this operation.

Step 6 After you submit the change, check whether a message is displayed in the **Billing Mode** column, indicating that the billing mode will be changed to pay-per-use after the instance expires.

Step 7 To cancel the change, choose **Billing > Renewal** to enter the Billing Center. On the **Renewals** page, locate the instance and click **More > Cancel Change to Pay-per-Use**.

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

----End

3.14.4 Unsubscribing from a Yearly/Monthly Instance

If you do not need a yearly/monthly instance any longer, unsubscribe from it.

Precautions

- Unsubscribed operations cannot be undone. Exercise caution when performing this operation. To retain data, create a manual backup before unsubscription. For details, see [Creating a Manual Backup](#).
- After an unsubscription request is submitted, resources and data will be deleted and cannot be retrieved. Ensure that the manual backup is complete before submitting the unsubscription request.

Unsubscribing from a Single Yearly/Monthly Instance

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases > GeminiDB**.
- Step 3** On the **Instances** page, locate the instance you want to unsubscribe from and choose **More > Unsubscribe** in the **Operation** column.

Figure 3-77 Unsubscribing from a yearly/monthly instance

Name/ID	DB Instance Type	Compatible API	Status	Billing Mode	Operation
nosql-03d2 a66dfec76871480b9a96217de81b5d0m12	Cluster	Redis 5.0	Available	Pay-per-use	Create Backup Restart Delete
nosql-edc4 f16290384d864b0390c4a35d4599e2b3in06	Cluster	Cassandra 3.11	Available	Pay-per-use	Log In Change to Yearly/Monthly More
nosql-f03 b10fad8208647b289a4e07790b021957in06	Cluster	Cassandra 3.11	Changing Instance class	Yearly/Monthly	Log In Renew More
nosql-7996 96a0ff8e98ac4c02929a51e2a24d51cin06	Cluster	Cassandra 3.11	Available Pending restart	Yearly/Monthly	Log In Renew More
nosql-79a8 8a51ab7a519443e9a22bd3c7a901e1afin06	Cluster	Cassandra 3.11	Available	Pay-per-use	Log In Renew More
nosql-3516 91f5994ce38c61519a9c7f4948ba3580in06	Cluster	Cassandra 3.11	Available	Yearly/Monthly	Log In Renew More

- Step 4** In the displayed dialog box, click **Yes**.
- Step 5** On the displayed page, confirm the unsubscription and select a reason. Then, click **Confirm**.
- For details, see [Unsubscription Rules](#).
- Step 6** In the displayed dialog box, click **Yes**.

NOTICE

1. After an unsubscription request is submitted, resources and data will be deleted and cannot be retrieved.
2. If you want to retain data, complete a manual backup before submitting the unsubscription request.

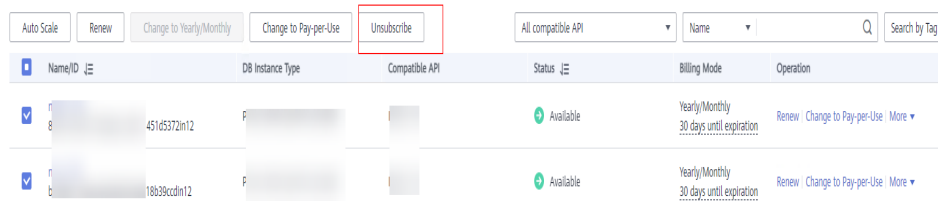
- Step 7** View the unsubscription result. After the instance order is successfully unsubscribed, the instance is no longer displayed in the instance list on the **Instances** page.

----End

Unsubscribing from Multiple Yearly/Monthly Instances

- Step 1** Log in to the management console.
- Step 2** In the service list, choose **Databases** > **GeminiDB**.
- Step 3** Choose **Instances** in the navigation pane on the left, select the instances you want to unsubscribe from and click **Unsubscribe** above the instance list.

Figure 3-78 Unsubscribing from multiple yearly/monthly instances



Name/ID	DB Instance Type	Compatible API	Status	Billing Mode	Operation
45165372m12	F		Available	Yearly/Monthly 30 days until expiration	Renew Change to Pay-per-Use More
18639ccdm12	F		Available	Yearly/Monthly 30 days until expiration	Renew Change to Pay-per-Use More

- Step 4** In the displayed dialog box, click **Yes**.
- Step 5** On the displayed page, confirm the unsubscription and select a reason. Then, click **Confirm**.
- For details, see [Unsubscription Rules](#).
- Step 6** In the displayed dialog box, click **Yes**.

NOTICE

- After an unsubscription request is submitted, resources and data will be deleted and cannot be retrieved.
- If you want to retain data, complete a manual backup before submitting the unsubscription request.

- Step 7** View the unsubscription result. After the instances are successfully unsubscribed from, they are no longer displayed in the instance list any longer on the **Instances** page.

----End

4 FAQs

4.1 Product Consulting

4.1.1 What Should I Pay Attention to When Using GeminiDB Cassandra?

1. DB instance operating systems (OSs) are invisible to you. Your applications can access a database only through an IP address and a port.
2. The backup files stored in OBS and the system containers are invisible to you. They are visible only in the GeminiDB Cassandra management system.
3. Precautions after purchasing DB instances:
After purchasing DB instances, you do not need to perform basic database O&M operations, such as applying HA and security patches, but you should still note:
 - a. The CPU, input/output operations per second (IOPS), and space are sufficient for the DB instances.
 - b. The DB instance has performance problems and whether optimization is required.

