GeminiDB

Cassandra

Issue 32

Date 2023-08-30





Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2024. All rights reserved.

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

Trademarks and Permissions

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

Notice

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

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

Contents

1 Service Overview	
1.1 What Is GeminiDB Cassandra API?	
1.2 Compatible APIs and Versions	3
1.3 Instance Specifications	3
1.4 Instance Statuses	4
1.5 Database Constraints	5
1.5.1 Basic Design	5
1.5.2 DB Instance Specifications and Performance	8
1.5.3 Database Objects	8
1.5.4 Database Usage	g
1.5.5 Access and Connection Pools	11
1.5.6 Batches	11
1.5.7 Queries	11
2 Billing	13
2.1 Billing Overview	13
2.2 Billing Modes	14
2.2.1 Overview	14
2.2.2 Yearly/Monthly Billing	15
2.2.3 Pay-per-Use Billing	20
2.3 Billing Items	24
2.4 Billing Examples	26
2.5 Billing Mode Changes	29
2.5.1 Overview	29
2.5.2 Pay-per-Use to Yearly/Monthly	30
2.5.3 Yearly/Monthly to Pay-per-Use	32
2.6 Renewing Subscriptions	34
2.6.1 Overview	34
2.6.2 Manually Renewing an Instance	36
2.6.3 Auto-renewing an Instance	39
2.7 Bills	41
2.8 Arrears	45
2.9 Billing Termination	47
2.10 Cost Management	49

2.10.1 Cost Composition	49
2.10.2 Cost Allocation	49
2.10.3 Cost Analysis	50
2.10.4 Cost Optimization	51
2.11 Billing FAQs	51
2.11.1 What Are the Differences Between Yearly/Monthly and Pay-per-Use Billing?	51
2.11.2 Can I Switch Between Yearly/Monthly and Pay-per-Use Billing?	52
3 Getting Started with GeminiDB Cassandra API	53
3.1 Service Overview	53
3.2 Buying an Instance	54
3.3 Instance Connections	62
3.3.1 Connection Methods	62
3.3.2 Connecting to an Instance Through DAS	63
3.3.3 Connecting to an Instance over a Private Network	64
3.3.4 Connecting to an Instance over a Public Network	67
3.3.5 Connecting to an Instance Using Java	71
3.3.6 Connecting to an Instance Using Go	72
3.4 Getting Started with Common Practices	73
4 Working with GeminiDB Cassandra API	75
4.1 Permissions Management	75
4.1.1 Creating a User and Assigning Permissions	75
4.1.2 Creating a Custom Policy	76
4.2 Instance Lifecycle	78
4.2.1 Restarting an Instance	78
4.2.2 Deleting a Pay-per-Use Instance	78
4.2.3 Recycling an Instance	79
4.3 Instance Modifications	80
4.3.1 Upgrading a Minor Version	80
4.3.2 Changing an Instance Name	82
4.3.3 Resetting the Administrator Password	83
4.3.4 Scaling Storage Space	84
4.3.5 Configuring Autoscaling	86
4.3.6 Changing vCPUs and Memory of an Instance	
4.3.7 Adding Nodes	93
4.3.8 Deleting Nodes	96
4.3.9 Managing Tags	96
4.3.10 Updating the OS of an Instance	99
4.4 Connection Management	
4.4.1 Configuring Security Group Rules	
4.4.2 Binding and Unbinding an EIP	
4.4.3 Viewing the IP Address and Port Number	
4.4.4 Changing a Database Port	106

4.4.5 Changing a Security Group	107
4.4.6 Configuring an SSL Connection	107
4.4.7 Downloading the SSL Certificate	109
4.5 Data Management	110
4.5.1 Importing and Exporting Data by Running COPY	110
4.6 Intra-region DR	130
4.6.1 Creating a DR Instance	130
4.6.2 Deleting the DR Relationship	135
4.7 Cross-region Dual-active DR	135
4.7.1 Overview	135
4.7.2 Creating a Dual-Active Relationship	136
4.7.3 Deleting a Dual-active Relationship	138
4.8 Data Backup	138
4.8.1 Overview	138
4.8.2 Managing Automated Backups	140
4.8.3 Configuring an Incremental Backup Policy	145
4.8.4 Setting a Cross-Region Backup Policy	148
4.8.5 Managing Manual Backups	152
4.8.6 Creating a Table-level Backup	154
4.8.7 Setting a Table-level Backup Policy	156
4.8.8 Viewing Incremental Backups	158
4.9 Data Restoration	159
4.9.1 Restoration Methods	159
4.9.2 Restoring Data to a New Instance	160
4.9.3 Restoring a Backup to a Specific Point in Time	161
4.10 Parameter Template Management	163
4.10.1 Creating a Parameter Template	163
4.10.2 Modifying a Parameter Template	164
4.10.3 Viewing Parameter Change History	167
4.10.4 Exporting a Parameter Template	168
4.10.5 Comparing Parameter Templates	169
4.10.6 Replicating a Parameter Template	170
4.10.7 Resetting a Parameter Template	172
4.10.8 Applying a Parameter Template	172
4.10.9 Viewing Application Records of a Parameter Template	172
4.10.10 Modifying a Parameter Template Description	173
4.10.11 Deleting a Parameter Template	173
4.11 Audit on Instance Operations	174
4.11.1 Key Operations Supported by CTS	
4.11.2 Querying Traces	176
4.12 Monitoring and Alarm Configuration	176
4.12.1 GeminiDB Cassandra Metrics	176

4.12.2 Configuring Alarm Rules	183
4.12.3 Viewing Metrics	189
4.12.4 Event Monitoring	190
4.12.4.1 Introduction to Event Monitoring	190
4.12.4.2 Viewing Event Monitoring Data	190
4.12.4.3 Creating an Alarm Rule for Event Monitoring	191
4.12.4.4 Events Supported by Event Monitoring	193
4.13 Log Management	201
4.13.1 Slow Query Logs	201
4.14 Enterprise Project	202
4.14.1 Overview	202
4.14.2 Managing Quotas	202
4.15 Billing Management	204
4.15.1 Renewing Instances	204
4.15.2 Changing the Billing Mode from Pay-per-Use to Yearly/Monthly	205
4.15.3 Changing the Billing Mode from Yearly/Monthly to Pay-per-Use	207
4.15.4 Unsubscribing from a Yearly/Monthly Instance	208
5 Best Practices	211
5.1 Performance of GeminiDB Cassandra and On-Premises Open Source Cassandra Clusters	
5.2 Buying and Connecting to a GeminiDB Cassandra Instance	
6 Performance White Paper	
6.1 Test Method	
6.2 Test Data	
7 FAQs	
7.1 Product Consulting	
7.1.1 What Should I Pay Attention to When Using GeminiDB Cassandra?	
•	234
7.2 Billing	
7.2.1 What Are the Differences Between Yearly/Monthly and Pay-per-use Billing Mode?	
7.2.2 Can I Switch Between Yearly/Monthly and Pay-per-Use Payments?	
7.3 Database Usage	
7.3.1 Why Does the Overall Instance Performance Deteriorate When QPS Increases After the Barberreased?	
7.3.2 What Can I Do if Error "field larger than field limit (131072)" Is Reported During Data Im	port?236
7.3.3 What Should I Pay Attention to When Creating a GeminiDB Cassandra Table?	237
7.3.4 How Do I Detect and Resolve BigKey and HotKey Issues?	241
7.3.5 How Do I Set Up a Materialized View?	
7.3.6 How Do I Use a Secondary Index?	248
7.3.7 How Do I Set Paging Query?	249
7.4 Database Connection	250
7.4.1 What Can I Do If Spark Failed to Connect to Cassandra?	250

A Change History	256
7.7 Instance Freezing, Release, Deletion, and Unsubscription	254
7.6.3 Can I Change the Region of a GeminiDB Cassandra Instance?	254
7.6.2 Can Different AZs Communicate with Each Other?	253
7.6.1 What Is AZ and How Can I Select an AZ?	253
7.6 Regions and AZs	
7.5.1 How Long Does GeminiDB Cassandra Store Backup Data?	253
7.5 Backup and Restoration	253
7.4.4 Can I Change the VPC of a GeminiDB Cassandra Instance?	252
7.4.3 How Can I Create and Connect to an ECS?	252
7.4.2 What Can I Do If an Error Occurs When I Use Java Driver and a Mapped IP Address to Database?	

1 Service Overview

1.1 What Is GeminiDB Cassandra API?

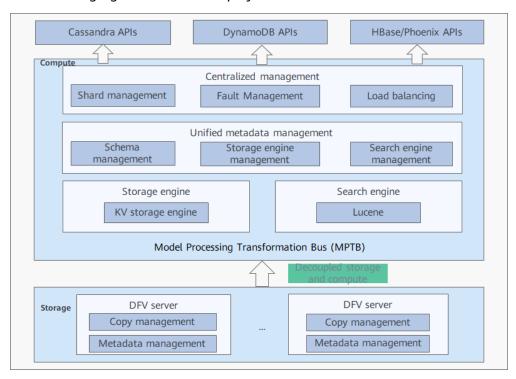
GeminiDB Cassandra API is a cloud-native NoSQL database compatible with Cassandra. It supports Cassandra Query Language (CQL), which gives you SQL-like syntax. GeminiDB Cassandra API improves performance and resolves open-source Cassandra pain points. It is secure, reliable, and compatible with diverse ecosystems.

- 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).
- Compatibility with various ecosystems
 - It is fully compatible with the open-source Cassandra.
 - It supports a SQL-like syntax and offers you a MySQL-like user experience.
 - It is compatible with DynamoDB APIs and allows you to smoothly migrate data from DynamoDB.
- Enhanced capabilities
 - Enhanced index feature makes it easy to handle high concurrency and complex queries.
 - Data recovery capabilities such as second-level flashback and Point-In-Time Recovery (PITR) ensure high data reliability.
- Superior performance
 - Wide table model enables it to store petabytes of structured and semistructured data.
 - The read-only performance is improved by several times and read/write performance is doubled.

- Free of pain points of the open-source Cassandra
 - Data consistency and garbage collection (GC) issues are resolved.
 - Storage scaling can be completed in seconds without affecting services.
 - Compute nodes can be added in minutes. There is only a few seconds of jitter.

Architecture

The following figure shows the deployment architecture of GeminiDB Cassandra.



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 ecommerce 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

CPU Type	Flavor	vCPUs	Memory (GB)	Min. Storage Space (GB)	Max. Storage Space
x86	geminidb.cassandra.la rge.4	2	8	100	12,000
	geminidb.cassandra.xl arge.4	4	16	100	24,000

CPU Type	Flavor	vCPUs	Memory (GB)	Min. Storage Space (GB)	Max. Storage Space
	geminidb.cassandra.xl arge.8	4	32	100	24,000
	geminidb.cassandra.2x large.4	8	32	100	48,000
	geminidb.cassandra.2x large.8	8	64	100	48,000
	geminidb.cassandra.4x large.4	16	64	100	96,000
	geminidb.cassandra.4x large.8	16	128	100	96,000
	geminidb.cassandra.6x large.8	24	192	100	144,000
	geminidb.cassandra.8x large.4	32	128	100	192,000
	geminidb.cassandra.8x large.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.
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.

Status	Description
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.
Configuring SSL	SSL is being enabled or disabled.
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 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

- 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 **Z00_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 Database Usage

Mandatory Constraints

General rules

• If the size of a request or any configuration value in it exceeds the preset alarm threshold, the client receives a warning and Cloud Eye generates an alarm. For some requests, events are also generated.

• If the size of a request or any configuration value in it exceeds the preset failure threshold, the request fails and an event is generated.

Specific constraints

- Data volume of a single table row:
 - When a database reads data, it will combine multiple writes of the same primary key into one row. When the amount of data in a single row exceeds the preset threshold, an alarm and event are generated.
- A single partition:
 - The database background scanning task periodically collects partitions, rows in a partition, and total size of all rows in a partition. When the amount of data in a single row exceeds the preset threshold, an alarm and event are generated.
- Number of elements of the collection type:
 - When your database reads data (of the types like Map, List, Set, and Tuple) in the collection column, the database will count all elements in the column. An alarm and event is generated when the number of elements exceeds the preset threshold.
- Amount of data returned by a query:
 - Before a response is returned to the client, the database checks the amount of data in the response. If the amount exceeds the preset threshold, an alarm is generated or the request fails.
- Number of tombstones returned by a query:
 - After a query request is submitted, the database checks the number of tombstones scanned. When the number exceeds the preset threshold, an alarm is generated or the request fails.

Table 1-6 Thresholds for GeminiDB Cassandra constraints

Constraint	Alarm Threshold	Failure Threshold
Amount of data in a single row	100Kb	-
A single partition	Rows: 100,000Amount of data: 100 MB	-
Number of elements of the collection type	500	-
Amount of data returned by a query	2Mb	100Mb
Number of tombstones returned by a query	1,000	100,000

Optional Constraints

- Do not use ALLOW FILTERING:
 - A warning is returned when the ALLOW FILTERING statement is executed.
- Use a proper limit value for a RANGE query:
 - Using a proper limit for prefetching can accelerate RANGE queries. You
 can check whether a limit value is proper based on the average limit
 values of all monitoring items.

1.5.5 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.

1.5.6 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.

1.5.7 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.

Not Using ALLOW FILTERING

If a query statement does not specify all primary keys and contains **ALLOW FILTERING**, the query will scan and filter the entire table. A table with a large data volume may cause the query to time out. **ALLOW FILTERING** is forbidden in later kernel versions.

◯ NOTE

Query timeout and excessive resource usage issues that occur when **ALLOW FILTERING** is used are not within commitments on SLAs

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=";

Ⅲ NOTE

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

$\mathbf{2}$ Billing

2.1 Billing Overview

In this document, you will learn about how instances are billed, how you can renew subscriptions and manage costs, and what happens if your account goes into arrears.

Billing Modes

There are yearly/monthly and pay-per-use billing modes. Each one has different advantages and disadvantages.

- Yearly/Monthly: You pay upfront for the amount of time you expect to use the instance for. You will need to make sure you have a top-up account with a sufficient balance or have a valid payment method configured first.
- Pay-per-use: You can start using the GeminiDB instance first and then pay as you go.

For details about the two billing modes, see Overview.

You can also change the billing mode later if it no longer meets your needs. For details, see **Overview**.

Billing Items

You will be billed for instance specifications, storage space, backup space, and EIP bandwidths. For details about the billing factors and formulas for each billed item, see **Billing Items**.

For more information about billing samples and the billing for each item, see **Billing Examples**.

Renewing Subscriptions

If you want to continue using an instance after it expires, you need to renew the instance subscription within the specified period. Otherwise, resources, such as compute and storage, will be automatically released, and data may be lost

You can renew your subscription manually or automatically. For details, see **Overview**.

Viewing Bills

You can choose **Billing & Costs** > **Bills** to check the instance transactions and bills. For details, see **Bills**.

Arrears

If there is not a sufficient account balance to pay for your bill and there is no other payment method configured, your account will go into arrears. If you want to continue using your cloud services, you will need to top up your account in a timely manner. For details, see **Arrears**.

• Stopping Billing

If you no longer need to use your GeminiDB Cassandra instance, you can unsubscribe from or delete it to stop the billing. For details, see **Billing Termination**.

Managing Costs

GeminiDB Cassandra costs include resource costs and O&M costs. You can allocate, analyze, and optimize GeminiDB costs to save more money. For details, see **Cost Management**.

2.2 Billing Modes

2.2.1 Overview

There are yearly/monthly and pay-per-use billing modes. Each one has different advantages and disadvantages.

- Yearly/Monthly is a prepaid billing mode. You pay in advance for a subscription term, and in exchange, you get a discounted rate. The longer the subscription term, the bigger the discount. Yearly/Monthly billing is a good option for long-term, stable services.
- Pay-per-use is a postpaid billing mode. You pay as you go and just pay for
 what you use. The instance usage is calculated by the second but billed every
 hour. Pay-per-use billing is a good option for scenarios where there are
 sudden traffic bursts, such as e-commerce promotions.

Table 2-1 lists differences between the two billing modes.

Table 2-1 Differences between billing modes

Billing Mode	Yearly/Monthly	Pay-per-use
Payment	Prepaid Billed by the subscription term you purchase	Postpaid Billed for what you use
Billing Method	Billed by the subscription term you purchase	Calculated by the second but billed every hour
Billed Items	Instance specifications (vCPUs and memory), storage space, backup space, and EIPs	Instance specifications (vCPUs and memory), storage space, backup space, and EIPs

Changing the Billing Mode	Yearly/Monthly can be changed to pay-per-use. The change takes effect only after the yearly/monthly subscription expires. For details, see Yearly/Monthly to Pay-per-Use.	Pay-per-use can be changed to yearly/monthly. For details, see Pay-per-Use to Yearly/Monthly.
Changing the Specificati ons	Supported	Supported
Applicatio n Scenarios	Recommended for resources expected to be in use long term. A cost-effective option for scenarios where the resource usage duration is predictable.	Recommended when the resource demands are likely to fluctuate and you want more flexibility.

2.2.2 Yearly/Monthly Billing

If you expect to use resources for a longer period, you can save money by selecting yearly/monthly billing. This section describes billing rules of yearly/monthly GeminiDB Cassandra resources.

Application Scenarios

If you want to ensure resource stability over a certain period of time, yearly/monthly billing is a good choice for the following types of workloads:

- Long-term workloads with stable resource requirements, such as official websites, online malls, and blogs.
- Long-term projects, such as scientific research projects and large-scale events.
- Workloads with predictable traffic bursts, for example, e-commerce promotions or festivals.
- Workloads with high data security requirements.

Billed Items

You are billed for the following items on a yearly/monthly basis.

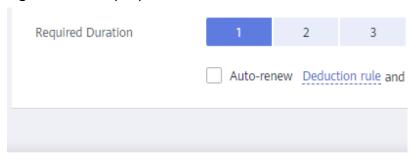
Table 2-2 Items billed on a yearly/monthly basis

Billing Item	Description
Instance specificatio ns	Instance specifications, including vCPUs and memory.

Billing Item	Description
Storage space	If the actual storage usage exceeds your purchased storage, you will be billed for additional storage on a pay-per-use basis.
Backup space	GeminiDB Cassandra provides backup storage up to 100% of your provisioned database storage at no additional charge.
	After the free backup space is used up, charges are applied based on the backup space pricing details. Pricing is listed on a per-hour basis, but bills are calculated based on the actual usage duration.
(Optional) Public network bandwidth	GeminiDB Cassandra instances are accessible from public networks, and you are billed for the generated public network traffic, but not for private network traffic.

If you want to purchase a 3-node (specifications of each node: 2 vCPUs | 8 GB) GeminiDB Cassandra instance with 100 GB of storage space. At the bottom of the instance buying page, price details (excluding the backup space fee) will be displayed.

Figure 2-1 Example price



Price \$433.06 USD ②

The price includes:

- Selected specifications for your instance
- Storage space

Ⅲ NOTE

The backup space fee is not included. For details about the backup price, see **Product Pricing Details**.

Backup Storage Space



Billed Usage Period

A yearly/monthly GeminiDB Cassandra instance is billed for the purchased duration (UTC+8). The billing starts when you activated or renewed the subscription, and ends at 23:59:59 of the expiry date.

For example, if you purchased a one-month GeminiDB Cassandra instance on March 08, 2023, 15:50:04, the billed usage period is from March 08, 2023, 15:50:04 to April 08, 2023, 23:59:59.

Billing Examples

Suppose you purchased a one-month GeminiDB Cassandra instance (instance specifications: 2 vCPUs | 8 GB; nodes: 3; storage: 100 GB; backup space: 110 GB (100 GB for free)) on March 08, 2023, 15:50:04, and renewed the subscription for one more month before the initial subscription expired. That would include two usage periods:

- March 08, 2023, 15:50:04 to April 08, 2023, 23:59:59
- April 08, 2023, 23:59:59 to May 08, 2023, 23:59:59
 - From April 08, 2023, 23:59:59 to May 01, 2023, 23:59:59, 50 GB of free backup space was used.
 - From May 01, 2023, 23:59:59 to May 08, 2023, 23:59:59, another 10 GB of backup space was used, which was billed for 168 hours.

You will be billed for both usage periods. GeminiDB Cassandra resources are billed individually as follows:

Table 2-3 Formulas for billing yearly/monthly GeminiDB Cassandra resources

Resource	Formula	Unit Price
Instance specifications (including vCPUs and memory)	Unit price of the instance specifications x Required duration x Number of nodes	For details about the unit price, see Cluster CPU/Memory on Product Pricing Details
Storage space	Storage space unit price x Required duration x Storage space (GB)	For details about the unit price, see Storage Space on Product Pricing Details .

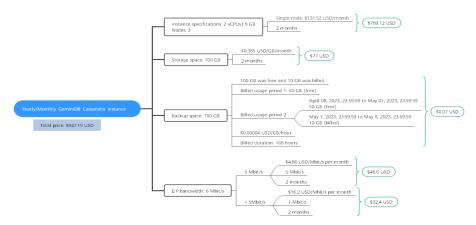
Resource	Formula	Unit Price
Backup space	Backup space unit price x Required duration x (Backup space – Storage space) (GB)	For details about the unit price, see Backup Storage Space on Product Pricing Details .
	NOTE The billed duration refers to the length of time the billed backup space was used for.	
Public network bandwidth	Billed by fixed bandwidth	For details, see Product Pricing Details .

Figure 2-2 shows how the total price is calculated.

NOTICE

Prices in the figure are just examples. Actual prices are subject to **Product Pricing Details**.

Figure 2-2 Total price for a yearly/monthly GeminiDB Cassandra instance



Price Change After Specification Change

If the specifications of a yearly/monthly GeminiDB Cassandra instance no longer meet your needs, you can change the specifications on the console. The system will recalculate the price and either bill or refund you the difference.

- If you upgrade your GeminiDB Cassandra instance specifications, you need to pay the difference in price.
- If you downgrade your GeminiDB Cassandra instance specifications, Huawei Cloud will refund you the difference.

You are not advised to downgrade your GeminiDB Cassandra instance to a lower specification because the instance performance may be affected. Suppose you

purchased a yearly/monthly GeminiDB Cassandra instance (2 vCPUs | 8 GB and 3 nodes) on April 08, 2023 and upgraded the instance specifications to 4 vCPUs | 16 GB and 3 nodes on April 18, 2023. The price for the original specifications was \$587.06 USD/month, and that for the new specifications was \$981.62 USD/month. The price difference will be calculated as follows:

Price difference for the specification upgrade = Price for the new specifications × Remaining period - Price for the original specifications × Remaining period

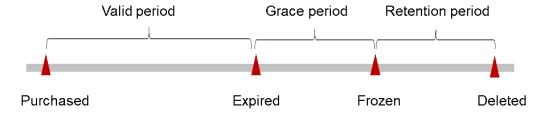
The remaining period in the formula is the remaining days of each calendar month divided by the maximum number of days in each calendar month. In this example, Remaining period = 12 (Remaining days in April)/30 (Maximum number of days in April) + 8 (Remaining days in May)/31 (Maximum number of days in May) = 0.6581. Cost of upgrade = $$981.62 \text{ USD} \times 0.6581 - $587.06 \text{ USD} \times 0.6581 = 259.66 USD .

For more details, see **Pricing of a Changed Specification**.

Impact of Expiration

Figure 2-3 shows the statuses a yearly/monthly GeminiDB Cassandra instance can go through throughout its lifecycle. After a GeminiDB Cassandra instance is purchased, it enters the valid period and runs normally during this period. If the instance is not renewed after it expires, before being deleted, it first enters a grace period and then a retention period.

Figure 2-3 Lifecycle of a yearly/monthly GeminiDB Cassandra instance



Expiration Reminder

The system will send you a reminder (by email, SMS, or in-app message) 7 days before a yearly/monthly GeminiDB Cassandra instance expires to remind you to renew the subscription.

Impact of Expiration

If your yearly/monthly GeminiDB Cassandra instance is not renewed after it expires, it changes to the **Expired** state and enters a grace period. During the grace period, you can access the GeminiDB Cassandra instance but cannot:

- Change instance specifications.
- Change the billing mode from yearly/monthly to pay-per-use.
- Unsubscribe from it.

If the yearly/monthly GeminiDB Cassandra instance is not renewed after the grace period ends, its status turns to **Frozen** and it enters a retention period. You cannot

perform any operations on the GeminiDB Cassandra instance while it is in the retention period.

If the yearly/monthly GeminiDB Cassandra instance is not renewed by the time the retention period ends, it will be released and data cannot be restored.

• For details about renewals, see Overview.

2.2.3 Pay-per-Use Billing

Pay-per-use billing means you pay nothing up front and are not tied into any contract or commitment. This section describes billing rules of pay-per-use GeminiDB Cassandra instances.

Application Scenarios

Pay-per-use billing is good for short-term, bursty, or unpredictable workloads that cannot tolerate any interruptions, such as applications for e-commerce flash sales, temporary testing, and scientific computing.

Billing Items

You are billed for the following items on a pay-per-use basis.

Table 2-4 Items billed on a pay-per-use basis

Table 2. Freeins state on a pay per use susis		
Billing Item	Description	
Instance specificatio ns	Instance specifications, including vCPUs and memory.	
Storage space	Instance storage space, which is billed hourly on a pay-per-use basis.	
Backup space	GeminiDB Cassandra provides backup storage up to 100% of your provisioned database storage at no additional charge.	
	After the free backup space is used up, charges are applied based on the backup space pricing details. Pricing is listed on a per-hour basis, but bills are calculated based on the actual usage duration.	
(Optional) Public network bandwidth	GeminiDB Cassandra instances are accessible from public networks, and you are billed for the generated public network traffic, but not for private network traffic.	

If you want to purchase a 3-node (specifications of each node: 2 vCPUs | 8 GB) GeminiDB Cassandra instance with 500 GB of storage space. At the bottom of the instance buying page, price details (excluding the backup space fee) will be displayed.

Figure 2-4 Example price

Price \$1.17 USD/hour ?

The price includes:

- Instance specifications (including vCPUs and memory)
- Selected storage space

□ NOTE

The backup space fee is not included. For details about the backup price, see **Product Pricing Details**.

Backup Storage Space



Billed Usage Period

Pay-per-use GeminiDB Cassandra instance usage is calculated by the second and billed every hour. The billing starts when ECS instance is created and ends when the instance is deleted.

□ NOTE

It takes a certain time to create an instance. The billing starts from the time when the instance is successfully created. You can view the two time points on the **Basic Information** page. You can view the time when the instance is created beside the **Created** field.

For example, if you purchased a pay-per-use GeminiDB Cassandra instance at 8:45:30 and deleted it at 8:55:30, you are billed for the 600 seconds from 8:45:30 to 8:55:30.

Billing Examples

Suppose you purchased a pay-per-use instance on April 18, 2023, 9:59:30, and deleted it on April 18, 2023, 10:45:46. Two usage periods will be billed:

- Usage of 30 seconds from 9:59:30 to 10:00:00
- Usage of 2,746 seconds from 10:00:00 to 10:45:46
 - The free backup space is used from 10:00:00 to 10:45:00.
 - Ten GB of billing backup space is used from 10:45:00 to 10:45:46 and the billed duration is 46 seconds.

The price displayed in the pricing details is per hour, so you need to divide it by 3,600 to obtain the price for each second and then multiply the per-second price by the total number of seconds. GeminiDB Cassandra instances are billed individually.

Table 2-5 Formulas for billing pay-per-use GeminiDB Cassandra resources

Resource	Formula	Unit Price
Compute resources (including vCPUs and nodes)	Unit price of instance specifications x Required duration	For details about the unit price, see Cluster CPU/Memory on Product Pricing Details
Storage space	Storage space unit price x Purchase duration	For details about the unit price, see Storage Space on Product Pricing Details.
Backup space	Backup space unit price x Required duration x (Backup space – Storage space) (GB) NOTE The billed duration refers to the length of time the billed backup space was used for.	For details about the unit price, see Backup Storage Space on Product Pricing Details.
Public network traffic	Tiered billing by fixed bandwidth • 0 Mbit/s to 5 Mbit/s (included): billed at a fixed unit price per Mbit/s • Greater than 5 Mbit/s: billed at a different price per Mbit/s	For details, see Bandwidth Price on Product Pricing Details page or Product Pricing Details .

Figure 2-5 shows how the total price is calculated.

NOTICE

Prices in the figure are just examples. Actual prices are subject to **Product Pricing Details**.

For pay-per-use billing, decimal numerals on the price calculator are rounded off and are accurate to two decimal places. If the fee is less than \$0.01 USD (after rounding off), \$0.01 USD will be displayed.

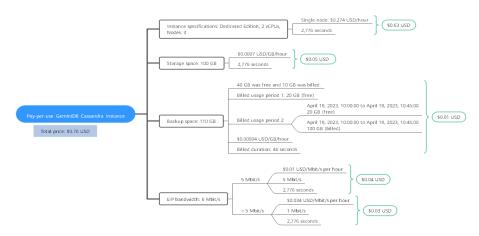


Figure 2-5 Total price for a pay-per-use GeminiDB Cassandra instance

Price Change After Specification Change

If you change the specifications of a pay-per-use GeminiDB Cassandra instance, the original order will become invalid and a new order will be placed. You will be billed based on the new specifications.

If you change instance specifications within a given hour, multiple records will be generated. Different records record the billing for different specifications.

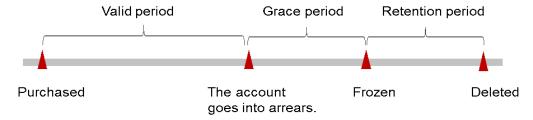
For example, if you purchased a pay-per-use instance (2 vCPUs | 8 GB) at 9:00:00 and changed the instance specifications to 4 vCPUs | 16 GB at 9:30:00, the following items will be billed:

- Specifications 2 vCPUs | 8 GB usage from 9:00:00 to 9:30:00
- Specifications 4 vCPUs | 16 GB usage from 9:30:00 to 10:00:00

Impact of Arrears

Figure 2-6 shows the statuses a pay-per-use GeminiDB Cassandra instance can go through throughout its lifecycle. After a GeminiDB Cassandra instance is purchased, it enters the valid period and runs normally during this period. If your account goes into arrears, the instance enters a grace period and then a retention period.

Figure 2-6 Lifecycle of a pay-per-use GeminiDB Cassandra instance



Arrears Reminder

The system will bill you for pay-per-use resources after each billing cycle ends. If your account goes into arrears, we will notify you by email, SMS, or in-app message.

Impacts of Arrears

When your account is in arrears due to automatic fee deduction for pay-per-use GeminiDB Cassandra instances, the account status turns to arrears. In arrears, the pay-per-use instance continues rendering service but the instance enters the grace period. You are still responsible for expenditures generated during the grace period. You can view the charges on the **Billing Center** > **Overview** page and pay any past due balance as needed.

If you do not bring your account balance current before the grace period expires, the GeminiDB Cassandra instance status turns to **Frozen** and it enters a retention period. You cannot perform any operations on a pay-per-use GeminiDB Cassandra instance in the **Frozen** status.

If you do not bring your account balance current before the retention period ends, your instance will be released, and data cannot be restored.

□ NOTE

- During the retention period, you cannot access or use your instance but the data stored in it can be retained. The retention period for Huawei Cloud International website is 15 days.
- During the grace period, you can access and use only some resources of your instance.
 The grace period for Huawei Cloud International website is 15 days.
- For details about top-up, see Topping Up an Account.

2.3 Billing Items

Billing

You will be billed for instance specifications, storage space, backup space, and EIP bandwidths. For details, see **Table 2-6**.

□ NOTE

The billed items marked with asterisks (*) are mandatory.

Table 2-6 Billing Items of a GeminiDB Cassandra Instance

Billing Item	Description	Billing Mode	Formula
* Specific ations	Billed by instance specifications, including vCPUs and memory. Computing and storage capabilities vary by the number of vCPUs and memory size.	Yearly/ Monthly and pay- per-use	Unit price x Required duration For details about the unit price, see Cluster CPU/Memory on Product Pricing Details.

Billing Item	Description	Billing Mode	Formula
* Storage space	Billed based on unified standards.	Yearly/ Monthly and pay- per-use	Unit price x Storage space x Required duration For details about the unit price, see Storage Space on Product Pricing Details.
Backup space	Billed based on unified standards.	Pay-per- use	Unit price x Billed backup space x Required duration For details about the unit price, see Backup Storage Space on Product Pricing Details. NOTE The billed duration refers to the length of time the billed backup space was used for.
Public network traffic	An EIP is required if a GeminiDB Cassandra instance needs to access the Internet. Billed by bandwidth, traffic, and the EIP reservation price. EIP for a yearly/monthly GeminiDB Cassandra instance: billed by bandwidth. EIP for a pay-per-use GeminiDB Cassandra instance: billed by bandwidth, traffic, or shared bandwidth, You are also charged for IP reservation if you do not bind the EIP to any instance.	Yearly/ Monthly and pay- per-use You can purchase a bandwidt h add-on package or a shared traffic package.	Tiered pricing based on fixed bandwidth. O Mbit/s to 5 Mbit/s (included): billed at a fixed unit price per Mbit/s. Greater than 5 Mbit/s: billed at a different price per Mbit/s. For details about the unit price, see Bandwidth Price on Product Pricing Details or Product Pricing Details.

Billing Examples

Suppose you purchased a one-month GeminiDB Cassandra instance (instance specifications: 2 vCPUs | 8 GB; nodes: 3; storage: 100 GB; backup space: 110 GB (100 GB for free)) on March 08, 2023, 15:50:04, and renewed the subscription for one more month before the initial subscription expired. That would include two usage periods:

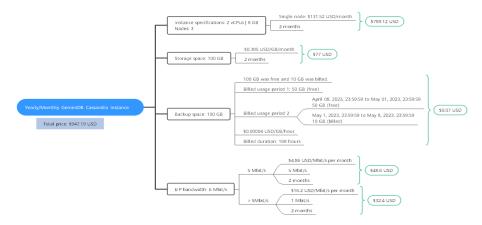
- March 08, 2023, 15:50:04 to April 08, 2023, 23:59:59
- April 08, 2023, 23:59:59 to May 08, 2023, 23:59:59
 - From April 08, 2023, 23:59:59 to May 01, 2023, 23:59:59, 50 GB of free backup space was used.
 - From May 01, 2023, 23:59:59 to May 08, 2023, 23:59:59, another 10 GB of backup space was used, which was billed for 168 hours.

Figure 2-7 shows how the total price is calculated.

NOTICE

Prices in the figure are only for reference. For details, see **Product Pricing Details**.

Figure 2-7 Total price for a yearly/monthly GeminiDB Cassandra instance



For more billing examples of a pay-per-use GeminiDB Cassandra instance, see **Billing Examples**.

2.4 Billing Examples

Billing Scenario

A user purchased a pay-per-use GeminiDB Cassandra instance at 15:30:00 on March 18, 2023. The instance configuration is as follows:

Specifications: 2 vCPUs | 8 GB

• Nodes: 3

Public network bandwidth: 6 Mbit/s

After a period of time, the user found that the current GeminiDB Cassandra instance specifications no longer met service requirements and updated the specifications to 4 vCPUs | 16 GB at 09:00:00 on March 20, 2023. Since the user wanted to use the instance long term, the user then changed the instance to yearly/monthly billing with a one-month duration at 10:30:00 on the same day. So how much will the user be billed for this GeminiDB Cassandra instance in March and April?

Billing Analysis

The total price of this GeminiDB Cassandra instance involves both pay-per-use and yearly/monthly usage:

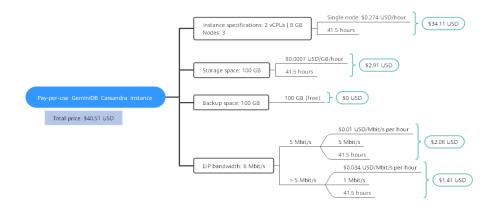
- Pay-per-use usage: March 18, 2023, 15:30:00 to March 20, 2023, 10:30:00
 - March 18, 2023, 15:30:00 to March 20, 2023, 9:00:00
 - Instance specifications: 2 vCPUs | 8 GB
 - Nodes: 3
 - Used storage space: 100 GB
 - Used backup space: 100 GB
 - Public network bandwidth: 6 Mbit/s
 - March 20, 2023, 9:00:00 to March 20, 2023, 10:30:00
 - Instance specifications: 4 vCPUs | 16 GB
 - Nodes: 3
 - Used storage space: 200 GB
 - Used backup space: 210 GB (billed on a pay-per-use basis from March 20, 2023, 10:00:00 to March 20, 2023, 10:30:00)
 - Public network bandwidth: 6 Mbit/s
- Yearly/Monthly: March 20, 2023, 10:30:00 to April 20, 2023, 23:59:59
 - Instance specifications: 4 vCPUs | 16 GB
 - Nodes: 3
 - Used storage space: 200 GB
 - Used backup space: 300 GB (billed on a pay-per-use basis from April 10, 2023, 23:59:59 to April 20, 2023, 23:59:59)
 - Public network bandwidth: 6 Mbit/s
 - Billed duration: one month

NOTICE

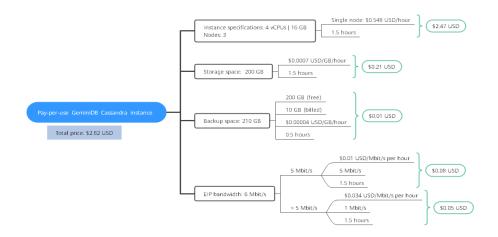
Unit prices in this example are used for reference only. The prices shown here are only estimates. As unit prices change from time to time, the prices shown here will differ from actual prices. For details, see the data released on the Huawei Cloud official website.

Pay-per-use

From March 18, 2023, 15:30:00 to March 20, 2023, 09:00:00, a GeminiDB Cassandra instance with specifications 2 vCPUs | 8 GB was used for 41.5 hours, so the price would be calculated as follows.

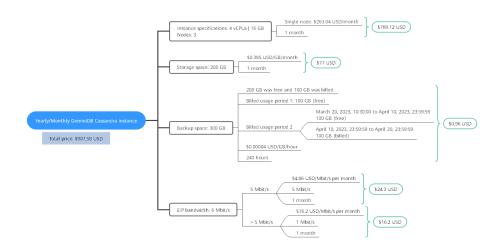


From March 20, 2023, 09:00:00 to March 20, 2023, 10:30:00, a GeminiDB Cassandra instance with specifications 4 vCPUs | 16 GB was used for 1.5 hours, so the price would be calculated as follows.



Yearly/Monthly

From March 20, 2023, 10:30:00 to April 20, 2023, 23:59:59, a GeminiDB Cassandra instance purchased using yearly/monthly billing was used for one month, so the price would be calculated as follows.



From March to April, the total price of this GeminiDB Cassandra instance is \$950.91 USD (40.51 + 2.82 + 907.58).

2.5 Billing Mode Changes

2.5.1 Overview

After purchasing a GeminiDB Cassandra instance, you can change the billing mode if it no longer meets your needs. **Table 2-7** lists changeable billing items of the GeminiDB Cassandra instance.

Table 2-7 Changeable billing items of GeminiDB Cassandra instances

Billing Item	Change Description	Reference
Instance specification s (vCPUs and nodes)	Changing the billing mode of a GeminiDB Cassandra instance includes the changes to compute resources (vCPUs and nodes). Change from pay-per-use to yearly/monthly to enjoy lower prices. Change from yearly/monthly to pay-per-use to use the GeminiDB Cassandra instance more flexibly. NOTE Such a change takes effect only after the yearly/monthly subscription ends.	 Pay-per-Use to Yearly/Monthly Yearly/Monthly to Pay-per-Use
EIP	 A yearly/monthly EIP can be changed to a pay-per-use EIP billed by bandwidth after the yearly/monthly subscription ends. A pay-per-use EIP billed by bandwidth can be changed to a yearly/monthly EIP. Pay-per-use EIPs billed by bandwidth can be changed to pay-per-use EIPs billed by traffic, and pay-per-use EIPs billed by traffic can be changed to pay-per-use EIPs billed by bandwidth. For details, see Figure 2-8. 	 Pay-per-Use to Yearly/Monthly Yearly/Monthly to Pay-per-Use



Traffic

Figure 2-8 EIP billing mode change

1: The change takes effect immediately.

Bandwidth

- 2: The change takes effect only after the yearly/monthly subscription period expires.
- x: The billing mode cannot be changed.

2.5.2 Pay-per-Use to Yearly/Monthly

Billed by

If you have a pay-per-use GeminiDB Cassandra instance that you expect to use for a long time, you can change it to yearly/monthly billing to reduce costs. Doing so will create an order. After you pay for the order, yearly/monthly billing will be applied immediately.

1

Traffic

Suppose you bought a pay-per-use GeminiDB Cassandra instance at 15:29:16 on April 18, 2023 and changed it to yearly/monthly billing at 16:30:30 on the same day. After you paid for the order, yearly/monthly billing was applied immediately. On the **Billing Center** > **Billing** page, three line items were generated.

- Pay-per-use expenditures for 15:29:16 to 16:00:00 on April 18, 2023
- Pay-per-use expenditures for 16:00:00 to 16:30:30 on April 18, 2023
- Yearly//monthly expenditure generated on April 18, 2023, 16:30:30

Constraints

Resources such as EIPs that are used by an instance may not support the change with this instance. For details about their billing mode change rules and handling methods, see **Table 2-8**.

Table 2-8 EIP billing mode change rules

Resourc e	Billing Mode	Billed By	Band width Type	Change to Yearly/ Monthly Billing with GeminiDB Cassandra Instance	Handling Measure
EIP	Pay- per-use	Bandwid th	Dedica ted	Supported	Change the EIP to yearly/monthly billing on the EIP console. For details, see Changing EIP Billing Mode.
EIP	Pay- per-use	Traffic	Dedica ted	Not supported	An EIP that is billed by traffic on a pay-per-use basis cannot be directly changed to be billed on a yearly/monthly basis. To change this: 1. Change the EIP to be billed by bandwidth on a pay-per-use basis.
					 Change the EIP to be billed on a yearly/monthly basis. For details, see Changing EIP Billing Mode.

Prerequisites

- The billing mode of the instance is pay-per-use.
- The instance status is **Available**.

Procedure

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

Figure 2-9 Changing pay-per-use to yearly/monthly



Alternatively, click the instance name to go to the **Basic Information** page. In the **Billing Information** area, click **Change to Yearly/Monthly** in the **Billing Mode** field.

Figure 2-10 Changing from pay-per-use to yearly/monthly

Billing Information

Billing Mode Pay-per-use Change to Yearly/Monthly

Created Jun 17, 2022 11:33:24 GMT+08:00

The billing mode of multiple instances can be changed in batches. Perform the following steps:

- 1. Select the instances whose billing mode you want to change.
- 2. Click Change to Yearly/Monthly above the instance list.
- **Step 4** On the displayed page, specify a subscription duration in month. The minimum duration is one month.

If you do not need to modify your settings, click Pay to go to the payment page.

- **Step 5** Select a payment method and click **Confirm**.
- **Step 6** View results on the **Instances** page.

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

----End

2.5.3 Yearly/Monthly to Pay-per-Use

After creating a yearly/monthly GeminiDB Cassandra instance, you can change it to pay-per-use for more flexibility, and you can recoup part of what you paid for the subscription.

Suppose you bought a yearly/monthly GeminiDB Cassandra instance at 15:29:16 on April 18, 2023 and changed it to pay-per-use billing at 16:30:00 on May 18, 2023. On the **Billing Center** > **Billing** page, bills information is generated as follows:

- Yearly/Monthly expenditures for 15:29:16 on April 18 to 23:59:59 on May 18, 2023
- Pay-per-use expenditures for 23:59:59 on May 18, 2023 to the end time of pay-per-use billing. A bill was generated every hour.

Ⅲ NOTE

The pay-per-use billing mode will take effect only after the yearly/monthly subscription has expired. Auto-renewal will not be in effect.