4.1.2 What Is GeminiDB Cassandra Instance Availability?

The formula for calculating the instance availability is as follows:

$$\text{DB instance availability} = (1 - \text{Failure duration} / \text{Total service duration}) \times 100\%$$

The failure duration refers to the total duration of faults that occur during the running of a DB instance after you buy the instance. The total service duration refers to the total running time of the DB instance.

4.2 Billing

4.2.1 What Are the Differences Between Yearly/Monthly and Pay-per-use Billing Mode?

Yearly/Monthly is a prepaid billing mode in which resources are billed based on the service duration. This cost-effective mode is ideal when the duration of resource usage is predictable. It is recommended for long-term users.

Pay-per-use is a post payment mode, so you can start or stop an instance at any time. Pricing is listed on a per-hour basis, but bills are calculated based on the actual usage duration.

4.2.2 Can I Switch Between Yearly/Monthly and Pay-per-Use Payments?

You can change the billing mode from yearly/monthly to pay-per-use or vice versa.

- If you want to change the billing mode from yearly/monthly to pay-per-use, see [Changing the Billing Mode from Yearly/Monthly to Pay-per-Use](#).
- If you want to change the billing mode from pay-per-use to yearly/monthly, see [Changing the Billing Mode from Pay-per-Use to Yearly/Monthly](#).

4.3 Database Usage


4.3.1 Why Does the Overall Instance Performance Deteriorate When QPS Increases After the Batch Size Is Decreased?

Symptom

The original **batch_size** was 100, and the size of a single row was about 400 bytes. **batch_size** was then changed to 10 because an alarm was triggered when the batch size reached 5 KB. To ensure the overall write performance, QPS was 10 times of the original QPS. However, the overall performance deteriorated after the changes.

Possible Cause

The number of concurrent clients is restricted by the Driver configuration parameters, including the number of hosts, number of sessions, **ConnectionsPerHost**, and **MaxRequestsPerConnection**.



For example, a user starts a cluster, creates a session for the cluster, and has three hosts. If **ConnectionsPerHost** is set to 2 and **MaxRequestsPerConnection** uses the default value 128, the maximum number of concurrent requests of the session is 768, and the maximum number of requests of a single node is 256.

For details about the parameters, see the [official document](#).

Solution

View [monitoring metrics](#) to observe the CPU usage, read/write pending, and read/write latency of a single node.

- If the load of a single node reaches the upper limit, you need to add nodes. For details, see [Adding Nodes](#).
- If the load of a single node is low, you need to adjust the configuration of Driver.
 - a. Increase the value of **ConnectionsPerHost**. Ensure that the total number of connections to the cluster does not exceed the configured alarm threshold.
 - b. Increase the value of **MaxRequestsPerConnection**. Ensure that the value does not exceed the load capability of a single node. Observe the CPU usage, read/write latency, and read/write pending.

4.3.2 What Can I Do if Error "field larger than field limit (131072)" Is Reported During Data Import?

Symptom

When you import data, the size of a single column exceeds 128 KB. As a result, the Python CSV single-column restriction is triggered.

Error message:

```
field larger than field limit (131072)
```

Possible Cause

When Python CSV reads a file, **csv.field_size_limit** limits the size of a single column.

Solution

Step 1 Run the following commands in the **cqlsh** directory to find the **cqlshrc** file:

```
touch cqlshrc  
rm -rf ~/.cassandra/cqlshrc*
```

Step 2 Add the following information in the **cqlshrc** file and save the file:

```
[csv]  
field_size_limit = 9223372036854775807
```

Step 3 Add the following parameters when connecting to an instance using cqlsh:

```
-cqlshrc=cqlshrc
```

Command example:

```
cqlsh 127.0.0.1 8635 -u rwuser -p password --cqlshrc=cqlshrc  
----End
```

4.3.3 What Should I Pay Attention to When Creating a GeminiDB Cassandra Table?

When you create tables in a GeminiDB Cassandra database, pre-allocate memory to guarantee database performance. GeminiDB Cassandra has a limit on the number of tables.

Precautions

- Half of node memory is allocated to the storage engine.
- An odd number of clusters can tolerate $N/2-1$ faulty nodes, and an even number of clusters can tolerate $N/2$ faulty nodes.

Calculating the Number of Tables

The memory required for creating tables depends on the instance specifications. Assume that an instance has 4 vCPUs and 16 GB memory and the size of a single table is 768 MB.

Maximum number of tables that can be created = Total available memory of the cluster / Memory required by a single table

- Cluster with an odd number of nodes
Available cluster memory = Node memory/2 x ($N/2 + 1$)
- Cluster with an even number of nodes
Available cluster memory = Node memory/2 x ($N/2$)

For example:

- Available memory of an instance with 3 nodes, 4 vCPUs, and 16 GB memory = $16/2 \times (3/2 + 1) = 16$ GB
Maximum number of created tables = $16 \times 1024 \text{ MB} / 768 \text{ MB} = 21$
- Available memory of an instance with 4 nodes, 4 vCPUs, and 16 GB memory = $16/2 \times (4/2) = 16$ GB
Maximum number of created tables = $16 \times 1024 \text{ MB} / 768 \text{ MB} = 21$
- Available memory of an instance with 5 nodes, 4 vCPUs, and 16 GB memory = $16/2 \times (5/2 + 1) = 24$ GB
Maximum number of created tables = $24 \times 1024 \text{ MB} / 768 \text{ MB} = 32$

For details about the mapping between the number of nodes (4 vCPUs, 16 GB) and the number of tables, see [Table 4-1](#).