Constraints

Resources such as EIPs that are used by an instance may not support the change with this instance. For details about their billing mode change rules and handling methods, see **Table 2-9**.

Table 2-9 EIP billing mode change rules

Resour ce	Billing Mode	Billed By	Bandwi dth Type	Change to Pay-Per-Use Billing with GeminiDB Cassandra Instance	Handling Measure
EIP	Yearly/ Monthl y	Bandwi dth	Dedicat ed	Not supported	Change the EIP to yearly/monthly billing on the EIP console. For details, see Changing EIP Billing Mode.
EIP	Yearly/ Monthl y	Traffic	Dedicat ed	Not supported	An EIP billed on a yearly/monthly basis cannot be directly changed to be billed by traffic on a pay-per-use basis. To change this: 1. Change the EIP to be billed by bandwidth on a pay-per-use basis. 2. Change the EIP to be billed by traffic on a pay-per-use basis. For details, see Changing EIP Billing Mode.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** On the **Instances** page, locate the instance whose billing mode you want to change and click **More** > **Change to Pay-per-Use** in the **Operation** column.

Name/ID D8 Instance Type Compatible API Status Billing Mode Operation

Tougle (SEA)

Additionable API Status Billing Mode Operation

Tougle (SEA)

Additionable API Status Billing Mode Operation

Cluster Reds 5.0 Available Pay-per-use Create Backup Restart Delete

Castandra 3.11 Available Pay-per-use Log in Change to Yearly/Monthly More

Tougle (SEA)

TOURS AND AVAILABLE PROPOSED 1979/96-198

Cluster Castandra 3.11 Available Pay-per-use Log in Change to Yearly/Monthly More

TOURS AVAILABLE PROPOSED 1979/96-198

Cluster Castandra 3.11 Available Perpending restart

TOURS AVAILABLE PROPOSED 1979/96-198

Considerable Proposed

Figure 2-11 Changing from yearly/monthly to pay-per-use

□ NOTE

The billing mode of multiple pay-per-use instances can be changed in batches. Perform the following steps:

- 1. Select the instances whose billing mode you want to change.
- 2. Click More > Change to Pay-per-Use in the Operation column
- **Step 4** 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.

 Auto renewal will be disabled after the billing mode of your instances change to pay-per-use. Exercise caution when performing this operation.
- **Step 5** 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 subscription expires.
- **Step 6** 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 Payper-Use**.
- **Step 7** In the displayed dialog box, click **Yes**.

----End

2.6 Renewing Subscriptions

2.6.1 Overview

When to Renew Subscriptions

If a yearly/monthly instance is about to expire but you want to continue using it, you need to renew the instance subscription within a specified period, or resources, such as vCPUs and memory, will be automatically released, and data will be lost and cannot be restored.

Only yearly/monthly instance subscriptions can be renewed. If you use pay-per-use instances, just ensure that your account has a valid payment method configured or a top-up account with a sufficient balance.

If you renew the instance before it expires, resources will be retained and you can continue using the instance. For details about statuses after instances have expired and the associated impacts, see **Impact of Expiration**.

How to Renew Subscriptions

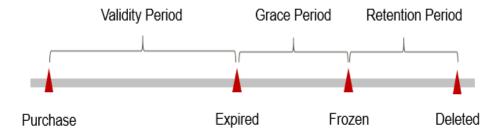
You can renew a yearly/monthly GeminiDB Cassandra instance manually or automatically.

Table 2-10 Renewing a yearly/monthly instance

Method	Description
Manually Renewing an Instance	You can renew a yearly/monthly instance anytime on the console before it is automatically deleted.
Auto-renewing an Instance	You can enable auto-renewal to automatically renew the instance before it expires. This prevents resources from being deleted in case you forget to renew a subscription.

You can select a method to renew a yearly/monthly instance based on the phase the instance is currently in.

Figure 2-12 Selecting a renewal method based on the instance's current phase



- An instance is in the **Provisioned** state after it is provisioned.
- When an instance subscription expires, the status will change from **Provisioned** to **Expired**.
- If an expired instance is not renewed, it enters a grace period. If it is not renewed by the time the grace period expires, the instance will be frozen and enter a retention period.
- If you do not renew the subscription before the retention period expires, your resources will be automatically deleted.

□ NOTE

- During the retention period, you cannot access or use your instance but the data stored in it can be retained. The retention period for Huawei Cloud International website is 15 days.
- During the grace period, you can access and use only some resources of your instance. The grace period for Huawei Cloud International website is 15 days.

You can enable auto-renewal any time before an instance expires. By default, the system will make the first attempt to charge your account for the renewal at 03:00, seven days before the expiry date. If this attempt fails, it will make another

attempt at 03:00 every day until the subscription is renewed or expired. You can change the auto-payment date for renewal as required.

2.6.2 Manually Renewing an Instance

You can renew a yearly/monthly instance anytime on the console before it is automatically deleted.

Renewing an Instance on the Console

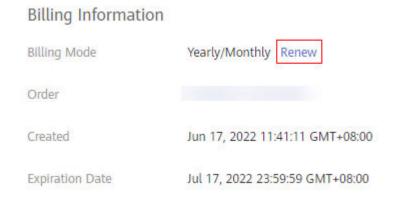
- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** On the **Instances** page, locate the instance that you want to renew and click **Renew** in the **Operation** column.

Figure 2-13 Renewal button



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 2-14 Renewal button



□ NOTE

To renew multiple yearly/monthly instances at a time, perform the following steps:

- 1. Select the yearly/monthly instances to be renewed.
- 2. Click **Renew** above the instance list.

Step 4 On the displayed page, renew the instances.

----End

Renewing a Subscription in Billing Center

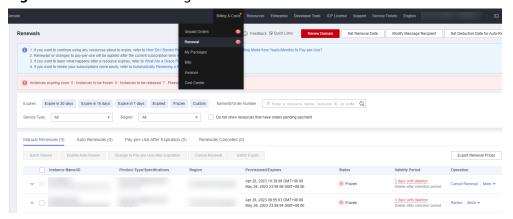
- **Step 1** Log in to the management console.
- **Step 2** Hover over **Billing & Costs** in the upper part of the console and choose **Renewal** from the drop-down list.

The **Renewals** page is displayed.

Step 3 Select the search criteria.

On the Manual Renewals, Auto Renewals, Pay-per-Use After Expiration, and Renewals Canceled pages, you can view the instances to be renewed.

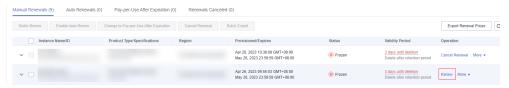
Figure 2-15 Renewal management



You can move all resources that need to be manually renewed to the **Manual Renewals** tab page. For details, see **Restoring to Manual Renewal**.

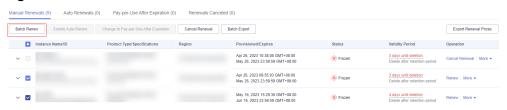
- Step 4 Manually renew resources.
 - Individual renewal: Locate an instance that you want to renew and click **Renew** in the **Operation** column.

Figure 2-16 Individual renewal



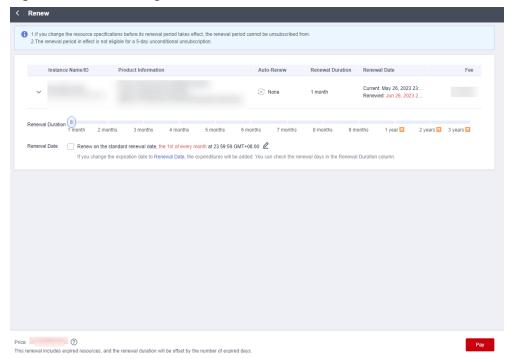
• Batch renewal: Select multiple instances that you want to renew and click **Batch Renew** in the upper left corner.

Figure 2-17 Batch renewal



Step 5 Select a renewal duration and optionally select Renew on the standard renewal date. For details, see Setting the Same Renewal Day for Yearly/Monthly Resources. Confirm the price and click Pay.

Figure 2-18 Confirming renewal



Step 6 Select a payment method and make your payment. Once the order is paid for, the renewal is complete.

----End

Setting the Same Renewal Day for Yearly/Monthly Resources

If the instances have different expiry dates, you can set the same renewal day, for example, the first day of each month, to make it easier to manage renewals.

In Figure 2-19, a user sets the same renewal day for two resources that will expire at different dates.

Figure 2-19 Setting the same renewal day for resources with different expiry dates



For more details, see **Setting a Renewal Date**.

2.6.3 Auto-renewing an Instance

Auto-renewal can prevent instances from being automatically deleted if you forget to manually renew them. The auto-renewal rules are as follows:

- The first auto-renewal date is based on when an instance expires and the billing cycle.
- The auto-renewal period of an instance depends on the subscription term.
 - Monthly subscriptions renew each month.
 - Yearly subscriptions renew each year.
- You can enable auto-renewal any time before an instance expires. By default, the system will make the first attempt to charge your account for the renewal at 03:00 seven days before the expiry date. If this attempt fails, it will make another attempt at 03:00 every day until the subscription is renewed or expired.
- After auto-renewal is enabled, you can still renew the instance manually if you want to. After a manual renewal is complete, auto-renewal is still valid, and the renewal fee will be deducted from your account seven days before the new expiry date.
- By default, the renewal fee is deducted from your account seven days before the new expiry date. You can change this auto-renewal payment date as required.

For more information about auto-renewal rules, see Auto-Renewal Rules.

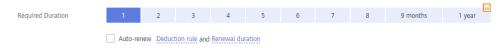
Prerequisites

Your yearly/monthly instance is not expired.

Enabling Auto-Renewal During Purchase

You can enable auto-renewal on the instance purchase page. For details, see **Buying an Instance**.

Figure 2-20 Enabling auto-renewal

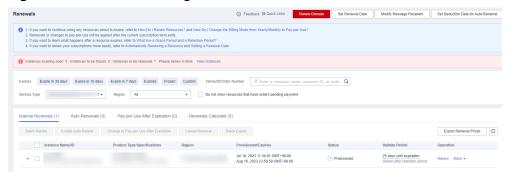


Enabling Auto-Renewal on the Renewals Page

- **Step 1** Log in to the management console.
- **Step 2** Hover over **Billing & Costs** in the upper part of the console and choose **Renewal** from the drop-down list.
- **Step 3** Select the search criteria.

- On the **Auto Renewals** page, you can view the resources that auto-renewal has been enabled for.
- You can enable auto-renewal for resources on the **Manual Renewals**, **Payper-Use After Expiration**, and **Renewals Canceled** pages.

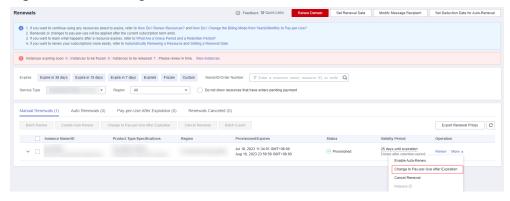
Figure 2-21 Renewal management



Step 4 Enable auto-renewal for yearly/monthly resources.

Enabling auto-renewal for a single instance: Locate the instance that you
want to enable auto-renewal for and choose More > Enable Auto-Renew in
the Operation column.

Figure 2-22 Enabling auto-renewal for an instance



Enabling auto-renewal for multiple instances at a time: Select the instances
that you want to enable auto-renewal for and click Enable Auto-Renew
above the list.

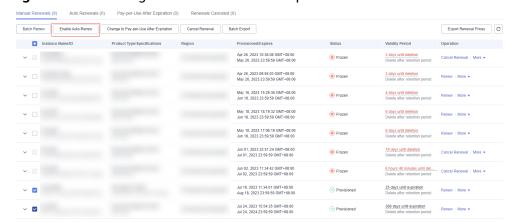
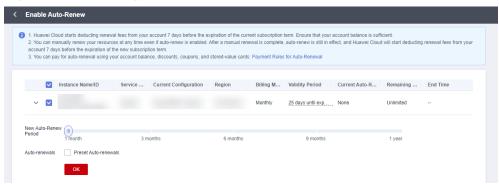


Figure 2-23 Enabling auto-renewal for multiple instances

Step 5 Select a renewal period, specify the auto-renewal times, and click **Pay**.





----End

2.7 Bills

You can view the resource usage and bills for different billing cycles on the **Bills** page in the Billing Center.

Bill Generation

Transaction records for yearly/monthly subscriptions are generated immediately after being paid for.

The usage of pay-per-use resources is reported to the billing system at a fixed interval. A pay-per-use resource is billed by the hour, day, or month, depending on the resource's usage type. The GeminiDB Cassandra instance usage is billed by the hour. For details, see **Bill Run for Pay-per-Use Resources**.

You are not charged immediately after a record is generated. For example, if a pay-per-use GeminiDB Cassandra instance (which is billed on an hourly basis) is deleted at 08:30, you will still have expenditures for the 08:00 to 09:00 hour. However, you will not likely be billed for the 08:00 to 09:00 hour until about 10:00.

On the **Bills** page of the Billing Center, select the **Bill Details** tab. **Expenditure Time** in the bill indicates the time when the pay-per-use resource is used.

Viewing Bills of a Specific Resource

[Method 1: Use the instance ID to search for a bill.]

- Step 1 Log in to the management console and choose Databases > GeminiDB Cassandra API.
- **Step 2** On the **Instances** page, locate the instance whose bill you want to view and click its name.
- **Step 3** Click the icon shown in the figure below to copy the instance ID.

Figure 2-25 Copying the instance ID



Step 4 On the top menu bar, choose **Billing & Costs** > **Bills**.

The **Bills** page is displayed.

Step 5 Choose Transactions and Detailed Bills > Bill Details. On the displayed page, select Resource ID as the filter criteria, enter the obtained instance ID, and click the Q icon.

Figure 2-26 Searching for a bill



By default, the bill details are displayed by usage and billing cycle. You can choose other display options as required. For details, see **Bill Details**.

----End

[Method 2: Use the resource name to search for a bill.]

- **Step 1** Log in to the management console and choose **Databases** > **GeminiDB Cassandra API**.
- **Step 2** On the **Instances** page, locate the instance whose bill you want to view and click its name.
- **Step 3** On the **Basic Information** > **Instance Information** page, obtain the instance name.

Figure 2-27 Copying the instance name

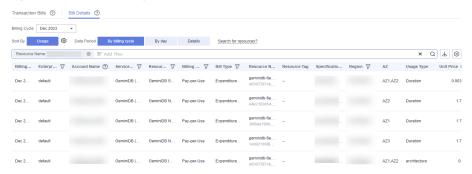


Step 4 On the top menu bar, choose **Billing & Costs** > **Bills**.

The **Bills** page is displayed.

Step 5 Choose **Transactions and Detailed Bills > Bill Details**. On the displayed page, select **Resource Name** as the filter criteria, enter the obtained instance ID, and click the cich.

Figure 2-28 Searching for a bill



By default, the bill details are displayed by usage and billing cycle. You can choose other display options as required. For details, see **Bill Details**.

----End

Scenario Example: Checking the Consistency of the Actual Usage and Billed Usage

Assume that you purchased a pay-per-use GeminiDB Cassandra instance at 10:09:06 on April 8, 2023 and deleted it later that day, at 12:09:06.

• Transaction Records

Pay-per-use GeminiDB Cassandra instance usage is calculated by the second and but billed on an hourly basis. You can check the transaction records against the actual usage. The billed resources are billed separately. For details, see **Table 2-11**.

Table 2-11 GeminiDB Cassandra transaction records

Service	GeminiDB Cassandra
Туре	

Resour ce Type	GeminiDB Cassandra storage
Billing Mode	Pay-per-use
Expend iture Time	For the period of time from 10:09:06 to 12:09:06 on April 08, 2023, 6 transaction records would be generated for the resource usage in the following periods: • 2023/04/08 10:09:06 - 2023/04/08 11:00:00 • 2023/04/08 11:00:00 - 2023/04/08 12:00:00 • 2023/04/08 12:00:00 - 2023/04/08 12:09:06
List Price	List price on the official website = Usage x Unit price x Capacity The GeminiDB Cassandra instance was used for 3,054 seconds in the first period, and the unit price can be obtained on the Pricing Details page. The list price for the first period = $(3054 \div 3600) \times 0.0007 \times 40 = \0.02375333 USD. Similarly, you can calculate the GeminiDB Cassandra instance list price for the other periods.
Discou nted Amoun t	Discounts offered for cloud services, for example, commercial discounts, partner authorized discounts, and promotional discounts. It is the discounted amount based on the list price.
Truncat ed Amoun t	Billing of Huawei Cloud is calculated to the 8th decimal place. However, the amount due is truncated to the 2nd decimal place. The third and later decimal places are referred to as the truncated amounts. Take the first period as an example. The truncated amount is \$0.00375333 USD.
Amoun t Due	Amount due = List price – Discount amount – Truncated amount Take the first period as an example. If the discount amount is 0, the amount due is \$0.02 USD (0.02375333 – 0 – 0.00375333).

• Bill details of the GeminiDB Cassandra instance

Bill details can display in multiple ways. By default, the bill details of a resource are displayed by usage and by billing cycle. **Table 2-12** illustrates the GeminiDB Cassandra instance bill details, which can be used to check against the actual usage.

Table 2-12 Bill details of the GeminiDB Cassandra instance

Service Type	GeminiDB Cassandra
Resour ce Type	GeminiDB Cassandra storage

Billing Mode	Pay-per-use
Resour ce Name/I D	Name and ID of a specific GeminiDB Cassandra instance Example: nosql-b388 and 21e8811a64bf4de88bc2e2556da17983in12
Specific ations	GeminiDB Cassandra storage
Usage Type	Duration for a pay-per-use GeminiDB Cassandra instance
Unit Price	When pay-per-use billing is used, the unit price is only provided if the amount is equal to the usage multiplied by the unit price. No unit price is provided in other pricing modes, for example, tiered pricing. You can search for the unit price for pay-per-use GeminiDB Cassandra instances on Product Pricing Details .
Unit	Displayed on the Product Pricing Details page. Example: USD/GB/hour.
Usage	Depends on the unit of the unit price, which is USD/GB/hour. Storage usage is billed by the hour. Example: 2 hours.
Usage Unit	Hour
List Price	List price on the official website = Usage x Unit price x Capacity The GeminiDB Cassandra instance is used for 2 hours in total, and the unit price is obtained on the Product Pricing Details page. The list price = $2 \times 0.0007 \times 40 = \0.056 USD.
Discou nted Amoun t	Discounts offered for cloud services, for example, commercial discounts, partner authorized discounts, and promotional discounts. It is the discounted amount based on the list price.
Amoun t Due	Amount that should be paid for used cloud services after discounts are applied.

2.8 Arrears

If your configured payment method is unable to pay for your bill, your account will be in arrears. You will need to update you payment method or to top up your account in a timely manner if you want to continue using your instance resources.

Arrears Reason

If you do not have yearly/monthly instances, your account falls into arrears any time your configured payment method is unable to pay for the used resources on the pay-per-use basis.

Arrears Impact

Yearly/Monthly

This is a pre-paid billing mode, so you can continue using yearly/monthly GeminiDB Cassandra resources even if your account is in arrears. However, you cannot perform operations such as purchasing GeminiDB Cassandra instances, upgrading instance specifications, and renewing subscriptions, because they will generate new expenditures.

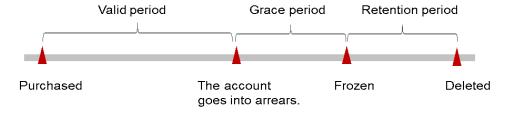
Pay-per-use

If your configured payment method is unable to pay a bill for pay-per-use resources, the resources enter a grace period. After you top up your account, Huawei Cloud will bill you for expenditures generated by the resources during the grace period. You can view the expenditures on the **Overview** page of the Billing Center.

If your account is still in arrears after the grace period ends, the resources enter the retention period and their status turns to **Frozen**. You cannot perform any operations on these resources.

After the retention period ends, the compute resources (vCPUs and memory) and EIPs will be released and cannot be restored.

Figure 2-29 Lifecycle of a pay-per-use instance



■ NOTE

The grace period and retention period are both 15 days.

Avoiding and Handling Arrears

Make sure you have a valid payment method configured as soon as possible after your account is in arrears. For details, see **Topping Up an Account**.

If a GeminiDB Cassandra instance is no longer used, you can delete it to avoid generating further expenditures.

To help make sure your account never falls into arrears, you can configure the **Balance Alert** on the **Overview** page of the Billing Center. Then, any time an expenditure quota drops to below the threshold you specify, Huawei Cloud automatically notifies you by SMS or email.

2.9 Billing Termination

Yearly/Monthly Resources

When you purchase a yearly/monthly resource, such as a yearly/monthly GeminiDB Cassandra instance, you make a one-time up-front payment. By default, the billing automatically stops when the purchased subscription expires.

- If you no longer need a resource, but the subscription has not yet expired, you
 can unsubscribe from it. Depending on what coupons were used for the
 purchase and on if 5-day unconditional unsubscription rules apply, Huawei
 Cloud may issue you a refund. For details about unsubscription rules, see
 Unsubscriptions.
- If you have enabled auto-renewal but no longer wish to automatically renew the subscription, disable it before the auto-renewal date (7 days before the expiration date by default) to avoid unexpected expenditures.

Pay-per-Use Resources

If pay-per-use resources, such as pay-per-use GeminiDB Cassandra instances, are no longer required, delete them in a timely manner.

Searching for Resources from Bills and Stopping Billing

To ensure that all related resources are deleted, you can search the billing records by resource ID, and then delete the resources you identify in this way.

[Method 1: Use the resource ID in the bill to search for the resource.]

Step 1 Log in to the management console. On the top menu bar, choose **Billing & Costs** > **Bills**.

The **Bills** page is displayed.

Step 2 Choose **Transactions and Detailed Bills** > **Bill Details**, and click the icon shown in the following figure to copy the resource ID.

Figure 2-30 Copying the resource ID



- Step 3 Log in to the management console and choose Databases > GeminiDB Cassandra API.
- **Step 4** Select the region where the resource is located, select **Instance ID** and enter the resource ID copied in **Step 2**, and click the Q icon to search for the resource.