Table 4-1 Upper limit on the number of tables

Instance Class	Number of Nodes	Number of Tables
4 vCPUs 16 GB	3	21
	4	21
	5	32
	6	32

Instance Class	Number of Nodes	Number of Tables
	7	42
	8	42
	9	53
	10	53
	11	64
	12	64

 **NOTE**

- A single table occupies 768 MB memory, and the default number of table tokens is 12. If tokens are separately set, calculate the number of tables using the following formula: $(768/12) \times \text{Number of tokens}$.
- The preceding formula is designed for common tables. If stream table is enabled, one stream table consumes resources 2.5 times more than common tables.

For details about the mapping between the number of nodes (8 vCPUs, 32 GB) and the number of tables, see [Table 4-2](#).

Table 4-2 Upper limit on the number of tables

Instance Class	Number of Nodes	Number of Tables
8 vCPUs 32 GB	3	22
	4	22
	5	34
	6	34
	7	45
	8	45
	9	56
	10	56
	11	68
	12	68

 **NOTE**

- A single table occupies 1440 MB memory, and the default number of table tokens is 12. If tokens are set separately, calculate the number of tables using the following formula: $(1440/12) \times \text{Number of tokens}$.
- The preceding formula is designed for common tables. If stream table is enabled, one stream table consumes resources 2.5 times more than common tables.

For details about the mapping between the number of nodes (16 vCPUs, 64 GB) and the number of tables, see [Table 4-3](#).

Table 4-3 Upper limit on the number of tables

Instance Class	Number of Nodes	Number of Tables
16 vCPUs 64 GB	3	45
	4	45
	5	68
	6	68
	7	91
	8	91
	9	113
	10	113
	11	136
	12	136

 **NOTE**

- A single table occupies 1440 MB memory, and the default number of table tokens is 12. If tokens are set separately, calculate the number of tables using the following formula: $(1440/12) \times \text{Number of tokens}$.
- The preceding formula is designed for common tables. If stream table is enabled, one stream table consumes resources 2.5 times more than common tables.

For details about the mapping between the number of nodes (32 vCPUs, 128 GB) and the number of tables, see [Table 4-4](#).

Table 4-4 Mapping between the number of nodes (32U128GB) and the number of tables

Instance Class	Number of Nodes	Number of Tables
32 vCPUs 128 GB	3	68
	4	68
	5	102

Instance Class	Number of Nodes	Number of Tables
	6	102
	7	136
	8	136
	9	170
	10	170
	11	204
	12	204

NOTE

- A single table occupies 1920 MB memory, and the default number of table tokens is 12. If tokens are separately set, calculate the number of tables using the following formula: $(1920/12) \times \text{Number of tokens}$
- The preceding formula is designed for common tables. If stream table is enabled, one stream table consumes resources 2.5 times more than common tables.

Parameters for Creating a Table

1. Throughput parameter: **ZOO_THROUGHPUT**, with the value of **big**.
 - Low throughput

```
CREATE TABLE test1 (k int,p int,s int static,v int,PRIMARY KEY (k, p)) WITH ZOO_THROUGHPUT = 'small';
```
 - Medium throughput

```
CREATE TABLE test2 (k int,p int,s int static,v int,PRIMARY KEY (k, p)) WITH ZOO_THROUGHPUT = 'medium';
```
 - High throughput

```
CREATE TABLE test3 (k int,p int,s int static,v int,PRIMARY KEY (k, p)) WITH ZOO_THROUGHPUT = 'big';
```
2. Number of table tokens: indicates the number of table tokens when a table is created. The number of tokens must be greater than 1.

```
CREATE TABLE test4 (k int,p int,s int static,v int,PRIMARY KEY (k, p)) WITH ZOO_TABLE_TOKENS = 24;
```
3. Table parameters: **ZOO_BUFFER_SIZE** and **ZOO_BUFFER_NUMBER** (not recommended).

When creating a table, you can specify the number of memtables in the storage layer and the size of each memtable.

- **ZOO_BUFFER_SIZE** is of the map type and specifies the CF name and value. The value ranges from 2 to 32.

```
CREATE TABLE test6 (k int,p int,s int static,v int,PRIMARY KEY (k, p)) WITH ZOO_BUFFER_SIZE = {'default': 16};
```
- **ZOO_BUFFER_NUMBER** is of the map type and specifies the CF name and value. The value ranges from 2 to 8.

```
CREATE TABLE test5 (k int,p int,s int static,v int,PRIMARY KEY (k, p)) WITH ZOO_BUFFER_NUMBER = {'default': 3};
```

 NOTE

If you need to adjust the table specifications after the table is created, for example, when the maximum number of the tables is reached, you can reduce the table specifications to create more tables by adjusting the following parameters.

- If you set the throughput of all created tables to medium, the number of tables can be doubled

```
ALTER TABLE keyspace_name.table_name WITH ZOO_THROUGHPUT = 'medium';
```
- If you set the throughput of all created tables to small, the number of tables can be tripled.

```
ALTER TABLE keyspace_name.table_name WITH ZOO_THROUGHPUT = 'small';
```

4.3.4 How Do I Detect and Resolve BigKey and HotKey Issues?

The Cassandra database is a highly scalable, high-performance, and distributed database. It is suitable for big data scenarios and can be used to manage a large amount of structured data. With continuous growth of service volume and data traffic, some service design defects are gradually exposed, which reduces the stability and availability of the cluster. For example, the primary key design is improper, or a single partition contains a large amount of data. As a result, the partition key is too large, the node load is unbalanced, and the cluster stability deteriorates. This type of problem is called BigKey. When the workload of access to a key exceeds the maximum workload that a server can handle, we can call it a HotKey. Generally, a BigKey is an indirect cause of a HotKey issue.

GeminiDB Cassandra is a cloud-native distributed NoSQL database with a decoupled compute and storage architecture and compatible with the Cassandra ecosystem. To solve the preceding issues, GeminiDB Cassandra provides real-time detection of BigKey and HotKey issues to help you design schemas and avoid service stability risks.

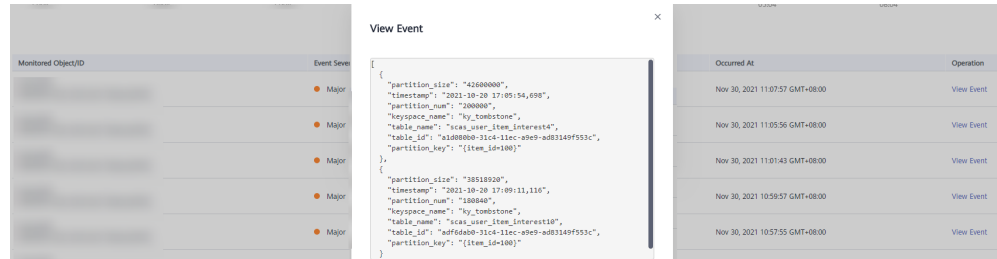
BigKey Issue

- Possible causes
The main cause of the BigKey issue is that the primary key design is improper. As a result, a single partition contains too many records or data. Once a partition becomes extremely large, the access to this partition increases the load of the server where the partition is located, and even causes the out of memory (OOM) issue.
- Troubleshooting
You can use either of the following methods to rectify BigKey issues:
 - Add caches and optimize the table structure.
 - Add a new partition key for hashing data. Split data to avoid too much data in a single partition.
- Check method
You can specify a threshold based on your service requirement. If any threshold is exceeded, a BigKey is generated.
 - a. The number of rows of a single partition key cannot exceed 100,000.
 - b. The size of a single partition cannot exceed 100 MB.

GeminiDB Cassandra supports BigKey detection and alarms. On the Cloud Eye console, you can configure BigKey alarms for instances. For details, see [Configuring Alarm Rules](#).

When a BigKey event occurs, the system sends a warning notification immediately. You can [view the event data](#) on the Cloud Eye page and handle the event in a timely manner to prevent service fluctuation.

Figure 4-1 BigKey alarm



The alarm is described as follows:

```
[
  {
    "partition_size": "1008293497", //Total size of oversized partition keys
    "timestamp": "2021-09-08 07:08:18,240", //Time when a BigKey is generated
    "partition_num": "676826", //Total number of rows for oversized partition keys
    "keyspace_name": "ssss", //keyspace name
    "table_name": "zzzz", //Table name
    "table_id": "024a1070-0064-11eb-bdf3-d3fe5956183b", //Table ID
    "partition_key": "{vin=TESTW3YWZD2021003}" //Partition key
  }
]
```

- Common cases and solutions

Case 1: The data volume of a cluster is too large. As a result, the cluster has large partition keys (more than 2,000 partition keys are checked), and the maximum size of a partition key has reached 38 GB. When services frequently access these large partition keys, the node load remains high, affecting the service request success rate.

The table structure is designed as follows.

```
CREATE TABLE movie (
  movieid text,
  appid int,
  uid bigint,
  accesstring text,
  moviename text,
  access_time timestamp,
  PRIMARY KEY (movieid, appid, uid, accesstring, moviename)
)
```

Table design analysis:

The **movie** table stores information about short videos. The partition key is movieid, and stores user information (uid). If movieid is a popular short video and tens of millions or even hundreds of millions of users like this short video, the size of the partition where the short video is located is large (38 GB).

Solution:

To solve the problem, perform the following steps:

a. Optimize the table structure.

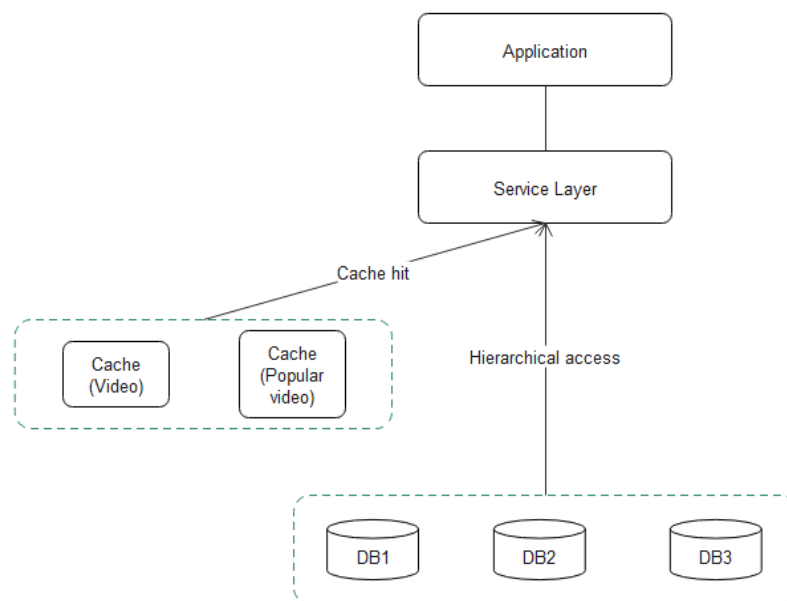
Create a table to store the short video information. Only public short video information is retained, and user information is not included. This ensures that the table does not generate large partition keys. Write the short video information to the table.

```
CREATE TABLE hotmovieaccess (
  movieid text,
  appid int,
  accesstring text,
  access_time timestamp,
  PRIMARY KEY (movieid, appid)
)
```

b. Add caches.

A service application first reads popular file information from the cache. If no information is found, the service application queries the database to reduce the number of database query times.