Figure 2-31 Searching for resources



Step 5 Locate the instance you want to delete and click **More** > **Delete** in the **Operation** column. Ensure that the resource is not found in the list.

□ NOTE

You are billed one hour after the resource usage is calculated, so a bill may still be generated after the pay-per-use resource is deleted. For example, if you delete an instance (which is billed on an hourly basis) at 08:30, the expenditures for that hour from 08:00 to 09:00 are usually not billed until about 10:00.

----End

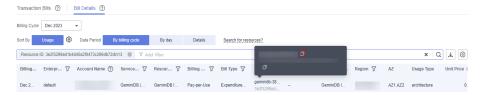
[Method 2: Use the resource name in the bill to search for the resource.]

Step 1 Log in to the management console. On the top menu bar, choose **Billing & Costs** > **Bills**.

The **Bills** page is displayed.

Step 2 Choose **Transactions and Detailed Bills** > **Bill Details**, and click the icon shown in the following figure to copy the resource name.

Figure 2-32 Copying the resource name



- **Step 3** Log in to the management console and choose **Databases** > **GeminiDB Cassandra API**.
- **Step 4** Enter the instance name copied in **Step 2** in the search box and click Q.

Figure 2-33 Searching for resources



Step 5 Locate the instance you want to delete and click **More** > **Delete** in the **Operation** column. Ensure that the resource is not found in the list.

You are billed one hour after the resource usage is calculated, so a bill may still be generated after the pay-per-use resource is deleted. For example, if you delete an instance (which is billed on an hourly basis) at 08:30, the expenditures for that hour from 08:00 to 09:00 are usually not billed until about 10:00.

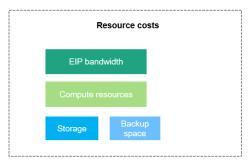
----End

2.10 Cost Management

2.10.1 Cost Composition

GeminiDB Cassandra costs consist of two parts:

- Resource costs: costs of compute and storage resources. For details, see Billing Modes.
- O&M costs: labor costs incurred during the use of GeminiDB Cassandra.





2.10.2 Cost Allocation

A good cost accountability system is a prerequisite for cost management. It ensures that departments, business teams, and owners are accountable for their respective cloud costs. An enterprise can allocate cloud costs to different teams or projects so as to have a clear picture of their respective costs.

Huawei Cloud **Cost Center** provides various tools for you to group costs in different ways. You can experiment with these tools and find a way that works best for you.

By linked account

The enterprise master account can manage costs by grouping the costs of its member accounts by linked account. For details, see **Viewing Costs by Linked Account**.

• By enterprise project

Before allocating costs, enable Enterprise Project Management Service (EPS) and plan your enterprise projects based on your organizational structure or service needs. When purchasing cloud resources, select an enterprise project so that the costs of resources will be allocated to the selected enterprise project. For details, see **Viewing Costs by Enterprise Project**.

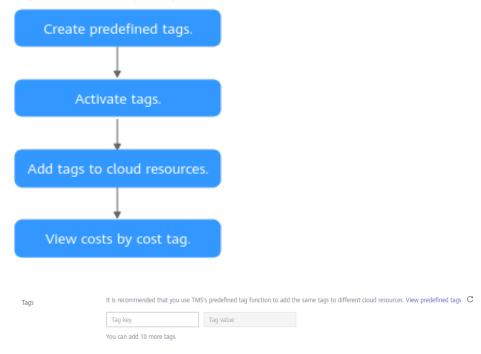
Figure 2-34 Selecting an enterprise project



By cost tag

You use tags to sort your Huawei Cloud resources in a variety of different ways, for example, by purpose, owner, or environment. The following is the process of managing costs by predefined tags (recommended).

Figure 2-35 Adding a tag



For details, see Viewing Costs by Cost Tag.

By cost category

You can use cost categories provided by **Cost Center** to split shared costs. Shared costs are the costs of resources (compute, network, storage, or resource packages) shared across multiple departments or the costs that cannot be directly split by cost tag or enterprise project. These costs are not directly attributable to a singular owner, and they cannot be categorized into a singular cost type. In this case, you can define cost splitting rules to fairly allocate these costs among teams or business units. For details, see **Viewing Cost By Cost Category**.

2.10.3 Cost Analysis

To precisely control and optimize your costs, you need a clear understanding of what parts of your enterprise incurred different costs. **Cost Center** visualizes your original costs and amortized costs using various dimensions and display filters for cost analysis so that you can analyze the trends and drivers of your service usage and costs from a variety of perspectives or within different defined scopes.

You can also use cost anomaly detection provided by **Cost Center** to detect unexpected expenses in a timely manner. In this way, costs can be monitored, analyzed, and traced.

For details, see Performing Cost Analysis to Explore Costs and Usage and Enabling Cost Anomaly Detection to Identify Anomalies.

2.10.4 Cost Optimization

You can identify resources with high costs based on the analysis results in the cost center, determine the causes of high costs, and take optimization measures accordingly.

Resource rightsizing

- View GeminiDB Cassandra monitoring metrics on Cloud Eye, such as the CPU, memory, and disk usage. If the current configuration is too high, you can reduce the configuration by changing specifications.
- Monitor idle GeminiDB Cassandra resources and delete idle instances in a timely manner.

Billing mode selection

Different types of services have different requirements on resource usage periods, so the most economical billing mode for one resource may not be the best option for another resource.

- For mature services that tend to be stable for the long term, select yearly/ monthly billing.
- For short-term, unpredictable services that experience traffic bursts and cannot afford to be interrupted, select pay-per-use billing.
- Monitor the lifecycle of instances and renew yearly/monthly resources that are about to expire in a timely manner.

2.11 Billing FAQs

2.11.1 What Are the Differences Between Yearly/Monthly and Pay-per-Use Billing?

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 billing is a postpaid payment mode. This billing mode allows you to make or cancel subscriptions at any time. Pricing is listed on a per-hour basis, but bills are calculated based on the actual usage duration.

2.11.2 Can I Switch Between Yearly/Monthly and Pay-per-Use Billing?

You can change the billing mode of your instance from yearly/monthly to pay-peruse or vice versa.

- For details about how to change the billing mode from yearly/monthly to a pay-per-use, see Yearly/Monthly to Pay-per-Use.
- For details about how to change the billing mode from pay-per-use to yearly/monthly, see Pay-per-Use to Yearly/Monthly.

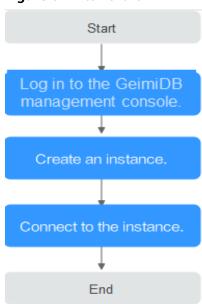
3 Getting Started with GeminiDB Cassandra API

3.1 Service Overview

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

Process

Figure 3-1 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.

3.2 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 Cassandra API.
- Step 3 On the Instances page, click Buy DB Instance.
- **Step 4** On the displayed page, specify a billing mode and instance specifications and click **Next**.

Figure 3-2 Billing mode and basic information



Table 3-1 Billing parameters

Parameter	Description
Billing Mode	Method that the instance is billed in. The value can be Yearly/Monthly or Pay-per-use.
	Yearly/Monthly
	 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.
	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.
	Pay-per-use
	 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 3-2 Basic information

Parameter	Description
Region	The region where the instance is deployed. 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.
DB Instance Name	 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 (_). After an instance is created, you can change its name. For details, see Changing an Instance Name.
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.
	An instance can be deployed in one or three AZs.
	If you want to deploy an instance in a single AZ, select one AZ.
	 If you want to deploy an instance across AZs for disaster recovery, select three AZs. Nodes of the instance are evenly distributed across the three AZs.

Figure 3-3 Specifications and storage

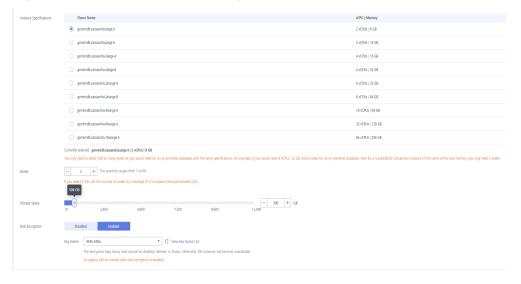


Table 3-3 Specifications and storage

Parameter	Description
Instance Specifications	Decoupled storage and compute and software-hardware synergy deliver twice or more the performance of an onpremises 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 .

Parameter	Description
Nodes	Number of nodes that the instance is deployed on.
	After an instance is created, you can add nodes. For details, see Adding Nodes.
Storage Space	Instance storage space. The range depends on the instance specifications. For details, see Instance Specifications .
	Select at least 1 GB each time you scale up the storage, and ensure that the increment is an integer.
	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:
	• 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.
	After an instance is created, you can scale up its storage if necessary. For details, see Scaling Storage Space .
	 NOTE 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.

Parameter	Description	
Disk Encryption	You can select to enable disk encryption based on service requirements.	
	Disabled: Data storage is not encrypted.	
	Enabled: Data store is encrypted. If you select this option, your data will be encrypted on disks and stored in ciphertext after you create a DB instance. When you download encrypted objects, the ciphertext will be decrypted into plain text and then sent to you. Disk encryption can improve data security and may have slight impacts on database writes and reads. Key Name: Select an existing key or create one.	
	NOTE	
	 This function is in the open beta test (OBT) phase. To use it, contact customer service. 	
	An agency will be created after disk encryption is enabled.	
	After an instance is created, the disk encryption status and the key cannot be changed.	
	 The key cannot be disabled, deleted, or frozen when being used. Otherwise, the database becomes unavailable. 	
	 For details about how to create a key, see "Creating a CMK" in Data Encryption Workshop User Guide. 	

Figure 3-4 Network configuration



Table 3-4 Network configuration

Parameter	Description
VPC	The virtual network where the instance is created. A VPC isolates networks for different services. You can select an existing VPC or create one.
	If there are no VPCs available, the system automatically allocates a VPC to you.
	For details, see "Creating a VPC" in the <i>Virtual Private Cloud User Guide</i> .
	NOTE
	 After a GeminiDB Cassandra instance is created, its VPC cannot be changed.
	 If you want to connect to a GeminiDB Cassandra instance from an ECS over a private network, ensure that the GeminiDB Cassandra instance and the ECS are in the same VPC. If they are not, create a VPC peering connection between them for access.
Subnet	A subnet where your instance is created. The subnet provides dedicated and isolated networks, improving network security. NOTE Create an IPv4 subnet or select an existing one. IPv6 subnets are not supported.
Security Group	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.
	If there are no security groups available, the system allocates one to you by default.
SSL	A security protocol. Secure Sockets Layer (SSL) certificates set up encrypted connections between clients and servers, preventing data from being tampered with or stolen during transmission.
	You can enable SSL to improve data security. After an instance is created, connect to the instance through SSL. For details, see SSL.
	NOTE After SSL is enabled, it cannot be disabled.

Figure 3-5 Database configuration

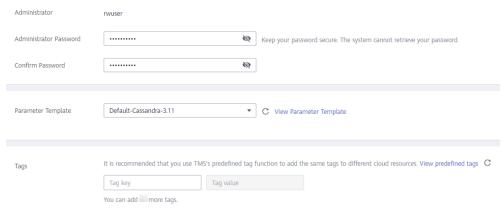


Table 3-5 Database configuration

Parameter	Description	
Administrator	Username of the administrator account. The default value is rwuser.	
Administrator Password	 Password of the administrator account. The password: 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. Keep your password secure. The system cannot retrieve it if it is lost. 	
Confirm Password	This password must be consistent with administrator password.	
Parameter Template	A template of parameters for creating an instance. The template contains API configuration values that are applied to one or more instances. After an instance is created, you can modify its parameters for	
	optimal performance. For details, see Modifying a Parameter Template .	
Enterprise	This parameter is provided for enterprise users.	
Project	An enterprise project groups cloud resources, so you can manage resources and members by project. The default project is default .	
	Select an enterprise project from the drop-down list. For more information about enterprise projects, see <i>Enterprise Management User Guide</i> .	

Figure 3-6 Tag configuration



Table 3-6 Tags

Parameter	Description
Tags	This setting is optional. Adding tags helps you better identify and manage your GeminiDB Cassandra instances.
	Each instance supports up to 10 tags by default. To use more tags, contact customer service to apply for a quota of 20 tags.
	A tag consists of a tag key and a tag value.
	• 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 (-).
	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 .

Figure 3-7 Required duration configuration



Table 3-7 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	 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

3.3 Instance Connections

3.3.1 Connection Methods

GeminiDB Cassandra can be accessed through Data Admin Service (DAS), private networks, and public networks.

Table 3-8 Connection methods

Method	Scenario	Description
DAS	You can log in to an instance on the console without using an IP address.	 Easy to use, secure, advanced, and intelligent By default, you have the permissions required for remote login. It is recommended that you use the DAS service to log in to DB instances. DAS is secure and convenient.
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

Method	Scenario	Description
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.	 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.	-

3.3.2 Connecting to an Instance Through DAS

DAS enables you to manage DB instances from a web-based console, simplifying database management and improving efficiency. You can connect and manage instances through DAS. By default, you have the permissions required for remote login. It is recommended that you use the DAS service to log in to instances. DAS is secure and convenient.

Procedure

- **Step 1** Log in to the management console.
- Step 2 In the service list, choose Databases > GeminiDB Cassandra API.
- **Step 3** On the **Instance Management** page, locate the target DB instance and click **Log In** in the **Operation** column.

Figure 3-8 Logging in to the database



Alternatively, click the target DB instance on the **Instance Management** page. On the displayed **Basic Information** page, click **Log In** in the upper right corner of the page.

Figure 3-9 Logging in to the database



Step 4 On the displayed login page, enter the administrator username and password and click **Log In**.

For details about how to manage databases through DAS, see **GeminiDB Cassandra Data Management**.

----End

3.3.3 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.

GeminiDB Cassandra API allows you to connect to an instance through an unencrypted connection or an encrypted connection (SSL). The SSL connection encrypts data and is more secure.

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.
 - If the CPU type is Kunpeng, download the Cassandra client installation package.
- 4. Before accessing an instance using an SSL connection, obtain the SSL certificate by referring to **Downloading the SSL Certificate**.

Non-SSL Connection

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 3-9 Parameter description

Parameter	Description
<db_host></db_host>	The private IP address of the instance to be connected. 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. If the instance you purchased has multiple nodes, select the private IP address of any node. Figure 3-10 Viewing the private IP address
<db_port></db_port>	Port number of the instance to be connected. The default port number is 8635. Replace it with the actual port number. Click the instance name to go to the Basic Information page and obtain the port number in the Network Information area. Figure 3-11 Viewing the port number Network Information VPC Subnet Security Group Jarabase Port B635
<db_user></db_user>	Database account. The default value is rwuser .

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

rwuser@cqlsh>

----End

SSL

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** Upload the SSL certificate to the ECS.
- **Step 4** 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 5 Run the following command to grant the execute permission on all files:

chmod +x *

Step 6 Connect to the DB instance in the directory where the cqlsh tool is located.

export SSL_CERTFILE=/<*PATH_OF_SSL_CERT_FILE*>

export SSL_VERSION=TLSv1_2

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

Example:

./cqlsh 192.168.1.8 8635 --ssl -u rwuser

Table 3-10 Parameter description

Parameter	Description
<path_of_ssl_cert_ FILE></path_of_ssl_cert_ 	SSL file path.

Parameter	Description
<db_host></db_host>	The private IP address of the instance to be connected. 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. If the instance you purchased has multiple nodes, select the private IP address of any node. Figure 3-12 Viewing the private IP address Private Pri
<db_port></db_port>	Port number of the instance to be connected. The default port number is 8635. Replace it with the actual port number. Click the instance name to go to the Basic Information page and obtain the port number in the Network Information area. Figure 3-13 Viewing the port number Network Information VPC Subnet Security Group Jroup 2 Subnet B635
<db_user></db_user>	Database account. The default value is rwuser .

Step 7 Check the results. If the following information is displayed, the connection is successful.

rwuser@cqlsh>

----End

Follow-up Operations

After logging in to the instance, you can create keyspaces, databases, or tables. For details, see **Buying and Connecting to a GeminiDB Cassandra Instance**.

3.3.4 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.

You can use a common connection or establish a secure connection with SSL if necessary.

Prerequisites

- Bind an EIP to the GeminiDB Cassandra instance node and set security group rules. For details, see <u>Binding and Unbinding an EIP</u> and <u>Configuring</u> <u>Security Group Rules</u>.
- 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.
 - If the CPU type is Kunpeng, download the Cassandra client installation package.
- 4. Before accessing an instance using an SSL connection, obtain the SSL certificate by referring to **Downloading the SSL Certificate**.

Non-SSL Connection

- **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 3-11 Parameter description

Parameter	Description	
<db_host></db_host>	Specifies the EIP bound to the instance to be connected. 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. If the instance you bought has multiple nodes, you can bind the EIP to any node to connect to the GeminiDB Cassandra instance.	
	Figure 3-14 Viewing EIPs 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.	
<db_port></db_port>	Port number of the instance to be connected. The default port number is 8635. Replace it with the actual port number. Click the instance name to go to the Basic Information page and obtain the port number in the Network Information area. Figure 3-15 Viewing the port number Network Information VPC Security Group Joun & Joun Pound Jou	
<db_user></db_user>	Database account. The default value is rwuser .	

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

rwuser@cqlsh>

----End

SSL Connection

- **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** Upload the SSL certificate to the ECS.
- **Step 4** 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 5 Run the following command to grant the execute permission on all files:

chmod +x *

Step 6 Connect to the DB instance in the directory where the cqlsh tool is located.

export SSL_CERTFILE=/<*PATH_OF_SSL_CERT_FILE*>

export SSL_VERSION=TLSv1_2

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

Example:

./cqlsh 192.168.1.8 8635 --ssl -u rwuser

Table 3-12 Parameter description

Parameter	Description	
<path_of_ssl_cert_fi LE></path_of_ssl_cert_fi 	The path of the SSL certificate.	
<db_host></db_host>	Specifies the EIP bound to the instance to be connected.	
	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.	
	If the instance you bought has multiple nodes, you can bind the EIP to any node to connect to the GeminiDB Cassandra instance.	
	Figure 3-16 Viewing EIPs	
	Note information	
	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.	

Parameter	Description		
<db_port></db_port>	Port number of the instance to be connected. The default port number is 8635. Replace it with the actual port number.		
	Click the instance name to go to the Basic Information page and obtain the port number in the Network Information area.		
	Figure 3-17 Viewing the port number		
	Network Information		
	VPC Security Group group ∠		
	Subnet Database Port 8635		
<db_user></db_user>	Database account. The default value is rwuser .		

Step 7 Check the results. If the following information is displayed, the connection is successful.

rwuser@cqlsh>

----End

Follow-up Operations

After logging in to the instance, you can create keyspaces, databases, or tables. For details, see **Buying and Connecting to a GeminiDB Cassandra Instance**.

3.3.5 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 > Purchasing 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 "Getting Started > Logging in to an ECS" in the **Elastic Cloud Server User Guide**.

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

3.3.6 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 <u>Buying an</u> <u>Instance</u>.
- For details about how to create an ECS, see "Getting Started > Purchasing 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.

```
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()
}()
```

3.4 Getting Started with Common Practices

After purchasing and connecting to a GeminiDB Cassandra DB instance, you can view common practices to better use it.

Table 3-13 Common practices

Practic	e	Description	
Datab ase	Basic Design	Describes basic design specifications of GeminiDB Cassandra.	
rules	Database Objects	Describes rules of using database objects.	
	Database Usage	Describes constraints of using databases.	
	Queries	Describes rules of querying GeminiDB Cassandra.	
Instan ce modif	Changing an Instance Name	Describes how to change the name of a GeminiDB Cassandra instance to identify different instances.	
icatio ns	Resetting the Administrator Password	Describes how to change your administrator password. For security reasons, change it periodically.	
	Changing vCPUs and Memory of an Instance	Describes how to change the CPU or memory of your instance to suit your service requirements.	
Data backu p	Managing Automated Backups	Describes how GeminiDB Cassandra API automatically creates backups for a DB instance during a backup window and saves the backups based on the configured retention period.	
	Managing Manual Backups	Describes how to create manual backups for a DB instance. These backups can be used to restore data for improved reliability.	
	Setting a Cross- Region Backup Policy	Describes how to set a cross-region backup policy for a DB instance. Then for disaster recovery, you can restore backups to a new instance in another region.	
	Creating a Table- level Backup	Describes how to create a table-level backup for a DB instance. If a database or table is deleted, maliciously or accidentally, backups can help restore your data.	
Data restor ation	Restoring Data to a New Instance	Describes how to restore an existing automated or manual backup to a new instance. The restored data is the same as the backup data.	
	Restoring a Backup to a Specific Point in Time	Describes how to use an automated backup to restore instance data to a specified point in time.	
Log mana geme nt	Slow Query Logs	Describes how to view slow query logs of a GeminiDB Cassandra database. The unit of the execution time is ms. You can identify the SQL statements that take a long time to execute and tune them based on slow query logs.	

4 Working with GeminiDB Cassandra API

4.1 Permissions Management

4.1.1 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.

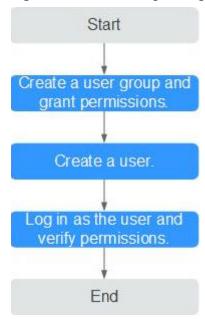
The following describes the procedure for granting permissions (see Figure 4-1).

Prerequisites

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

Process Flow

Figure 4-1 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 **GeminiDB 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 **GeminiDB 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.

4.1.2 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 **Permissions Policies and Supported Actions**.

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

- Visual editor: Select cloud services, actions, resources, and request conditions. 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

• 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 **GeminiDB FullAccess** 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:

• 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:

4.2 Instance Lifecycle

4.2.1 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 Cassandra API.
- **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

4.2.2 Deleting a Pay-per-Use 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 Cassandra API**.
- **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.