The overall optimization logic is as follows:



- i. The service applications query the cache first. If the data to be queried already exists in the cache, the results are directly returned.
- ii. If the data is not in the cache, the popular video cache, the **hot** table, and the **hotmovieaccess** table will be accessed in sequence.
- iii. If the **hotmovieaccess** table contains the results, the results are directly returned. If the **hotmovieaccess** table does not contain any record, the **movie** table is queried.
- iv. Cache the query results.

Case 2: The **movie_meta** table is created by month, and each table stores only the data of the current month. The initial design can reduce or avoid

large partition keys. Due to frequent service writes, a large number of popular video records are stored, generating large partitions.

```
CREATE TABLE movie_meta202110 (  
    path text,  
    moviename text,  
    movieid text,  
    create_time timestamp,  
    modify_mtime timestamp,  
    PRIMARY KEY (path, moviename)  
)
```

Solution:

A random number (0 to 999) is added to the new partition key. The information stored in the original partition is randomly and discretely stored to 1,000 partitions. After the new partition key is used, no new partition key whose size exceeds 100 MB is formed. The old partition key data whose size exceeds 100 MB expires as time goes by.

HotKey Problem

- Hazards of HotKey:

In daily life, when the hot news is clicked, viewed, and commented for tens of thousands of times in an application, large number of requests will be generated. In this case, the same key is frequently accessed within a short period of time. As a result, the CPU usage and load of the node where the key is located suddenly increase, affecting other requests on the node and decreasing the service success rate. Such scenarios include promotion of popular products and Internet celebrity live streaming. In these read-intensive scenarios, HotKey issues will be generated.

The HotKey issue has the following impacts:

- a. The traffic is centralized and reaches the upper limit of the physical NICs.
- b. Too many requests may cause the cache service to break down.
- c. The database breaks down, causing service avalanche.

- Troubleshooting

To solve the HotKey issue, perform the following steps:

- a. HotKeys must be considered in design to prevent them from being generated in a database.
- b. Add caches in the service side to reduce HotKey issues. Multi-level cache should be used to solve the HotKey issue (such as Redis + local level-2 cache).
- c. Disable hotspot keys. For example, configure a whitelist for HotKeys on the service side to shield HotKeys as required.

- Check method


A key whose access frequency is greater than 100,000 times per minute is defined as a HotKey.

HotKey events are classified into the following types: One is the Writes event, indicating a write hotspot, and the other is the Reads event, indicating a read hotspot.

GeminiDB Cassandra provides HotKey monitoring and alarms. On the Cloud Eye console, you can configure HotKey alarms for instances. For details, see [Configuring Alarm Rules](#).

When a BigKey event occurs, the system sends a warning notification immediately. You can [view the event data](#) on the Cloud Eye page and handle the event in a timely manner to prevent service fluctuation.

Figure 4-2 HotKey alarm



HotKey alarm description:

```
{
  "sampler_type": "WRITES", //Sampling type. The value can be WRITES or READS. WRITES
  indicates write, and READS indicates read.
  "partition_num": "2969", //Hotspot times of a partition key
  "keyspace_name": "performance", //Keyspace name
  "table_id": "a10f3bb0-3626-11ec-bbdf-63e05bbb4391", //Table ID
  "table_name": "stresstable", //Table name
  "partition_key": "85897376" //The value of the hotspot partition key.
}
```

Summary

If you use Cassandra for online services, you must follow related rules to minimize risks in the development and design phase. Proper design can reduce the probability of most risks.

- The design of any table must consider whether HotKey or BigKey will be generated and whether load skew will occur.
- A data expiration mechanism must be established. Data in a table cannot increase indefinitely without being deleted or expired.
- In read-intensive scenarios, a cache mechanism needs to be added to handle read hotspots and improve query performance.
- A threshold must be set for each primary key and row. Otherwise, the database performance and stability will be affected. If the threshold is exceeded, optimize the settings in a timely manner.

4.3.5 How Do I Set Up a Materialized View?

Concept

A materialized view is a standard CQL table that automatically maintains the consistency between the data that meets certain conditions and the data in the base tables.

Constraints

- The primary key of a materialized view must contain all primary keys of the base table. Static columns cannot be included in a materialized view.
- All columns that are part of the view primary key are restricted by the "IS NOT NULL" restriction, meaning that they cannot be null.
- In a materialized view, a CQL row must be mapped from the base table to another row of the view, meaning that the rows of the view and base table correspond to each other.
- The WHERE condition of the SELECT statement does not constrain non-primary key columns in a view, except the IS NOT NULL condition.

Figure 4-3 Example value

```
cs1sh>> CREATE MATERIALIZED VIEW mv AS SELECT v1, c1, pk2 FROM tv WHERE v2 > 2 AND v1 IS NOT NULL AND pk1 IS NOT NULL AND pk2 IS NOT NULL AND c1 IS NOT NULL AND c1 IS NOT NULL PRIMARY KEY ( (id, pk1), c1, c2, pk2);
InvalidRequest: Error from server: code=2000 [Invalid query] message=Non-primary key columns cannot be restricted in the SELECT statement used for materialized view creation (got restrictions on: v2)
```

- Static columns, counter, superColumn, and duration types are not supported.

Setting Up a Materialized View

1. Insert a record into the base table and query the result.

Example:

```
CREATE TABLE person (
  id int,
  name text,
  addr text,
  age int,
  email text,
  PRIMARY KEY (id, name));
```

Insert a record.

```
insert into person(id, name, age, addr, email) values (0, 'ruby', 26, 'beijing', 'ruby@email.com');
```

Query the result.

Figure 4-4 Querying the result

```
cqlsh:ks> SELECT * FROM person ;

 id | name | addr | age | email
----+----+----+----+----
  0 | ruby | beijing | 26 | ruby@email.com
(1 rows)
```

2. Create a materialized view.

```
CREATE MATERIALIZED VIEW person_addr AS
SELECT * from person WHERE id IS NOT NULL AND addr IS NOT NULL
AND name IS NOT NULL
primary key (addr, id, name);
```

The `system_schema.views` table records the association between views and base tables.

Figure 4-5 Mapping between views and base tables

```
cqlsh:ks> SELECT * FROM system_schema.views WHERE keyspace_name = 'ks' and view_name = 'person_addr';
@ Row 1
-----
keyspace_name | ks
view_name     | person_addr
base_table_id | 74445d39-ebc5-11e9-8065-91e8e817a0b6
base_table_name | person
bloom_filter_fp_chance | 0.01
caching        | {'keys': 'ALL', 'rows_per_partition': 'NONE'}
cdc            | null
comment       |
compaction     | {'class': 'org.apache.cassandra.db.compaction.SizeTieredCompactionStrategy', 'max_threshold': '32', 'min_threshold': '4'}
compression    | {'chunk_length_in_kb': '64', 'class': 'org.apache.cassandra.io.compress.LZ4Compressor'}
crc_check_chance | 1
dlocal_read_repair_chance | 0.1
default_time_to_live | 0
extensions     | {}
gc_grace_seconds | 864000
id             | d8d9fc40-ebc5-11e9-8065-91e8e817a0b6
include_all_columns | True
max_index_interval | 2048
memtable_flush_period_in_ms | 0
min_index_interval | 128
read_repair_chance | 0
speculative_retry | 99PERCENTILE
where_clause   | id IS NOT NULL AND addr IS NOT NULL AND name IS NOT NULL

(1 rows)
cqlsh:ks>
```

The query results that do not meet the condition are not displayed, for example, IS NOT NULL.

3. Insert a record in which the **addr** value is **null**.
insert into person(id, name, age, addr, email) values (1, 'mike', 30, null, 'mike@email.com');
 Query the data in the the base cassandra table and materialized view.

Figure 4-6 Querying the result

```
cqlsh:ks> SELECT * FROM person;

id | name | addr | age | email
-----+-----+-----+-----+-----
1 | mike | null | 30 | mike@email.com
0 | ruby | beijing | 26 | ruby@email.com

(2 rows)
cqlsh:ks> SELECT * FROM person_addr ;

addr | id | name | age | email
-----+-----+-----+-----+-----
beijing | 0 | ruby | 26 | ruby@email.com

(1 rows)
cqlsh:ks>
```

4. Delete the materialized view.
DROP MATERIALIZED VIEW person_addr;

Figure 4-7 Deleting a view

```
cqlsh:ks> DROP MATERIALIZED VIEW person_addr ;
cqlsh:ks>
```

4.3.6 How Do I Use a Secondary Index?

Concept

In a GeminiDB Cassandra database, a primary key is the primary index, which can be used to query records. If you want to query records without the primary key, you can use secondary indexes.

Secondary Index Principles

A secondary index creates a hidden indexed table. The primary key becomes one of the columns in the hidden table.

Assume that there is a **playlists** table. The table structure is as follows:

```
CREATE TABLE playlists (
  id int,
  song_id int,
  song_order int,
  album text,
  artist text,
  title text,
  PRIMARY KEY (id, song_id));
```

The query result is as follows.

Figure 4-8 Querying the result

id	song_id	album	artist	song_order	title
1	1	01 01010101	01 0101	1	01 01010101

If an index is created for the **artist** field, the hidden table structure is as follows.

Figure 4-9 Querying the result

artist	id
01 0101	1

(1 rows)

artist is the primary key of the index table. **id** and **song_id**, functioning as the primary key of the original table, become common columns.

In Which Scenario Is the Index Not Recommended?

- Too many duplicate values exist in a column.
For example, if a table contains 100 million records and the values of **artist** are the same, you are not advised to index the **artist** column.
- The **counter** column cannot be indexed.
- Columns that are frequently updated or deleted.

How Do I Use an Index?

1. Creating an index

CREATE INDEX artist_names ON playlists(artist);

Note: If the original table contains a large amount of data, indexed data needs to be rebuilt before queries.

You can query the **IndexInfo** table to check whether the index is recreated. If the name of the created index exists, it indicates that the indexed data has been rebuilt.

Figure 4-10 Querying the result

```
cqlsh:ks> SELECT *from system."IndexInfo";
```

table_name	index_name
ks	artist_names

(1 rows)

2. Query records by indexed column.

Figure 4-11 Querying the result

```
cqlsh:ks> SELECT *from playlists where artist ='Jay Chou';
```

id	song_id	album	artist	song_order	title
1	1	周杰伦	Jay Chou	1	周杰伦

(1 rows)

 **NOTE**

Each table can have multiple indexes, but the write performance may be affected.

4.3.7 How Do I Set Paging Query?

Specifying the Number of Rows Fetched in Each Page

The fetch size specifies how many rows will be fetched at once. When you create a cluster connection, you can set a fetch size for it.

```
Cluster cluster = Cluster.builder()
    .addContactPoint(contactPoint)
    .withPort(8636)
    .withQueryOptions(new QueryOptions().setFetchSize(20))
    .build();
```

After the setting is successful, for all sessions spawned with this configuration, the configured number of rows is fetched from the server at a time. When the cache (20 rows) is exhausted, the system triggers a request for fetching another 20 rows from the server and there can be a waiting period.