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

4.2.3 Recycling an Instance

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

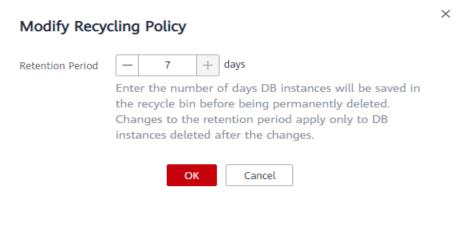
Precautions

- 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 Cassandra API**.
- **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 4-2 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 Cassandra API**.
- **Step 3** On the **Recycling Bin** page, locate the instance to be rebuilt and in the **Operation** column, click **Rebuild**.

Figure 4-3 Rebuilding an instance



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

----End

4.3 Instance Modifications

4.3.1 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 4-4 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:
 - $600 + (N \times 60) \le Total upgrade duration (s) \le 600 + (N \times 120)$

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 Cassandra API**.
- **Step 3** On the **Instances** page, locate the instance you want to upgrade and click **Upgrade Minor Version** in the **Compatible API** column.

Figure 4-5 Patch installation



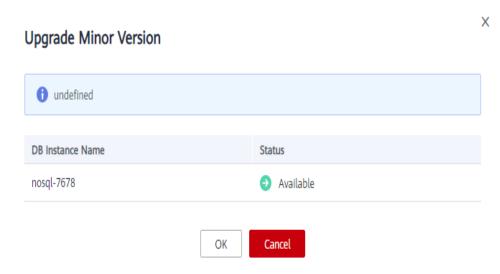
Alternatively, click the instance name to go to the **Basic Information** page. In the **DB Information** area, click **Upgrade Minor Version** in the **Compatible API** field.

Figure 4-6 Patch installation



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

Figure 4-7 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

4.3.2 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 Cassandra API.
- **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**.

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 Cassandra API**.
- **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 \times .

□ 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

4.3.3 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.
- For two instances with an intra-region DR or cross-region dual-active relationship, make sure that they have the same administrator passwords.
- 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 Cassandra API**.
- **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: $\sim !@#\%^*-=+?$

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 Cassandra API**.
- **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: $\sim !@#\%^*-=+?$

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

4.3.4 Scaling Storage Space

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

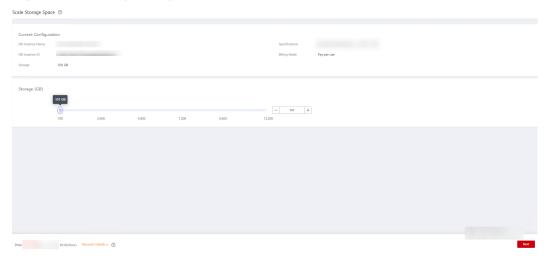
Precautions

- Storage space of an instance can be scaled up or down.
- When you scale down storage space of your instance, make sure that new storage space is 25% more than the used space and rounded up.
- Scaling up storage does not interrupt your services. After storage scale-up is complete, you do not need to restart your instance.
- Scaling down storage does not interrupt your services. After storage scale-down is complete, you do not need to restart your instance.
- If your instance is a yearly/monthly instance and its storage is used up, you
 will be billed on a pay-per-use basis for any additional storage. Scale up your
 storage space, and you can pay for the additional storage at yearly/monthly
 rates
- To view storage details on the **Instances** page, contact customer service.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **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 4-8 Scaling storage space



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 scaling task is ongoing, the instance status is **Scaling storage space**.
 - 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

4.3.5 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.

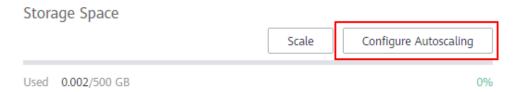
Precautions

- 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 Cassandra API**.
- **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 4-9 Configure Autoscaling



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

Configure Batch Autoscaling Configure Autoscaling Trigger If Available Storage Drops To Increment Storage Space Maximum storage supported by the current instance specifications The upper limit for autoscaling can only be set to the maximum storage supported by the current instance specification. If available storage drops to 10% or less, your storage will autoscale by 15% (in increments of 10 GB) of your allocated storage. If your account balance is insufficient, autoscaling will fail. All stateme... Enter a DB instance name o Q Autoscaling No Cancel

Figure 4-10 Configuring autoscaling parameters

Table 4-1 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 Cassandra API.
- **Step 3** Select instances and click **Auto Scale**.

Figure 4-11 Storage autoscaling



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

Figure 4-12 Configuring autoscaling parameters

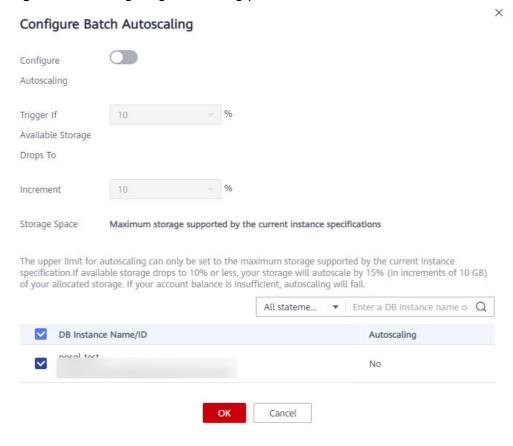


Table 4-2 Parameter description

Parameter	Description
Configure Autoscaling	The autoscaling toggle button.

Parameter	Description
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

4.3.6 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.

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.

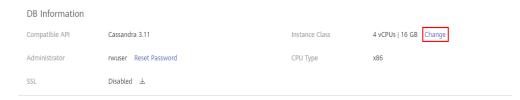
- To view vCPUs and nodes on the **Instances** page, contact customer service.
- When you change the CPU or memory of an instance, the change takes effect on each node one by one. During the period, read/write operations may fail and the latency may increase. Change the CPU or memory during off-peak hours.

Method 1

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.

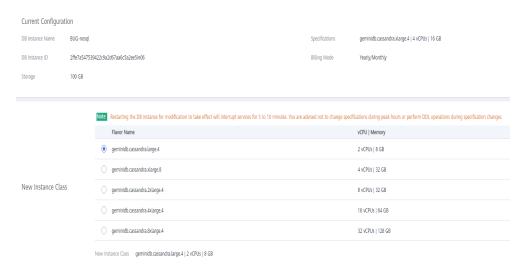
- **Step 3** On the **Instances** page, locate the instance whose vCPUs and memory you want to change and click its name.
- **Step 4** In the **DB Information** area on the **Basic Information** page, click **Change** next to the **Instance Class** field.

Figure 4-13 Changing instance specifications



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

Figure 4-14 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

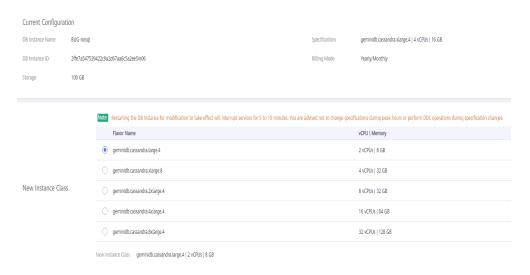
- **Step 1** Log in to the management console.
- Step 2 In the service list, choose Databases > GeminiDB Cassandra API.
- **Step 3** On the **Instances** page, locate the instance whose specifications you want to change and choose **More** > **Change Specifications** in the **Operation** column.

Figure 4-15 Changing specifications



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

Figure 4-16 Changing specifications



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

4.3.7 Adding Nodes

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

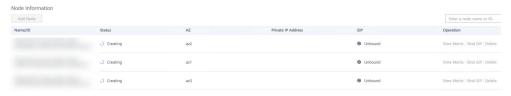
Precautions

- 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 Cassandra API**.
- **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 4-17 Basic information



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

Figure 4-18 Adding nodes



◯ 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

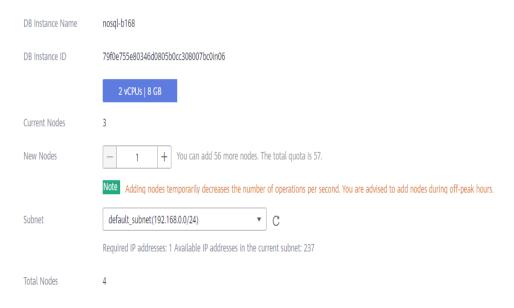
- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** On the **Instance Management** page, locate the instance you want to add nodes for and choose **More** > **Add Node** in the **Operation** column.

Figure 4-19 Adding nodes



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

Figure 4-20 Adding nodes



Ⅲ 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

4.3.8 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 Cassandra API**.
- **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

4.3.9 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 4-3**.
- By default, each instance supports up to 10 tags. To use more tags, contact customer service to apply for a quota of 20 tags.
- The tag name must comply with the naming rules described in **Table 4-3**.

Table 4-3 Naming rules

Parameter	Requirement	Example Value
Tag key	 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	 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 Cassandra API**.
- **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 Cassandra API**.
- **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

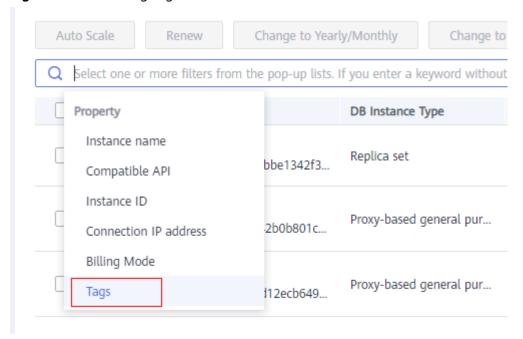
- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** On the **Instances** page, locate the instance whose tags you want to delete and click its name. The **Basic Information** page is displayed.
- **Step 4** In 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

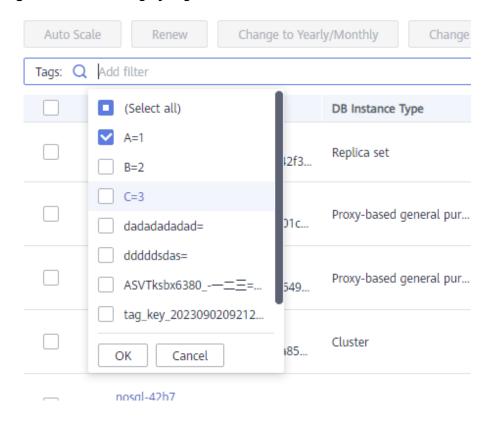
- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** On the **Instances** page, select **Tags** in the search box.

Figure 4-21 Selecting tags



Step 4 Select the tag to be queried and click **OK** to query information about instances associated with the tag.

Figure 4-22 Searching by tag



----End

4.3.10 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.

4.4 Connection Management

4.4.1 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 4-4**.

Table 4-4 Parameter description

Scenario	Description
Connecting to an instance over a private network	 Check whether the ECS and GeminiDB Cassandra instance are in the same security group: If yes, no security group rules need to be configured. If no, configure security group rules for them, respectively.
	 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	Add inbound rules for the security group associated with the GeminiDB Cassandra instance. For details, see Procedure .

Procedure

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **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.

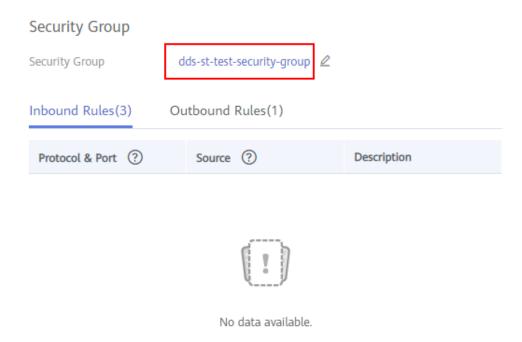
Figure 4-23 Security group



Method 2

On the **Basic Information** page, choose **Connections** in the navigation pane on the left. In the **Security Group** area on the right, click the name of the security group. The **Security Group** page is displayed.

Figure 4-24 Security group



Step 5 Add an inbound rule.

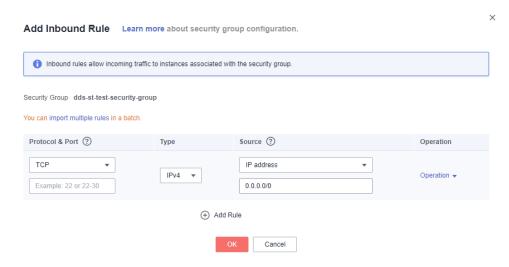
1. Click the **Inbound Rules** tab.

Figure 4-25 Inbound rules



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

Figure 4-26 Adding a rule



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

Table 4-5 Inbound rule settings

Parame ter	Description	Example Value
Protoco l & Port	 Network protocol. Available options are All, TCP, UDP, ICMP, or GRE 	ТСР
	 Port: The port or port range that allows the access to the ECS. Range: 1 to 65535 	
Туре	IP address type. This parameter is available only after the IPv6 function is enabled. - 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: - 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

Parame ter	Description	Example Value
Descrip tion	(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

4.4.2 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.

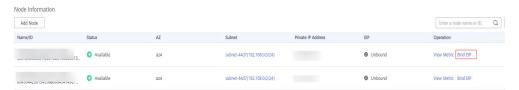
Precautions

- 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

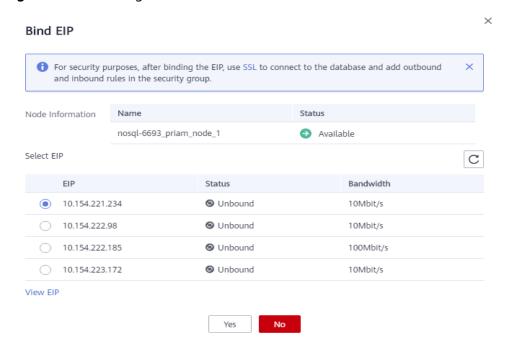
- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** On the **Instances** page, locate the GeminiDB Cassandra instance that you want to bind an EIP to and click its name.
- **Step 4** On the **Basic Information**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

Figure 4-27 Binding an 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.

Figure 4-28 Selecting an EIP



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 Cassandra API**.
- **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 4-29 Unbinding an EIP



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

To bind an EIP to the instance again, see **Binding an EIP**.

----End

4.4.3 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 Cassandra API**.
- **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 4-30 Viewing IP addresses



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

Figure 4-31 Viewing the port number



In the navigation pane on the left, click **Connections** to view private IP addresses, EIPs, and port number of the instance.

Figure 4-32 Viewing the IP addresses and port number



4.4.4 Changing a Database Port

Scenarios

GeminiDB Cassandra API allows you to change the database port of an instance to ensure security.

The database port cannot be changed when the instance is in any of the following statuses:

- Frozen
- Restarting
- Adding nodes
- Changing specifications
- Scaling storage space
- Deleting nodes

Precautions

To change the database port of your instance, contact customer service.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** On the **Instances** page, locate the instance whose database port you want to change and click its name.
- **Step 4** In the navigation pane on the left, choose **Connections**.
- **Step 5** In the **Basic Information** area, click $\stackrel{\checkmark}{=}$ to the right of the **Database Port** field.

The database port number can range from 2100 to 9500 but cannot be 2180, 2887, 3887, 7000, 7001, 7199, 8018, 8079, 8091, 8092, 8479, 8484, 8636, and 8999.

- To submit the change, click ✓. This operation takes about 1 to 5 minutes.
- To cancel the change, click X.
- **Step 6** View the change result at the **Basic Information** area.

----End

4.4.5 Changing a Security Group

You can change security groups of your GeminiDB Cassandra instances.

Precautions

- If you are adding nodes to an instance, the security group of the instance cannot be changed.
- This function is in the open beta test (OBT) phase. To use this function, contact customer service.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** On the **Instances** page, locate the instance whose security group you want to change and click its name.
- **Step 4** In the navigation pane on the left, choose **Connections**.
- **Step 5** In the **Security Group** area, click beside the security group name and select the required security group.
 - To submit the change, click ✓. This process takes about 1 to 3 minutes.
 - To cancel the change, click X.
- **Step 6** View the change result.

----End

4.4.6 Configuring an SSL Connection

Secure Socket Layer (SSL) is an encryption-based Internet security protocol for establishing an encrypted link between a server and a client. It provides privacy, authentication, and integrity to Internet communications.

- Authenticates users and servers, ensuring that data is sent to the correct clients and servers.
- Encrypts data to prevent it from being intercepted during transfer.
- Ensures data integrity during transmission.

After SSL is enabled, you can establish an encrypted connection between your client and the instance you want to access to improve data security.

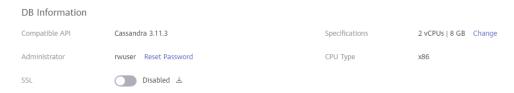
Precautions

- After you enable or disable SSL, the established connection is interrupted. Restart the instance for the change to take effect.
- Enabling SSL will prolong network connection response and increase CPU usage. So, evaluate impacts on service performance before enabling SSL.
- The SSL function provided by GeminiDB Cassandra supports only TLS 1.3 or later.

Enabling SSL

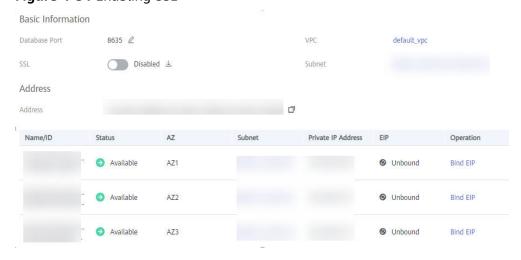
- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** On the **Instances** page, locate the instance that you want to enable SSL for.
- **Step 4** In the **DB Information** area, click to enable SSL.

Figure 4-33 Enabling SSL



Alternatively, choose **Connections** in the navigation pane on the left. In the **Basic**Information area, click in the **SSL** field to enable SSL.

Figure 4-34 Enabling SSL



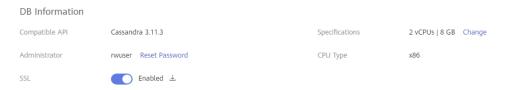
After SSL connection is enabled, you can connect to the GeminiDB Cassandra instance through SSL connections. For details, see **SSL**.

----End

Disabling SSL

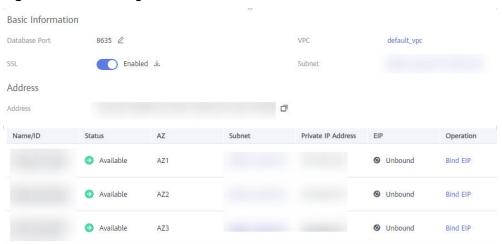
- **Step 1** Log in to the management console.
- Step 2 In the service list, choose Databases > GeminiDB Cassandra API.
- **Step 3** On the **Instances** page, locate the instance that you want to disable SSL for.
- **Step 4** In the **DB Information** area, click to disable SSL

Figure 4-35 Disabling SSL



Alternatively, choose **Connections** in the navigation pane on the left. In the **Basic Information** area, click in the **SSL** field to disable SSL.

Figure 4-36 Disabling SSL



After SSL connection is disabled, you can connect to the GeminiDB Cassandra instance an unencrypted connection. For details, see **Non-SSL Connection**.

----End

4.4.7 Downloading the SSL Certificate

Secure Sockets Layer (SSL) certificates set up encrypted connections between clients and servers, preventing data from being tampered with or stolen during transmission.

To improve data security, you can enable SSL when creating an instance. After the instance is created, you can connect to it using an SSL certificate.

This section describes how to obtain an SSL certificate.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** On the **Instances** page, locate the instance that you want to connect to and click its name.
- **Step 4** In the **DB Information** area, click in the **SSL** field to download the SSL certificate.

Figure 4-37 Downloading the SSL certificate



----End

4.5 Data Management

4.5.1 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 Download Binary Package.

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 4-6**.

Table 4-6 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.
I	Or. Use vertical bars () to separate elements. You can input any element. Do not enter only vertical bars.
	Repeatable. The ellipsis () indicates that you can repeat syntax elements multiple times as required.
'Literal string'	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.
<datatype1,dat atype2></datatype1,dat 	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 (,).
cql_statement;	End a CQL statement. Semicolons (;) end all CQL statements.

Symbol	Description
[]	Use two hyphens () to separate command line options from command arguments. This syntax is useful when parameters may be mistaken for command arguments.
' <schema> </schema> '	Search CQL only; single quotation marks (') enclose the entire XML schema declaration.
@xml_entity='x ml_entity_type'	Search CQL only; identify entities and literal values to overwrite XML elements in schemas and solrConfig files.

COPY Usage Suggestions

Table 4-7 Description

Co mm and	Paramet er	Description	Default Value	Applicability
TO/ FRO M	DELIMITE R	A single character used to separate fields.	English comma,	-
TO/ FRO M	QUOTE	A single character that contains a field value.	п	-
TO/ FRO M	ESCAPE	Escapes a single character using the QUOTE character.	\	-
TO/ FRO M	HEADER	Boolean value (true false), indicating the name of the column in the first row.	FALSE	-
		True matches the field name with the imported column name and inserts the column name into the first row of the exported data.		
TO/ FRO M	NULL	Filled value of the field whose query result is empty. You can set this parameter as required.	Empty string ()	-

Co mm and	Paramet er	Description	Default Value	Applicability
TO/ FRO M	DATETIM EFORMA T	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/ FRO M	MAXATTE MPTS	Maximum number of retry times when an error occurs.	5	-
TO/ FRO M	REPORTF REQUEN CY	Frequency of displaying the status, in seconds.	0.25	-
TO/ FRO M	DECIMAL SEP	Delimiter character for decimal values.	English full stop.	-
TO/ FRO M	THOUSA NDSSEP	Separator of a thousand array.	None	-
TO/ FRO M	BOOLSTY LE	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	-

Co mm and	Paramet er	Description	Default Value	Applicability
TO/ FRO M	NUMPRO CESSES	Number of working processes.	16	The default value of this parameter is the number of kernels on the computer minus one. There is no maximum value for this parameter. 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.

Co mm and	Paramet er	Description	Default Value	Applicability
TO/ FRO M	CONFIGFI LE	Specifies a cqlshrc configuration file to set the WITH option. NOTE Command line options always overwrite the cqlshrc file.	None, user- defined	-
TO/ FRO M	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/ FRO M	ORIGIN	Check whether the database to be imported or exported is an opensource Cassandra database. If the open-source Cassandra is used, the value is True . If GeminiDB Cassandra is used, the value is False .	False	-
FRO M	CHUNKSI ZE	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.