Obtaining the Next Page in Advance

If you need to manually fetch more rows in advance to avoid waiting and save them to the current result set, refer to the following code. When the result set has 10 rows left, submit a parallel request for fetching more rows from the server.

```
ResultSet rs = session.execute("select * from space3.table3;");
for (Row row : rs) {
    if (rs.getAvailableWithoutFetching() == 10 && !rs.isFullyFetched()){
        System.out.println("pre-fetch more rows. ");
        rs.fetchMoreResults();
    }
    System.out.println(row);
}
```

Saving and Reusing the Paging State

1. Save the current paging state.


```
PagingState pagingState = resultSet.getExecutionInfo().getPagingState();
String string = pagingState.toString();
byte[] bytes = pagingState.toBytes();
```
2. Load and reuse the current paging state.


```
PagingState pagingState = PagingState.fromString(string);
Statement st = new SimpleStatement("your query");
st.setPagingState(pagingState);
ResultSet rs = session.execute(st);
```

Note: The paging state can only be collected, stored, and reused. They cannot be modified or applied to other query statements.

NOTE

GeminiDB Cassandra API does not support offset queries, which means that you cannot skip any part of the result set and cannot fetch results within the specified index range. If you want to use offset queries, you can emulate them on the client side. You will get all results in order, but you can delete results that you do not need. For more advanced usage and introduction, see [DataStax Java Driver 3.11](#).

4.4 Database Connection

4.4.1 What Can I Do If Spark Failed to Connect to Cassandra?

Symptom

You used Spark to connect to the open-source Cassandra, data can be read properly, but an error was reported during the connection.

Error message is as follows.

```
at co.mega.tetris.analyzer.history.VehicleHistoryToGn5.main(VehicleHistoryToGn5.scala:12)
at co.mega.tetris.analyzer.history.VehicleHistoryToGn5.main(VehicleHistoryToGn5.scala)
Caused by: java.util.NoSuchElementException: No value present
at java.util.Optional.get(Optional.java:135)
at com.datastax.spark.connector.rdd.partitioner.CassandraPartitionGenerator.$anonfun$describeRing$1(CassandraPartitionGenerator.scala:49)
at com.datastax.spark.connector.cql.CassandraConnector.$anonfun$withSessionDo$1(CassandraConnector.scala:112)
at com.datastax.spark.connector.cql.CassandraConnector.closeResourceAfterUse(CassandraConnector.scala:129)
at com.datastax.spark.connector.cql.CassandraConnector.withSessionDo(CassandraConnector.scala:111)
at com.datastax.spark.connector.rdd.partitioner.CassandraPartitionGenerator.describeRing(CassandraPartitionGenerator.scala:48)
at com.datastax.spark.connector.rdd.partitioner.CassandraPartitionGenerator.partitions(CassandraPartitionGenerator.scala:80)
at com.datastax.spark.connector.rdd.CassandraTableScanRDD.getPartitions(CassandraTableScanRDD.scala:273)
at org.apache.spark.rdd.RDD.$anonfun$partitions$2(RDD.scala:276)
at scala.Option.getOrElse(Option.scala:189)
at org.apache.spark.rdd.RDD.partitions(RDD.scala:272)
at org.apache.spark.rdd.MapPartitionsRDD.getPartitions(MapPartitionsRDD.scala:49)
at org.apache.spark.rdd.RDD.$anonfun$partitions$2(RDD.scala:276)
at scala.Option.getOrElse(Option.scala:189)
at org.apache.spark.rdd.RDD.partitions(RDD.scala:272)
at org.apache.spark.SparkContext.runJob(SparkContext.scala:2152)
at org.apache.spark.internal.io.SparkHadoopWriter$.write(SparkHadoopWriter.scala:78)
... 39 more
```


Configuration Details

The following shows the components and account details.

- Component configuration details

Table 4-5 Configuration details

Component	Version
spark-cassandra-connector	2.5.1
spark	2.5.1
Open-source Cassandra	3.11
scala	2.12

- User: **user1** (created by user **rwuser**)

Possible Cause

- **user1** does not have the permission to query the keyspace system.
- The Spark version is incorrect.

Solution

1. Grant the keyspace system query permission to **user1** as user **rwuser**.
2. Use spark-cassandra-connector 2.4.1.

4.4.2 What Can I Do If an Error Occurs When I Use Java Driver and a Mapped IP Address to Connect to a Database?

Symptom

When you use Java Driver to connect to a GeminiDB Cassandra instance, a session was established using the mapped IP address, rather than the database private IP address, over port 8635. However, an error was found in the connection log, and connection information of port 9042 was displayed.

Figure 4-12 Log information

```

2021-09-22 16:20:53 [main] INFO com.datastax.driver.core.ClockFactory - Using java.lang.System clock to generate timestamps.
2021-09-22 16:20:53 [main] INFO com.datastax.driver.core.NettyUtil - Found Netty-native-epoll transport in the classpath, using it.
2021-09-22 16:20:54 [main] WARN com.datastax.driver.core.Cluster - You listed 192.168.0.54:8635 in your contact points, but it wasn't found in the control host's
system.peers at startup
2021-09-22 16:20:54 [main] WARN com.datastax.driver.core.Cluster - You listed 192.168.0.153:8635 in your contact points, but it wasn't found in the control host's
system.peers at startup
2021-09-22 16:20:54 [main] INFO com.datastax.driver.core.policies.DCAwareRoundRobinPolicy - Using data-center name 'datacenter1' for DCAwareRoundRobinPolicy (if this is
incorrect, please provide the correct datacenter name with DCAwareRoundRobinPolicy constructor)
2021-09-22 16:20:54 [main] INFO com.datastax.driver.core.Cluster - New Cassandra host /192.168.0.54:9042 added
2021-09-22 16:20:54 [main] INFO com.datastax.driver.core.Cluster - New Cassandra host /192.168.0.54:8635 added
2021-09-22 16:20:54 [main] INFO com.datastax.driver.core.Cluster - New Cassandra host /192.168.0.153:9042 added
2021-09-22 16:20:54 [JanusGraph Cluster:nic-worker-0] WARN com.datastax.driver.core.HostConnectionPool - Error creating connection to /192.168.0.54:9042
com.datastax.driver.core.exceptions.TransportException: [/192.168.0.54:9042] Cannot connect
at com.datastax.driver.core.Connection$1.operationComplete(Connection.java:224)

```

Possible Cause

Java Driver was not used correctly, as shown in [Figure 4-13](#). Do not use `addContactPointsWithPorts` when using Java Driver and do not map each IP address.