Co mm and	Paramet er	Description	Default Value	Applicability
FRO M	INGESTR ATE	Approximate import rate per second.	100000	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.

Co mm and	Paramet er	Description	Default Value	Applicability
FROM	MAXBAT CHSIZE	Maximum size of a batch file to be imported.	20	The value of this parameter can be as large as possible but cannot exceed the upper limit. MAXBATCHSI ZE x The size of a single row < batch_size_fa il_threshold_i n_kb. If the batch size is too large, an alarm will be reported and rejected. Set the following parameters in cassandra.ya ml: batch_size_w arn_threshol d_in_kb (The current value is 5.) batch_size_fa il_threshold_i n_kb (The current value is 50.)

Co mm and	Paramet er	Description	Default Value	Applicability
FRO M	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.
FRO M	MAXROW S	Maximum number of rows. The value -1 indicates that there is no upper limit.	-1	-
FRO M	SKIPROW S	Number of rows to skip.	0	-
FRO M	SKIPCOLS	A comma-separated list of column names to skip.	None, user- defined	-
FRO M	MAXPARS EERRORS	Maximum number of global parsing errors. The value -1 indicates that there is no upper limit.	-1	-
FRO M	MAXINSE RTERROR S	Maximum number of global insertion errors. The value -1 indicates that there is no upper limit.	-1	-

Co mm and	Paramet er	Description	Default Value	Applicability
FRO M	ERRFILE	A file that stores all rows that are not imported.	import_ ks _ table .err	-
		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.		
FRO M	TTL	The time to live is in seconds. By default, data does not expire.	3600	-
ТО	ENCODIN G	Output character string type.	UTF-8	-

Co mm and	Paramet er	Description	Default Value	Applicability
ТО	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 (for example, Huawei Cloud obsfs is used), you can increase the value.

Co mm and	Paramet er	Description	Default Value	Applicability
TO	PAGETIM EOUT	The page times out to obtain the result.	10	The value is an integer, indicating the timeout interval for obtaining each page. The unit is second. The default value is 10 seconds. • 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 automaticall y 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

Co mm and	Paramet er	Description	Default Value	Applicability
				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.
ТО	BEGINTO KEN	Minimum token for exporting data.	None, user- defined	The value is a string, indicating the minimum token to be considered during data export. Records with smaller tokens
				will not be exported. The default value is empty, indicating that there is no minimum token.

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

Co mm and	Paramet er	Description	Default Value	Applicability
ТО	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.

Co mm and	Paramet er	Description	Default Value	Applicability
ТО	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 MAXOUTPUTSIZ E.
то	dataform ats	Output file format. Currently, this parameter can only be set to json.	None, user- defined	-
то	DATATYP E	The file format can be Parquet or ORC.	None, user- defined	-
ТО	RESULTFI LE	The exported file containing detailed results.	None, user- defined	You are advised to add this parameter when exporting data to improve statistics efficiency.
ТО	wherecon dition	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.

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 4-38 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:
6 rows exported in 0.201 seconds.
Processed: 6 rows; Rate:
                                                        40 rows/s
                             20 rows/s; Avg. rate:
Results
                         : success
                        : 6
Total operation
Total operation : 6
Total operation time : 0.201 seconds
Operation rate
                        : 40.468307430069224 rows/s
Total ranges
                         : 25
Success ranges
                         : 25
                         : 0
Failed ranges
Num processes
                         : 15
Max attempts
                         : 5
Ranges Results:
ranges
                                             result
                                                        exported rows
(-7, 683212743470724096)
                                             success
                                                        1
(3074457345618258593, 3757670089088982528)
                                                        1
                                             SUCCESS
(2220441416279853312, 3074457345618258593)
                                             success
                                                        1
(-6148914691236517207, -5465701947765792768) success
(-854015929338405120, -7)
                                             success
                                                        1
(-2391244602147534336, -1537228672809129307)
```

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.

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 4-39 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
                         : 0
Failed ranges
                        : 15
Num processes
Max attempts
                         : 5
Ranges Results:
ranges
                                             result
                                                        exported rows
(-854015929338405120, -7)
                                                        1
                                             success
(-7, 683212743470724096)
                                                        1
                                             success
(3074457345618258593, 3757670089088982528)
                                                        1
                                             success
(-6148914691236517207, -5465701947765792768) success
                                                       1
(2220441416279853312, 3074457345618258593)
                                                       1
                                             success
(-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.

- 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 4-40 Querying data

5. Import the cyclist_firstname.csv file.

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

Figure 4-41 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 4-42 Import succeeded

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

7. Import the cyclist_lastname.csv file.

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

Figure 4-43 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).
```

8. Check whether the data is updated.

SELECT * FROM cycling.cyclist_name;

The query result is displayed,

Figure 4-44 Import succeeded

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

----End

Helpful Links

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

4.6 Intra-region DR

4.6.1 Creating a DR Instance

GeminiD 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 Cassandra API**.
- **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 4-8 Basic information

Parameter	Description
Billing	Select Yearly/Monthly or Pay-per-use .
Mode	Yearly/Monthly
	 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.
	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.
	Pay-per-use
	 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.
Region	The region is the same as that of the primary instance.
DB Instance	The instance name:
Name	The 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	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.
	An instance can be deployed in one or three AZs.
	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 4-9 Specifications and storage

Parameter	Description
Instance Specifications	vCPUs and memory of the instance. 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. After an instance is created, you can change its vCPUs and memory by referring to Changing vCPUs and Memory of an Instance.
Nodes	Specify the number of nodes based on service requirements. After an instance is created, you can add nodes by referring to Adding Nodes.
Storage Space	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. 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: • 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. After an instance is created, you can scale up its storage space by referring to Scaling Storage Space. NOTE • 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 4-10 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. Access from the 192.168.0.0/24 CIDR block in the security group should be allowed to ensure that DR instances can be created and work properly.
SSL	A security protocol. Secure Sockets Layer (SSL) certificates set up encrypted connections between clients and servers, preventing data from being tampered with or stolen during transmission.
	You can enable SSL to improve data security. After an instance is created, connect to it through SSL.

Table 4-11 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 4-12 Tags

Parameter	Description
Tags The setting is optional. Adding tags helps you better ide and manage your instances. Each instance supports up tags. To use more tags, contact customer service to app quota of 20 tags.	
	A tag consists of a tag key and a tag value.
	• 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 (-).
	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 .

Table 4-13 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.	
Auto-renew	 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

4.6.2 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 Cassandra API**.
- **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

4.7 Cross-region Dual-active DR

4.7.1 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**.

4.7.2 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

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

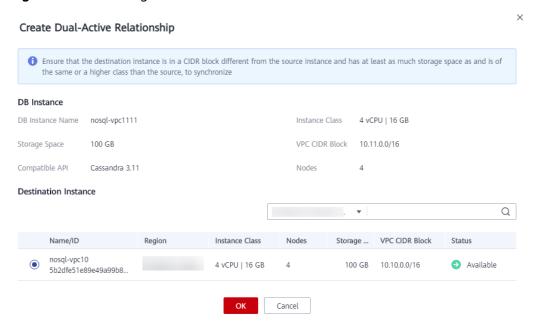
Procedure

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** 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.

Figure 4-45 Creating a dual-active relationship

Step 4 On the **Create Dual-Active Relationship** dialog box, locate the destination instance as the dual-active DR instance.

Figure 4-46 Selecting the destination instance



Ⅲ 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 5 Click OK.

- **Step 6** 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

4.7.3 Deleting a Dual-active Relationship

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

Procedure

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** 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.

Figure 4-47 Deleting a dual-active relationship



- Step 4 If you have enabled operation protection, in the displayed dialog box, 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 **Deleting dual-active relationship**. After the relationship is deleted, the instance status changes to **Available**.

----End

4.8 Data Backup

4.8.1 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 4-14 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.
Creating a Table- level Backup	If a database or table is deleted maliciously or accidentally, you can use backups to restore data.

How Backup Works

GeminiDB Cassandra provides a dedicated node (seed) responsible for managing backups. As shown in the following figure, a GeminiDB Cassandra cluster chooses the seed node for backing up data. The node takes snapshots of data in seconds and then stores them as compressed backups in OBS buckets, without using any store space of your instance. The CPU usage may increase 5% to 15% because uploading backups consumes CPU resources.

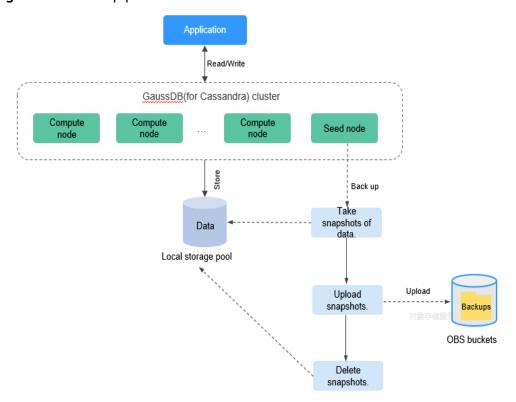


Figure 4-48 Backup process

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.

4.8.2 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.

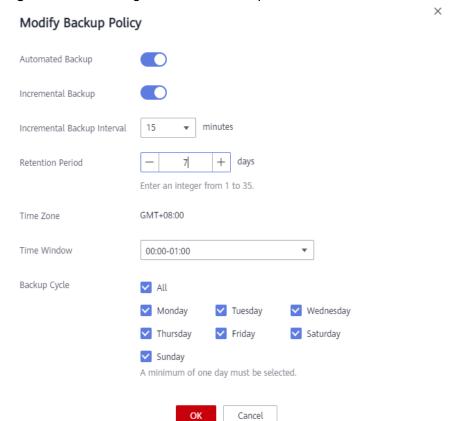


Figure 4-49 Enabling automated backup

- Incremental Backup: Incremental backup is enabled by default. You can manually enable or disable incremental backup. After you enable this function, it will take effect at the next full backup. Selecting Create a backup immediately after the incremental backup policy is modified is suggested. If you select this option, the system starts a full backup and the incremental backup policy takes effect. If you do not select this option, the policy takes effect at the next full backup. After incremental backup is enabled, the system stores backup data in OBS.
- **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 00:00–01:00. 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.

Ⅲ 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.

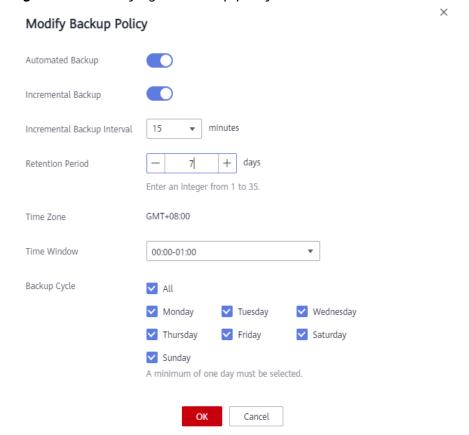
- After the DB instance is created, you can modify the automated backup policy as needed. You can change the time window after the DB instance is created. The system backs up data based on the automated backup policy you have set.
- After the automated backup policy is disabled, any automated backups 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 Cassandra API.
- **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 one 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 4-50 Modifying the backup policy



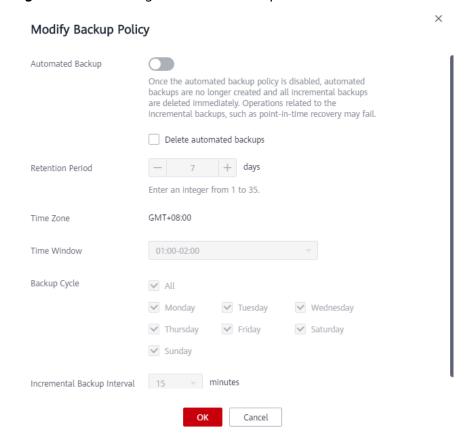
Step 5 Check or manage the generated backups on the **Backups** or **Backups & Restorations** page.

----End

Disabling an Automated Backup Policy

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **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 one the left, and click **Modify Backup Policy**.
- Step 5 In the displayed dialog box, click to disable automated backup and click OK.

Figure 4-51 Disabling automated backup



When your 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

- Log in to the management console.
- b. In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- 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 Cassandra API**.
- 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**.

4.8.3 Configuring an Incremental Backup Policy

GeminiDB Cassandra allows you to manually configure incremental backup based on your service requirements.

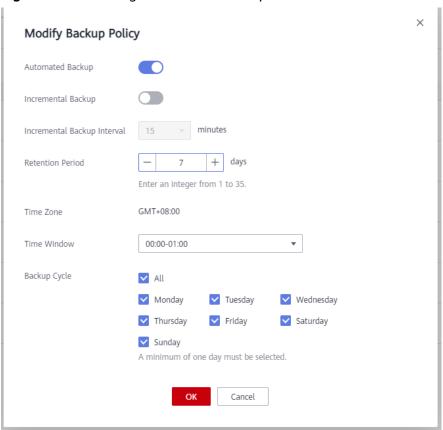
Precautions

• To enable or disable incremental backup, contact customer service.

Enabling Incremental Backup

- **Step 1** Log in to the management console.
- Step 2 In the service list, choose Databases > GeminiDB Cassandra API.
- **Step 3** On the **Instances** page, click the instance you want to enable incremental backup for and click its name.
- **Step 4** Choose **Backups & Restorations** in the navigation pane one the left, and click **Modify Backup Policy**. In the displayed dialog box, enable incremental backup.

Figure 4-52 Enabling incremental backup



Step 5 (Optional) Select **Create a backup immediately after the incremental backup policy is modified**. If you select this option, the system starts a full backup and the incremental backup policy takes effect immediately. If you do not select this option, the policy takes at the next full backup.

Modify Backup Policy Automated Backup Incremental Backup 10 minutes Incremental Backup Interval Create a backup immediately after the incremental backup policy is modified. + days Retention Period Enter an integer from 1 to 35. GMT+08:00 Time Zone Time Window 01:00-02:00 • Backup Cycle ✓ All Monday Tuesday Wednesday Thursday Friday Saturday Sunday Cancel

Figure 4-53 Selecting Create a backup immediately after the incremental backup policy is modified

----End

Disabling Incremental Backup

- **Step 1** Log in to the management console.
- Step 2 In the service list, choose Databases > GeminiDB Cassandra API.
- **Step 3** On the **Instances** page, click the instance you want to disable incremental backup for and click its name.
- **Step 4** Choose **Backups & Restorations** in the navigation pane one the left, and click **Modify Backup Policy**. In the displayed dialog box, disable incremental backup.

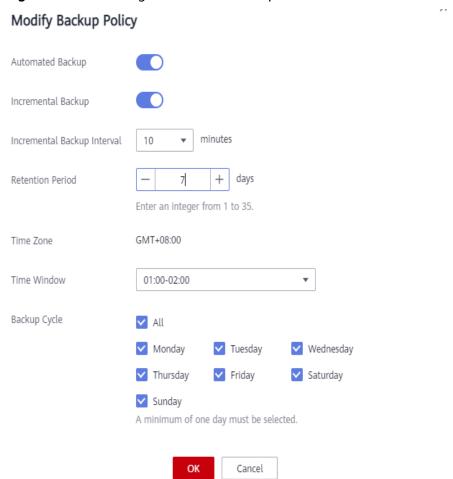


Figure 4-54 Disabling incremental backup

Step 5 After incremental backup is disabled, existing incremental backups will be deleted.

----End

4.8.4 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 Cassandra API**.
- **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 4-55 Setting a cross-region backup policy

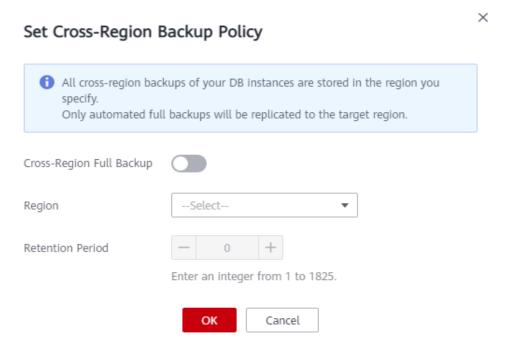


Table 4-15 Parameter description

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 7 Click OK.

Step 8 Choose **Backups** > **Cross-Region Backups** and manage the created backup.

Figure 4-56 Cross-region backups



All cross-region backups are displayed by default.

- To modify the cross-region backup policy, click **Set Cross-Region Backup** in the **Operation** column.
- To view all cross-region backups, click View Cross-Region Backup. To restore
 a backup to a new instance, click Restore in the Operation column. For
 details, see Restoring Data to a New Instance.

Figure 4-57 Restoring a cross-region backup



□ NOTE

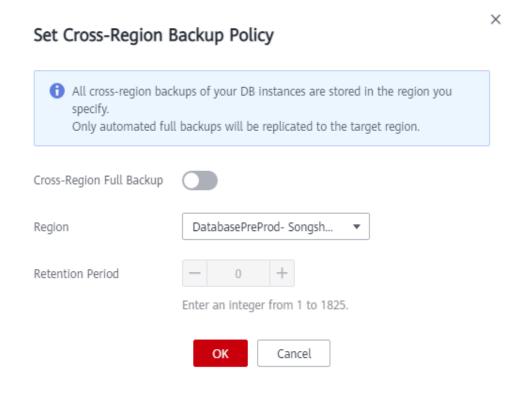
- Cross-region full backup replicates only automated full backups to another region for full restoration. Any of the full backups can be restored to a new instance that has no relationships with the original instance.
- The new instance uses the same parameter group as the original instance.
- During the instance restoration, backup files are downloaded from OBS buckets to the data directory of the new instance. The download bandwidth of OBS is 40 MB/s.
- If the original instance has cross-region backup disabled, the restoration may fail.
- Cross-region backup is not supported for instances that has disk encryption enabled.

----End

Disabling Cross-Region Backup

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **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, disable **Cross-Region Full Backup**.

Figure 4-58 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

4.8.5 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 Cassandra API.
- 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**.

Create Backup

DB Instance Name nosql-6693

* Backup Name backup-ee06

Description

?
/256

Figure 4-59 Creating a manual backup

Table 4-16 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 >!<"&'=

Cancel

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.

OK

----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 Cassandra API**.
- 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 Cassandra API**.
- 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**.

4.8.6 Creating a Table-level Backup

GeminiDB Cassandra API allows you to create table-level backup for your instance. If a database or table is deleted maliciously or accidentally, you can use backups to restore data.

Precautions

To use this function, contact customer service to apply for the required permissions.

Procedure

- **Step 1** Log in to the management console.
- Step 2 In the service list, choose Databases > GeminiDB Cassandra API.
- **Step 3** On the **Instances** page, locate the instance that you want to create a table-level backup for and click its name.
- **Step 4** In the navigation pane on the left, choose **Backups & Restorations**.
- **Step 5** On the displayed page, click **Create Table-level Backup**.
- **Step 6** On the displayed page, configure related parameters.

Name of the control o

Figure 4-60 Creating a table-level backup

Table 4-17 Parameter description

Parameter	Description
Backup Name	The backup name can contain 4 to 64 characters and must start with a letter. The name is case-sensitive and can contain only letters, digits, hyphens (-), or underscores (_).
Description	The description can include up to 256 characters and cannot contain carriage return characters and the following special characters: >!<"&'=
Databases and Tables	You can select the databases and specify tables therein that you want to back up.

Step 7 Click OK.

Step 8 Choose **Backups and Restorations** > **Table-level Backup** and manage the created backup.

Alternatively, click **Backups** in the navigation pane on the left, choose **Intraregion Backups** > **Table-level Backups**, and manage the created backup.

Figure 4-61 Managing the created table-level backup



Click View Tables to view tables contained in the backup file.

- Click **Restore** in the **Operation** column to restore the backup to a new instance. For details, see **Restoring Data to a New Instance**.
- Click **Delete** in the **Operation** column to delete the created backup.

NOTICE

Deleted backups cannot be recovered.

----End

4.8.7 Setting a Table-level Backup Policy

GeminiDB Cassandra API allows you to enable the table-level backup policy and back up table data based on the policy. If a database or table is deleted maliciously or accidentally, you can use backups to restore data.

Precautions

To use this function, contact customer service to apply for the required permissions.

Enabling or Modifying a Table-level Backup Policy

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** On the **Instances** page, locate the instance that you want to enable the table-level backup policy for and click its name.
- **Step 4** In the navigation pane on the left, choose **Backups & Restorations**.
- **Step 5** On the displayed page, click **Configure Table-level Backup Policy**.
- **Step 6** On the displayed page, configure related parameters.

Figure 4-62 Setting a table-level backup policy

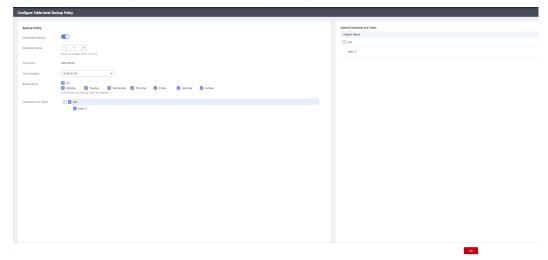


Table 4-18 Parameter description

Parameter	Description
Automated Backup	After automated backup is enabled, the system will back up table data based on the preset backup policy.
Retention Period	 Automated backup files are stored for 7 days by default. The backup retention period ranges from 1 to 732 days. 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 01:00-02:00 or 12:00-13:00. The backup time is in GMT format. If the DST or standard time is switched, the backup time segment changes with the time zone.
Backup Cycle	All: Each day of the week is selected. The system automatically backs up data every day. This option is selected by default.
	Alternative, you can select one or more days in a week. The system automatically backs up data at the specified time.
Databases and Tables	Select the tables that you want to back up.

□ 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.

Step 7 Click OK.

Step 8 Choose **Backups and Restorations** > **Table-level Backups** and manage the created backup.

Alternatively, click **Backups** in the navigation pane on the left, choose **Intraregion Backups** > **Table-level Backups**, and manage the created backup.

Figure 4-63 Managing the created table-level backup



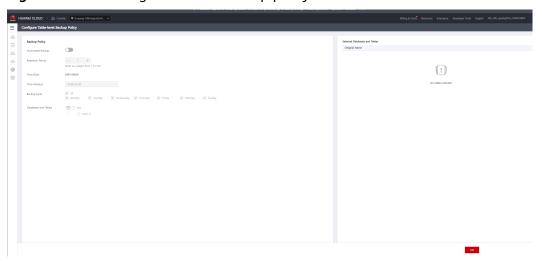
- Click View Tables to view tables contained in the backup file.
- Click **Restore** in the **Operation** column to restore the backup to a new instance. For details, see **Restoring Data to a New Instance**.

----End

Disabling Table-level Backup Policy

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** On the **Instances** page, locate the instance that you want to disable the table-level backup policy for and click its name.
- **Step 4** In the navigation pane on the left, choose **Backups & Restorations**.
- **Step 5** On the displayed page, click **Configure Table-level Backup Policy**.
- **Step 6** On the displayed page, configure related parameters.

Figure 4-64 Disabling table-level backup policy



■ NOTE

After automated backup is disabled, any table-level backup task in progress stops immediately. After table-level backup policy is disabled, all table-level backup files of the instance are retained. The retention duration depends on the **Retention Period** value that you specify when you enable table-level backup policy.

Step 7 Click OK.

----End

4.8.8 Viewing Incremental Backups

GeminiDB Cassandra API allows you to view incremental backup files and their size.

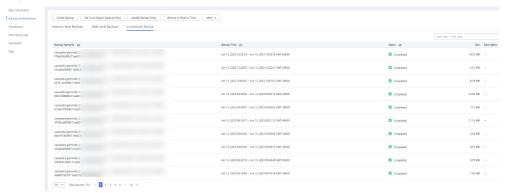
Precautions

- To use this function, contact customer service.
- You can view incremental backup files and their size of an instance only after you enable this function. Otherwise, no data is displayed.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** On the **Instances** page, locate the instance whose incremental backups you want to view and click its name.
- **Step 4** In the navigation pane on the left, choose **Backups & Restorations**.
- **Step 5** On the **Backups & Restorations** page, click **Incremental Backups**.

Figure 4-65 Incremental backup



Step 6 View incremental backups and their size.

----End

4.9 Data Restoration

4.9.1 Restoration Methods

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

Table 4-19 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.

4.9.2 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 Cassandra API**.
- **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.

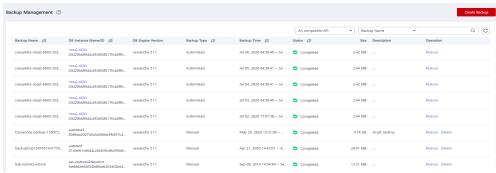
Figure 4-66 Backups and restorations



Method 2

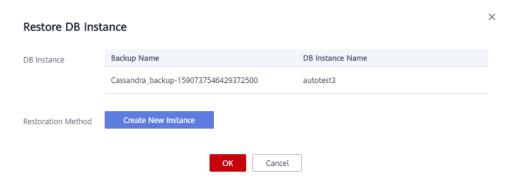
On the **Backups** page, locate the backup that you want to restore and click **Restore** in the **Operation** column.

Figure 4-67 Backup management



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

Figure 4-68 Restoring data to a new instance



- 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

4.9.3 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.
- Keep your account balance above zero so that backup data can be restored to a new instance.
- 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

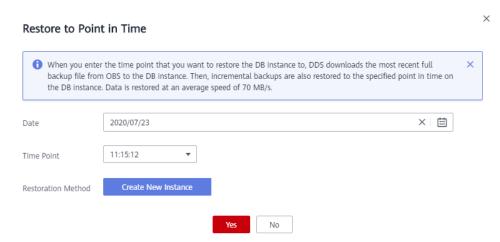
- **Step 1** Log in to the management console.
- Step 2 In the service list, choose Databases > GeminiDB Cassandra API.
- **Step 3** On the **Instances** page, click the instance that you want to perform a PITR for.
- **Step 4** In the navigation pane on the left, choose **Backups & Restorations**.
- **Step 5** On the **Backups & Restorations** page, click **Restore to Point in Time**.

Figure 4-69 Restoring data to a point in time



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

Figure 4-70 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 <u>Buying an</u> <u>Instance</u>.

----End

4.10 Parameter Template Management

4.10.1 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 Cassandra API.
- **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**.

X Create Parameter Template ★ Compatible API Cassandra MongoDB InfluxDB ★ DB Engine Version 3.11 ★ Parameter Template Name paramsGroup-2b45 (?) (?) Description Enter a parameter template description. You can create 100 more parameter templates. The parameter template quota is shared by all DB instances in a project. Cancel

Figure 4-71 Creating a parameter template

- **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

4.10.2 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 Cassandra API.
- **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 4-72 Modifying parameters in a parameter template



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

Figure 4-73 Preview Change



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

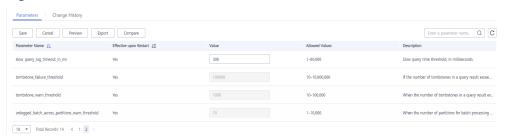
- The modifications take effect only after you apply the parameter template to DB instances. For details, see **Applying a Parameter Template**.
- The change history page displays only the modifications of the last seven days.

----End

Modifying Parameters of a Specified Instance

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **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 4-74 Parameters



- 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

4.10.3 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 Cassandra API.
- **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 4-75 Viewing change history of a customer parameter template



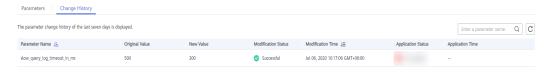
You can apply the parameter template to instances by referring to **Applying a Parameter Template**.

----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 Cassandra API.
- **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 4-76 Viewing parameter change history of an instance



----End

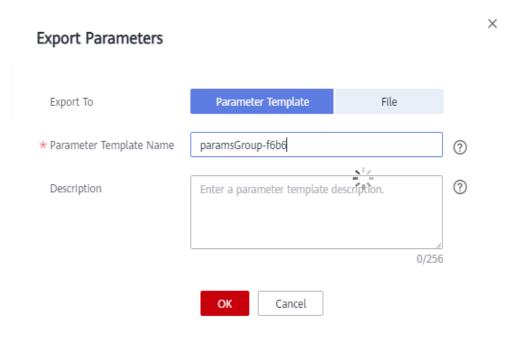
4.10.4 Exporting a Parameter Template

- 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 Cassandra API**.
- **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 4-77 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**.

- **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

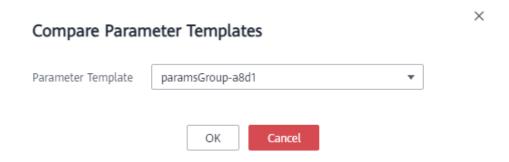
4.10.5 Comparing Parameter Templates

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 Cassandra API.
- **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 4-78 Comparing two parameter templates



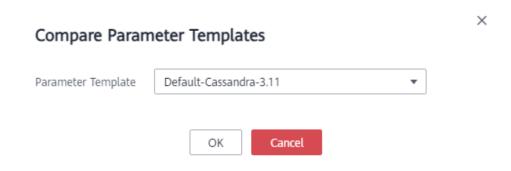
- If their parameters are different, the different parameter names and values are displayed.
- If their parameters are the same, no data is displayed.

----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 Cassandra API**.
- **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 4-79 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

4.10.6 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.

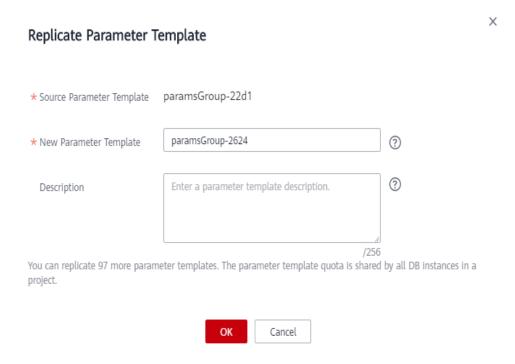
Procedure

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **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 4-80 Replicating a parameter template



- **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

4.10.7 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 Cassandra API**.
- **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

4.10.8 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 Cassandra API**.
- **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

4.10.9 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 Cassandra API.
- **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

4.10.10 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 Cassandra API**.
- **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 X 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

4.10.11 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 Cassandra API**.
- **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

4.11 Audit on Instance Operations

4.11.1 Key Operations Supported by CTS

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

Table 4-20 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	NoSQLExtendInstanceVo- lume
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	NoSQLUpdateConfigura- tions
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
Enabling SSL	instance	NoSQLSwitchSSL
Changing the security group of an instance	instance	NoSQLModifySecurityGroup
Configuring autoscaling	instance	NoSQLModifyAutoEnlarge- Policy
Creating a dual-active relationship	instance	NoSQLBuildBiactiveInstance
Exporting parameter template information for an instance	instance	NoSQLSaveConfigurations
Modifying the recycling policy	instance	NoSQLModifyRecyclePolicy

4.11.2 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 on 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 guery 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

4.12 Monitoring and Alarm Configuration

4.12.1 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

You can view metrics on instance nodes by referring to Viewing Metrics.

Table 4-21 GeminiDB Cassandra metrics

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
nosql0 05_disk _usage	Storage Space Usage	Storage space usage of the monitored object. Unit: Percent	0–100	GeminiDB Cassandra instances	1 minute
nosql0 06_disk _total_ size	Total Storage Space	Total storage space of the monitored object. Unit: GB	≥ 0	GeminiDB Cassandra instances	1 minute
nosql0 07_disk _used_s ize	Used Storage Space	Used storage space of the monitored object. Unit: GB	≥ 0	GeminiDB Cassandra instances	1 minute
nosql0 09_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
nosql0 10_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

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cassan dra001 _cpu_u sage	CPU Usage	CPU usage of an instance Unit: Percent	0–100	GeminiDB Cassandra instance nodes	1 minute
cassan dra002 _mem_ usage	Memory Usage	Memory usage of the instance Unit: Percent	0–100	GeminiDB Cassandra instance nodes	1 minute
cassan dra003 _bytes_ out	Network Output Throughpu t	Outgoing traffic in bytes per second Unit: byte/s	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra004 _bytes_ in	Network Input Throughpu t	Incoming traffic in bytes per second Unit: byte/s	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra014 _conne ctions	Active Node Connection s	Total number of connections attempting to connect to Cassandra instance nodes Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra015 _read_l atency	Average Read Latency	Average amount of time consumed by read requests Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cassan dra016 _write_ latency	Average Write Latency	Average amount of time consumed by write requests Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra037 _pendi ng_writ e	Suspended Write Tasks	Number of write tasks in waiting status Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra038 _pendi ng_rea d	Suspended Read Tasks	Number of read tasks in waiting status Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra044 _range _slice_l atency	Scan Duration	Average amount of time consumed by scan operations Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra049 _dropp ed_mut ation	Dropped Writes	Average number of dropped writes Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra052 _dropp ed_rea d	Dropped Reads	Average number of dropped reads Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra092 _load_i nfo	Data Volume on a Node	Data volume on a node Unit: byte	≥ 0	GeminiDB Cassandra instance nodes	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cassan dra093 _write_ count_l atency	Accumulat ed Write Requests	Total number of write requests initiated by a node Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra094 _write_ 1min_r ate	Average Write Rate in the Last Minute	Average write rate in the last minute Unit: count/s	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra095 _write_ p75_lat ency	p75 Write Latency	p75 write latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra096 _write_ p95_lat ency	p95 Write Latency	p95 write latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra097 _write_ p99_lat ency	p99 Write Latency	p99 write latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra098 _read_c ount_la tency	Accumulat ed Read Requests	Total number of read requests initiated by a node Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra099 _read_ 1min_r ate	Average Read Rate in the Last Minute	Average read rate in the last minute Unit: count/s	≥ 0	GeminiDB Cassandra instance nodes	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cassan dra100 _read_ p75_lat ency	p75 Read Latency	p75 read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra101 _read_ p95_lat ency	p95 Read Latency	p95 read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra102 _read_ p99_lat ency	p99 Read Latency	p99 read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra103 _range _slice_c ount_la tency	Accumulat ed Range Read Requests	Accumulate d range read requests Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra104 _range _slice_1 min_rat e	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
cassan dra105 _range _slice_p 75_late ncy	p75 Range Read Latency	p75 range read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra106 _range _slice_p 95_late ncy	p95 Range Read Latency	p95 range read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cassan dra107 _range _slice_p 99_late ncy	p99 Range Read Latency	p99 range read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra163 _write_ p999_l atency	p999 Write Latency	p999 write latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra164 _read_ p999_l atency	p999 Read Latency	p999 read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra165 _large_ partitio n_num	Big Keys	Number of big keys on the current node Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra166 _write_ max_la tency	Maximum Write Latency	Maximum write latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra167 _read_ max_la tency	Maximum Read Latency	Maximum read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassan dra168 _imbal ance_t able_n um	Tables with Uneven Data Distributio n	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

4.12.2 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.

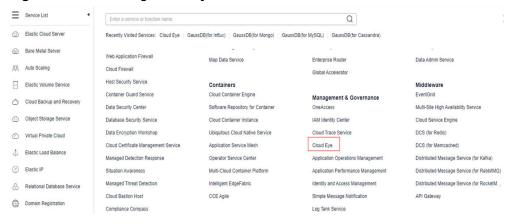
□ NOTE

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 & Governance, click Cloud Eye.

Figure 4-81 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 4-82 Creating an alarm rule



Step 5 Set alarm parameters.

1. Configure basic alarm information.

Figure 4-83 Configuring basic information for an alarm rule



Table 4-22 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 4-84 Configuring objects to be monitored



Table 4-23 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 .	-

Parameter	Description	Example Value
Dimension	Metric dimension of the alarm rule. Select Cassandra - Cassandra Nodes.	-
Monitoring Scope	Monitoring scope the alarm rule applies to. NOTE - If you select All resources, an alarm notification will be sent when any instance meets an alarm policy, and existing alarm rules will be automatically applied for newly purchased resources. - If you select Resource groups and any resource in the group meets the alarm policy, an alarm notification will be sent. - To specify Specific resources, click Select Specified Resources, select one or more resources, and click OK.	All resources
Group	This parameter is mandatory when Monitoring Scope is set to Resource groups.	-

3. Configure an alarm policy.

Figure 4-85 Configuring an alarm policy



Table 4-24 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	Policy for triggering an alarm. You can configure the threshold, consecutive periods, alarm interval, and alarm severity based on service requirements. - Metric Name: specifies the the metric that the alarm rule is created for. The following metrics are recommended: Storage Space Usage,	Take the CPU usage as an example. The alarm policy configured in Figure 4-85 indicates that a major alarm
	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 Storage Space. CPU Usage and Memory Usage,	notification will be sent to users every 10 minutes if the original CPU 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.	reaches 80% or above for three consecutive periods.
	For more metrics, see GeminiDB Cassandra Metrics .	
	 Alarm Severity: specifies the severity of the alarm. Valid values are Critical, Major, Minor, and Informational. 	
	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.	

4. Configure alarm notification information.

Figure 4-86 Configuring alarm notification information



Table 4-25 Parameter description

Parameter	Description	Example Value
Alarm Notification	Whether to notify users when alarms are triggered. Notifications can be sent by email, text message, or HTTP/HTTPS message. 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.	Enabled Alarm Notification.
Notification Recipient	Select Notification group or Topic subscription .	-
Notification Group	Notification group the alarm notification is to be sent to.	-
Notification Object	Object that receives alarm notifications. You can select the account contact or a topic. - 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. For details, see Creating a Topic and Adding Subscriptions.	-
Notification Window	Cloud Eye sends notifications only within the notification window specified in the alarm rule. For example, if Notification Window is set to 00:00-8:00 , Cloud Eye sends notifications only within 00:00-08:00.	-

Parameter	Description	Example Value
Trigger Condition	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.	-

5. Configure advanced settings.

Figure 4-87 Advanced settings

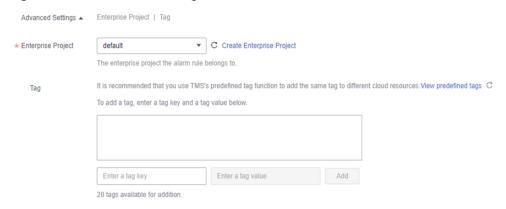


Table 4-26 Parameter description

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. For details about how to create an enterprise project, see Creating an Enterprise Project.	default
Tag	A tag is a key-value pair. Tags identify cloud resources so that you can easily categorize and search for your resources. You are advised to create predefined tags on TMS. For details about how to create predefined tags, see Creating Predefined Tags. - A key can contain a maximum of 128 characters, and a value can contain a maximum of 225 characters.	-
	A maximum of 20 tags can be added.	

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

4.12.3 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 Cassandra API**.
- **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 4-88 Viewing metrics



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

4.12.4 Event Monitoring

4.12.4.1 Introduction to Event Monitoring

Event monitoring provides event data reporting, query, and alarm reporting. You can create alarm rules for both system and custom events. When a specific event occurs, Cloud Eye generates and sends an alarm for you.

Key operations on GeminiDB Cassandra resources are monitored and recorded by Cloud Eye as events. Events include operations performed by specific users on specific resources, such as changing instance names and specifications.

Event monitoring provides an API for reporting custom events, which helps you collect and report abnormal events or important change events generated by services to Cloud Eye.

Event monitoring is enabled by default and allows you to view monitoring details of system events and custom events. For details about system events, see **Events Supported by Event Monitoring**.

4.12.4.2 Viewing Event Monitoring Data

Scenarios

Event monitoring provides event data reporting, query, and alarm reporting. You can create alarm rules for both system and custom events. When a specific event occurs, Cloud Eye generates and sends an alarm for you.

Event monitoring is enabled by default. You can view monitoring details about system events and custom events.

This topic describes how to view the event monitoring data.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** On the **Instances** page, locate the instance whose event monitoring data you want to view. In the **Node Information** area on the **Basic Information** page, click **View Metric** in the **Operation** column.
- **Step 4** Click to return to the Cloud Eye console.
- **Step 5** In the navigation pane on the left, choose **Event Monitoring**.

On the displayed **Event Monitoring** page, all system events generated in the last 24 hours are displayed by default.

You can also click **1h**, **3h**, **12h**, **1d**, **7d**, or **30d** to view events generated in different time periods.

Step 6 Locate an event and click **View Event** in the **Operation** column to view its details.

----End

4.12.4.3 Creating an Alarm Rule for Event Monitoring

Scenarios

This topic describes how to create an alarm rule for event monitoring.

Procedure

- **Step 1** Log in to the management console.
- Step 2 Click in the upper left corner of the page. Under Management & Governance, click Cloud Eye.
- **Step 3** In the navigation pane on the left, choose **Event Monitoring**.
- **Step 4** On the event list page, click **Create Alarm Rule** in the upper right corner.
- **Step 5** On the **Create Alarm Rule** page, configure the parameters.

Table 4-27 Parameter description

Parameter	Description
Name	Specifies the name of the alarm rule. The system generates a random name, but you can change it if needed.
Description	(Optional) Provides supplementary information about the alarm rule.
Enterprise Project	You can select an existing enterprise project or click Create Enterprise Project to create one.
Alarm Type	Specifies the alarm type corresponding to the alarm rule.
Event Type	Specifies the event type of the metric corresponding to the alarm rule.
Event Source	Specifies the service the event is generated for. Select GeminiDB.
Monitoring Scope	Specifies the monitoring scope for event monitoring.
Method	Specifies the event creation method.
Alarm Policy	Event Name indicates the instantaneous operations users performed on system resources, such as login and logout. For details about events supported by Event Monitoring, see Events Supported by Event Monitoring.
	You can select a trigger mode and alarm severity as needed.

Click to enable alarm notification. The validity period is 24 hours by default. If the topics you require are not displayed in the drop-down list, click **Create an SMN topic**.

Table 4-28 Alarm notification parameters

Parameter	Description
Alarm Notification	Specifies whether to notify users when alarms are triggered. Notifications can be sent by email, text message, or HTTP/ HTTPS message.
Notification Object	Specifies the object an alarm notification is to be sent to. You can select the account contact or a topic.
	Account contact is the mobile phone number and email address provided for registration.
	Topic is used to publish messages and subscribe to notifications. If the required topic is unavailable, create one first and add subscriptions to it. For details, see Creating a Topic and Adding Subscriptions.
Validity Period	Cloud Eye sends notifications only within the validity period specified in the alarm rule.
	If you set Validity Period to 08:00-20:00 , Cloud Eye sends notifications only within 08:00-20:00.
Trigger Condition	Specifies the condition for triggering the alarm notification.

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

----End

4.12.4.4 Events Supported by Event Monitoring

Table 4-29 Events Supported by Event Monitoring for GeminiDB

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
NoSQ L	Instance creation failure	NoSQL Createl nstance Failed	Maj or	The instance quota or underlying resources are insufficient.	Release the instances that are no longer used and try to provision them again, or submit a service ticket to adjust the quota.	Instan ces fail to be create d.
	Specificati ons change failure	NoSQL Resizel nstance Failed	Maj or	The underlying resources are insufficient.	Submit a service ticket to ask O&M personnel to coordinate resources, and then try again.	Servic es are interr upted.
	Node adding failure	NoSQL AddNo desFail ed	Maj or	The underlying resources are insufficient.	Submit a service ticket to ask O&M personnel to coordinate resources, delete the node that failed to be added, and add a new one.	None
	Node deletion failure	NoSQL Delete NodesF ailed	Maj or	Releasing underlying resources failed.	Delete the node again.	None

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Storage space scale-up failure	NoSQL ScaleU pStorag eFailed	Maj or	The underlying resources are insufficient.	Submit a service ticket to ask O&M personnel to coordinate resources, and then try again.	Servic es may be interr upted.
	Password resetting failure	NoSQL ResetPa ssword Failed	Maj or	Resetting the password times out.	Reset the password again.	None
	Parameter template change failure	NoSQL Updatel nstance Param GroupF ailed	Maj or	Changing a parameter template times out.	Change the parameter template again.	None
	Backup policy configurat ion failure	NoSQL SetBack upPolic yFailed	Maj or	The database connection is abnormal.	Configure the backup policy again.	None
	Manual backup creation failure	NoSQL Create Manual Backup Failed	Maj or	The backup files fail to be exported or uploaded.	Submit a service ticket to O&M personnel.	Data canno t be backe d up.
	Automate d backup creation failure	NoSQL CreateA utomat edBack upFaile d	Maj or	The backup files fail to be exported or uploaded.	Submit a service ticket to O&M personnel.	Data canno t be backe d up.
	Instance status abnormal	NoSQL FaultyD BInstan ce	Maj or	This event is a key alarm event and is reported when an instance is faulty due to a disaster or a server failure.	Submit a service ticket.	The datab ase servic e may be unava ilable.