Figure 4-13 Incorrect usage of the Java Driver

```

public static void connectToCluster() {
    Cluster cluster = Cluster.builder()
        .addContactPointsWithPorts(new InetSocketAddress("192.168.0.96", 8635))
        .addContactPointsWithPorts(new InetSocketAddress("192.168.0.54", 8635))
        .addContactPointsWithPorts(new InetSocketAddress("192.168.0.153", 8635))

        .addContactPointsWithPorts(new InetSocketAddress("124.70.177.38", 38635))
        .addContactPointsWithPorts(new InetSocketAddress("124.70.177.38", 28635))
        .addContactPointsWithPorts(new InetSocketAddress("124.70.177.38", 18635))
        .withReconnectionPolicy(new ConstantReconnectionPolicy(100L))
        .withCredentials(USER, PASSWORD)
        .withoutJMXReporting()
        .build();

    Session session = cluster.connect();

    System.out.println(" ");
    String queryCQL = "SELECT peer,data_center,host_id,rpc_address FROM system.peers ";
    ResultSet rs = session.execute(queryCQL);
    List<Row> dataList = rs.all();
    System.out.println(" ");
    System.out.println(dataList.toString());
    System.out.println(" ");
    System.out.println(" ");

    cluster.close();
    System.out.println("connectToCluster finished");
}

public static void main(String[] args) {
    connectToCluster();
}

```

Solution

Use the private IP address provided by the GeminiDB Cassandra database and change the port to port 8635.

The following figure shows the IP address and port.

```
Cluster cluster = Cluster.builder().addContactPoint(address: "192.168.0.96").withPort(8635).build();
```

4.4.3 How Can I Create and Connect to an ECS?

- To create an ECS, see *Elastic Cloud Server User Guide*.
 - The ECS to be created must be in the same VPC with the GeminiDB Cassandra instance to which it connects.
 - Configure the security group rules to allow the ECS to access to the instance.
- To connect to an ECS, see "Logging in to an ECS" *Getting Started with Elastic Cloud Server User Guide*.

4.4.4 Can I Change the VPC of a GeminiDB Cassandra Instance?

After a GeminiDB Cassandra instance is created, the VPC where the instance resides cannot be changed.

However, you can change a VPC by restoring the full backup of your instance to the VPC you want to use. For details, see [Restoring Data to a New Instance](#).

4.5 Backup and Restoration

4.5.1 How Long Does GeminiDB Cassandra Store Backup Data?

Automated backup data is kept based on the backup retention period you specified. There is no limit for the manual backup retention period. You can delete manually backup files as needed.

4.6 Instance Freezing, Release, Deletion, and Unsubscription

Why Are My GeminiDB Cassandra Instances Released?

If your subscriptions have expired but not been renewed, or you are in arrears due to insufficient balance, your instances enter a grace period. If you do not renew the subscriptions or top up your account after the grace period expires, your instances will enter a retention period and become unavailable. If you still do not renew them or top up your account after the retention period ends, your instances will be released and your data stored will be deleted.

Why Are My GeminiDB Cassandra Instances Frozen?

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

Can I Still Back Up Data If My Instances Are Frozen?

No. If your instances are frozen because your account is in arrears, go to top up your account to unfreeze your instances and then back up instance data.

How Do I Unfreeze My Instances?

If your instances are frozen because your account is in arrears, you can unfreeze them by renewing them or topping up your account. The frozen instances can be renewed, released, or deleted. Yearly/Monthly instances that have expired cannot be unsubscribed from, while those that have not expired can be unsubscribed from.

What Impacts Does Instance Freezing, Unfreezing or Release Have on My Services?

- After an instance is frozen:
 - It cannot be accessed, and your services will be interrupted. For example, if a GeminiDB Cassandra instance is frozen, it cannot be connected.

- No changes can be performed on it if it is a yearly/monthly instance.
- It can be unsubscribed from or deleted manually.
- After it is unfrozen, you can connect to it again.
- Releasing an instance means deleting it. Before the deletion, GeminiDB Cassandra API determines whether to **move the instance to the recycle bin** based on the recycling policy you specified.

How Do I Renew My Instances?

After a yearly/monthly instance expires, you can renew it on the Renewals page. For details, see [Renewal Management](#).

Can My Instances Be Recovered After They Are Released or Unsubscribed From?

If your instance is moved to the recycle bin after being deleted, you can recover it from the recycle bin by referring to [Recycling an Instance](#). If the recycling policy is not enabled, you cannot recover it.

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

How Do I Delete a GeminiDB Cassandra Instance?

- To delete a pay-per-use instance, see [Deleting Instance](#).
- To delete a yearly/monthly instance, see [Unsubscribing from a Yearly/Monthly Instance](#).

A Change History

Release Date	Description
2023-03-27	This issue is the second official release. In Instance Specifications , added the description of instances, each with a vCPUs to memory ratio of 1:4.
2023-02-19	This issue is the first official release.