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Instance status recovery	NoSQL DBInsta nceRec overed	Maj or	If a disaster occurs, NoSQL provides an HA tool to automatically or manually rectify the fault. After the fault is rectified, this event is reported.	No further action is required.	None
	Node status abnormal	NoSQL FaultyD BNode	Maj or	This event is a key alarm event and is reported when a database node is faulty due to a disaster or a server failure.	Check whether the database service is available and submit a service ticket.	The datab ase servic e may be unava ilable.
	Node status recovery	NoSQL DBNod eRecov ered	Maj or	If a disaster occurs, NoSQL provides an HA tool to automatically or manually rectify the fault. After the fault is rectified, this event is reported.	No further action is required.	None
	Primary/ standby switchove r or failover	NoSQL Primary Standb ySwitch ed	Maj or	This event is reported when a primary/ secondary switchover or a failover is triggered.	No further action is required.	None

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Occurrenc e of hotspot partitionin g keys	HotKey Occurs	Maj or	Hotspot data is stored in one partition because the primary key is improper. Improper application design causes frequent read and write operations on a key.	1. Choose a proper partition key. 2. Add service cache so that service applications read hotspot data from the cache first.	The servic e reque st succes s rate is affect ed, and the cluste r perfor manc e and stabili ty deteri orates .
	BigKey occurrenc e	BigKey Occurs	Maj or	The primary key design is improper. There are too many records or too much data in a single partition, causing load imbalance on nodes.	1. Choose a proper partition key. 2. Add a new partition key for hashing data.	As more and more data is stored in the partiti on, cluste r stabili ty deteri orates .

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Insufficien t storage space	NoSQL RiskyDa taDiskU sage	Maj or	The storage space is insufficient.	Scale up storage space. For details, see section "Scaling Up Storage Space" in the user guide of GeminiDB.	The instan ce is set to readonly and data canno t be writte n to the instan ce.
	Data disk expanded and being writable	NoSQL DataDi skUsag eRecov ered	Maj or	The data disk has been expanded and becomes writable.	No further action is required.	None
	Index creation failure	NoSQL Createl ndexFai led	Maj or	The service load exceeds what the instance specifications can take. In this case, creating indexes consumes more instance resources. As a result, the response is slow or even frame freezing occurs, and the creation times out.	Select matched instance specifications based on service load. Create indexes during off- peak hours. Create indexes in the background. Select indexes as required.	The index fails to be create d or is incom plete. Delet e the index and create a new one.

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Write speed decrease	NoSQL Stalling Occurs	Maj or	The write speed is close to the maximum write speed allowed by the cluster scale and instance specifications. As a result, the database flow control mechanism is triggered, and requests may fail.	1. Adjust the cluster scale or node specifications based on the maximum write rate of services. 2. Measure the maximum write rate of services.	The succes s rate of servic e reque sts is affect ed.
	Data write stopped	NoSQL Stoppin gOccur s	Maj or	The data write is too fast, reaching the maximum write capability allowed by the cluster scale and instance specifications. As a result, the database flow control mechanism is triggered, and requests may fail.	1. Change the cluster scale or node specifications based on the maximum write rate of services. 2. Measure the maximum write rate of services.	The succes s rate of servic e reque sts is affect ed.
	Database restart failure	NoSQL Restart DBFaile d	Maj or	The instance status is abnormal.	Submit a service ticket to O&M personnel.	The instan ce status may be abnor mal.

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Restoratio n to new instance failure	NoSQL Restore ToNewl nstance Failed	Maj or	The underlying resources are insufficient.	Submit a service ticket to ask O&M personnel to coordinate resources, and then add new nodes.	Data canno t be restor ed to a new instan ce.
	Restoratio n to existing instance failure	NoSQL Restore ToExistI nstance Failed	Maj or	The backup file fails to be downloaded or restored.	Submit a service ticket to O&M personnel.	The curren t instan ce may be unava ilable.
	Backup file deletion failure	NoSQL DeleteB ackupF ailed	Maj or	The backup files fail to be deleted from OBS.	Delete the backup files again.	None
	Failure to display slow query logs in plaintext	NoSQL SwitchS lowlog PlainTe xtFailed	Maj or	The DB API does not support this function.	Refer to the GeminiDB User Guide to check whether that the DB API supports the display of slow query logs in plaintext. Submit a service ticket to O&M personnel.	None

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	EIP binding failure	NoSQL BindEip Failed	Maj or	The node status is abnormal, an EIP has been bound to the node, or the EIP to be bound is invalid.	Check whether the node is normal and whether the EIP is valid.	The instan ce canno t be access ed from a public netwo rk.
	EIP unbinding failure	NoSQL Unbind EipFaile d	Maj or	The node status is abnormal or the EIP has been unbound from the node.	Check whether the node and EIP status are normal.	None
	Parameter modificati on failure	NoSQL Modify Parame terFaile d	Maj or	The parameter value is invalid.	Check whether the parameter value is within the valid range and submit a service ticket to O&M personnel.	None
	Parameter template applicatio n failure	NoSQL ApplyP aramet erGrou pFailed	Maj or	The instance status is abnormal. So, the parameter template cannot be applied.	Submit a service ticket to O&M personnel.	None
	Enabling or disabling SSL failure	NoSQL SwitchS SLFaile d	Maj or	Enabling or disabling SSL times out.	Try again or submit a service ticket. Do not change the connection mode.	The SSL conne ction mode canno t be chang ed.

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Too much data in a single row	LargeR owOcc urs	Maj or	If there is too much data in a single row, queries may time out, causing faults like OOM error.	1. Limit the write length of each column and row so that the key and value length of each row does not exceed the preset threshold. 2. Check whether there are abnormal writes or coding, causing large rows.	If there are too many record s in a single row, cluste r stabili ty will deteri orate as the data volum e increa ses.

4.13 Log Management

4.13.1 Slow Query Logs

GeminiDB Cassandra API allows you to view slow query logs of databases. The unit of the execution time is ms. You can identify the SQL statements that take a long time to execute and tune them based on slow query logs.

Viewing and Exporting Log Details

- **Step 1** Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- **Step 3** On the **Instances** page, locate the instance whose log details you want to view and click its name.
- **Step 4** In the navigation pane on the left, choose **Slow Query Logs**.
- **Step 5** On the **Slow Query Logs** page, set search criteria and click **Search** to view log information.
 - Select All nodes and view slow query logs of all nodes. Alternatively, select a specific node to view its slow query logs.

- View slow query logs of a specific node in different time ranges.
- View slow query logs of the following types of SQL statements:
 - SELECT
- **Step 6** On the **Log Details** page, click in the upper right corner of the log list to export log details.
 - You can view the exported CSV file to your local PC.
 - Up to 2,000 logs can be exported at a time.

----End

4.14 Enterprise Project

4.14.1 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*.

4.14.2 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 Cassandra API**.
- **Step 3** In the navigation pane on the left, choose **My Quotas** to view quota details of the current enterprise project.

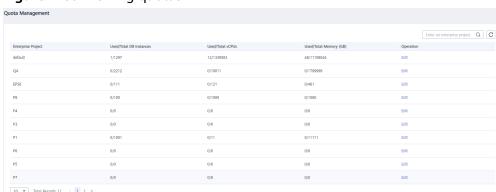


Figure 4-89 Viewing quotas

Table 4-30 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 Cassandra API**.
- **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 4-90 Modifying quotas

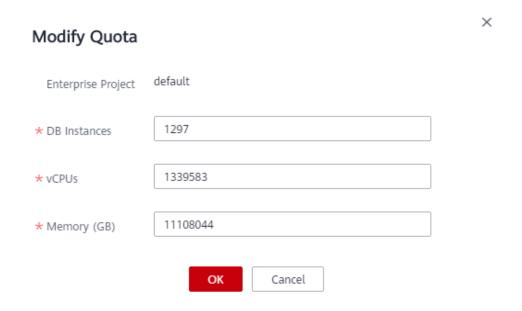


Table 4-31 Quota management

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

----End

4.15 Billing Management

4.15.1 Renewing Instances

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

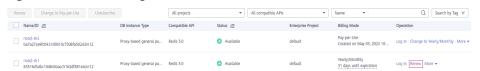
Precautions

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

Renewing a Yearly/Monthly Instance

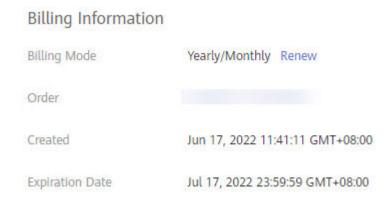
- **Step 1** Log in to the management console.
- Step 2 In the service list, choose Databases > GeminiDB Cassandra API.
- **Step 3** On the **Instances** page, locate the instance that you want to renew and click **Renew** in the **Operation** column.

Figure 4-91 Renewing an instance



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 4-92 Renewing an instance



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

Figure 4-93 Renewing instances



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

----End

4.15.2 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 Cassandra API.
- **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 4-94 Changing the billing mode from pay-per-use to yearly/monthly



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.

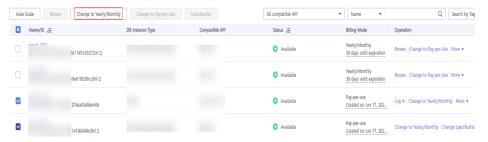
In the upper right corner of the DB instance list, click G to refresh the list. The instance status will become **Available** after the change is successful. The billing mode becomes to **Yearly/Monthly**.

----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 Cassandra API**.
- **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 4-95 Changing the billing mode of multiple instances



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 \square to refresh the list. The instance status will become **Available** after the change is successful. The billing mode changes to **Yearly/Monthly**.

----End

4.15.3 Changing the Billing Mode from Yearly/Monthly to Payper-Use

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

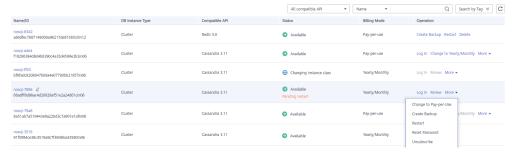
Precautions

- The billing mode of a yearly/monthly instance can only be changed to payper-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 Cassandra API**.
- **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 4-96 Changing yearly/monthly to pay-per-Use



- **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 payper-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**.

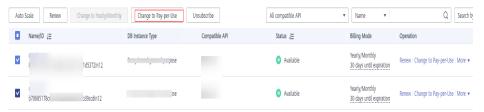
Step 7 In the displayed dialog box, click **Yes**.

----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 Cassandra API**.
- **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 4-97 Changing the billing mode of multiple instances to pay-per-use



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

----End

4.15.4 Unsubscribing from a Yearly/Monthly Instance

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

Precautions

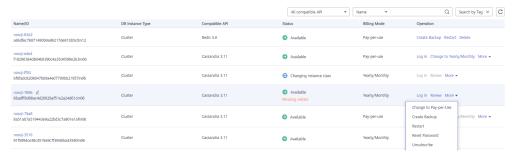
• Unsubscribed operations cannot be undone. Exercise caution when performing this operation. To retain data, create a manual backup before unsubscription. For details, see **Creating a Manual Backup**.

 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 Cassandra API**.
- **Step 3** On the **Instances** page, locate the instance you want to unsubscribe from and choose **More** > **Unsubscribe** in the **Operation** column.

Figure 4-98 Unsubscribing from a yearly/monthly instance



- **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

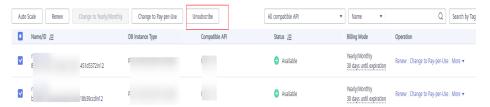
- 1. After an unsubscription request is submitted, resources and data will be deleted and cannot be retrieved.
- 2. If you want to retain data, complete a manual backup before submitting the unsubscription request.
- **Step 7** View the unsubscription result. After the instance order is successfully 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 Cassandra API.
- **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 4-99 Unsubscribing from multiple yearly/monthly instances



- **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 instances are successfully unsubscribed from, they are no longer displayed in the instance list any longer on the **Instances** page.

----End

5 Best Practices

5.1 Performance of GeminiDB Cassandra and On-Premises Open Source Cassandra Clusters

This section describes how the performance of an open-source Cassandra cluster compares to a GeminiDB Cassandra cluster. The test environment, test model, and test steps will all be described.

Test Environment

Open-source Cassandra test environment

Table 5-1 Test environment description

Name	Open-source Cassandra Cluster			
Version	3.11.5			
Nodes	3			
os	CentOS 7.4			
ECS Specification s	 General computing-plus 4 vCPUs 16 GB General computing-plus 8 vCPUs 32 GB General computing-plus 16 vCPUs 64 GB General computing-plus 32 vCPUs 128 GB 			

• GeminiDB Cassandra test environment

Table 5-2 Test environment description

Name	GeminiDB Cassandra Cluster			
Region	CN-Hong Kong			

Nodes	3				
AZ	AZ 3				
Version	3.11				
Instance Specifications	4 vCPUs 16 GB8 vCPUs 32 GB				
	16 vCPUs 64 GB32 vCPUs 128 GB				

Load Test Tool Environment

Load test tool specifications

Table 5-3 Specifications description

Name	est client ECS			
vCPUs	16			
Memory	64 GB			
OS	CentOS 7.4			

• Load test tool information

Table 5-4 Load test tool information

Test Tool	YCSB
Version	0.12.0
Download Address	https://github.com/brianfrankcooper/YCSB curl -Olocation https://github.com/ brianfrankcooper/YCSB/releases/download/0.12.0/ ycsb-0.12.0.tar.gz

Testing Models

Table 5-5 Testing models

Service Model	Description
_read95_update5	95% read and 5% update
_update50_read50	50% update and 50% read
_read65_update25_insert10	65% read, 25% update, and 10% write

Service Model	Description		
_insert90_read10	90% write and 10% read		

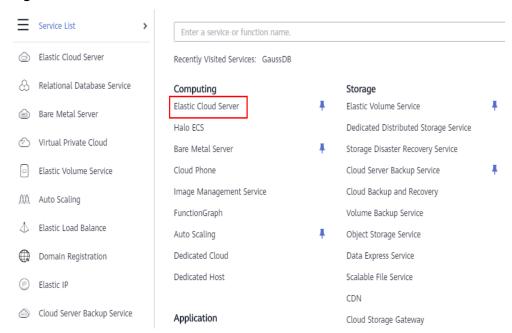
Test Procedure

Testing open-source Cassandra

Step 1 Purchase an ECS.

- 1. Log in to the management console.
- 2. Choose **Computing** > **Elastic Cloud Server**.

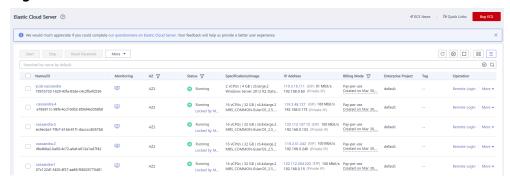
Figure 5-1 Elastic Cloud Server



- 3. Click **Buy ECS** in the upper right corner of the page and configure related parameters as follows:
 - Region: CN-Hong Kong
 - AZ: AZ3
 - Specifications: General computing-plus | c6.xlarge.4
 - Image: Public image and CentOS 7.6 64bit(40 GB)
 - Data Disk: Ultra-high I/O and 200 GB
 - Network: Select a VPC and subnet.
 - Other parameters: Set other parameters as needed. You can ignore optional parameters.
- 4. Repeat the preceding steps to create five ECSs named Cassandra-1 (192.168.0.15), Cassandra-2 (192.168.0.240), Cassandra-3 (192.168.0.153), Cassandra-4 (192.168.0.175) and ycsb-Cassandra (192.168.0.60).

ECSs Cassandra-1, Cassandra-2, and Cassandra-3 are for initializing Cassandra clusters. ECS Cassandra-4 is for capacity expansion. ECS ycsb-Cassandra serves as the load test server.

Figure 5-2 ECS details



5. After those ECSs are created, log in to them using the remote login option provided on the management console.

Figure 5-3 Logging in to an ECS



6. Install Java Runtime Environment:

yum install jre

- 7. Install the Cassandra service and create a data directory.
 - Download the Cassandra installation package:
 wget https://archive.apache.org/dist/cassandra/3.11.5/apache-cassandra-3.11.5-bin.tar.gz
 - b. Decompress the installation package:
 - tar -zxvf apache-Cassandra-3.11.5-bin.tar.gz -C /root/
 - c. Change the installation directory:
 - mv /root/apache-Cassandra-3.11.5 /usr/local/Cassandra
 - d. Configure environment variables:
 - echo "export PATH=/usr/local/Cassandra/bin:\$PATH" >> /etc/profile
 - e. Apply the variables:
 - source /etc/profile
 - f. Create a data directory:
 - mkdir /data
 - g. Confirm that the installation was successful.
 - cqlsh

Figure 5-4 Successful installation



Step 2 Configure an open-source Cassandra cluster.

- Log in to ECSs Cassandra-1, Cassandra-2, and Cassandra-3.
- Go to the /usr/local/Cassandra/conf directory and modify the Cassandratopology.properties file as follows:
 - Comment out the content in the area marked by ① in Figure 5-5.
 - Add the content in the area marked by ② in Figure 5-5.

Figure 5-5 Modifying the configuration file

```
Cassandra Node IP=Data Center:Rack
#192.168.1.100=DC1:RAC1
#192.168.2.200=DC2:RAC2
#10.0.0.10=DC1:RAC1
#10.0.0.11=DC1:RAC1
#10.0.0.12=DC1:RAC2
#10.20.114.10=DC2:RAC1
#10.20.114.11=DC2:RAC1
#10.21.119.13=DC3:RAC1
#10.21.119.10=DC3:RAC1
#10.0.0.13=DC1:RAC2
#10.21.119.14=DC3:RAC2
#10.20.114.15=DC2:RAC2
192.168.0.153=MY
192.168.0.240=M
192.168.0.15=MYI
# default for unknown nodes
default=DC1:r1
```

◯ NOTE

The Cassandra-topology.properties configuration files of Cassandra-1, Cassandra-2, and Cassandra-3 must be the same.

Modify the **Cassandra.yaml** file as follows:

```
data_file_directories:
- /data
commitlog_directory: /usr/local/Cassandra/commitlog
saved_caches_directory: /usr/local/Cassandra/saved_caches
seed provider:
# Addresses of hosts that are deemed contact points.
# Cassandra nodes use this list of hosts to find each other and learn
```

```
# the topology of the ring. You must change this if you are running
# multiple nodes!
- class_name: org.apache.Cassandra.locator.SimpleSeedProvider
parameters:
# seeds is actually a comma-delimited list of addresses.
# Ex: "<ip1>,<ip2>,<ip3>"
- seeds: "192.168.0.153,192.168.0.240,192.168.0.15" ## Enter IP addresses of the three nodes in the cluster.
listen_address: 192.168.0.153 #IP address of each node
rpc_address: 192.168.0.153 #IP address of each node
```

4. Run the following command on Cassandra-1, Cassandra-2, and Cassandra-3 to start the Cassandra cluster:

Cassandra -R &

Step 3 Add nodes to the open-source Cassandra cluster.

- 1. Log in to Cassandra-4.
- 2. Go to the /usr/local/cassandra/conf directory and edit the Cassandra-topology.properties file as follows:
 - Comment out the content in the area marked by ① in Figure 5-6.
 - Add the content in the area marked by ② in Figure 5-6.

Figure 5-6 Editing the configuration file

```
# Cassandra Node IP=Data Center:Rack
#192.168.1.100=DC1:RAC1
#192.168.2.200=DC2:RAC2
#10.0.0.10=DC1:RAC1
#10.0.0.11=DC1:RAC1
#10.0.0.12=DC1:RAC2
#10.20.114.10=DC2:RAC1
#10.20.114.11=DC2:RAC1
#10.21.119.13=DC3:RAC1
#10.21.119.10=DC3:RAC1
#10.0.0.13=DC1:RAC2
#10.21.119.14=DC3:RAC2
#10.20.114.15=DC2:RAC2
192.168.0.153=|
192.168.0.240=
192.168.0.15=
92.168.0.175=
```

3. Modify the Cassandra.yaml file as follows:

```
data_file_directories:
- /data
commitlog_directory: /usr/local/Cassandra/commitlog
```

```
saved_caches_directory: /usr/local/Cassandra/saved_caches
seed_provider:
# Addresses of hosts that are deemed contact points.
# Cassandra nodes use this list of hosts to find each other and learn
# the topology of the ring. You must change this if you are running
# multiple nodes!
- class_name: org.apache.Cassandra.locator.SimpleSeedProvider
parameters:
# seeds is actually a comma-delimited list of addresses.
# Ex: "<ip1>,<ip2>,<ip3>"
- seeds: "192.168.0.153,192.168.0.240,192.168.0.15" ## Enter IP addresses of the three seed nodes in
the cluster, which must be the same as the values entered in step 1.
                               #IP address of each node
listen address: 192.168.0.175
rpc_address: 192.168.0.175
                                #IP address of each node
```

- 4. Log in to Cassandra-1.
- 5. Stop compaction on all nodes:

nodetool disableautocompaction

6. Stop the ongoing compaction task:

nodetool stop COMPACTION

7. Limit migration traffic of the node:

nodetool setstreamthroughput 32

In the preceding command, the value of **nodetool setstreamthroughput 32** is set to **32 MB/s** to reduce the impact of migration on services.

- 8. Log in to Cassandra-4.
- 9. Start the Cassandra service:

Cassandra -R &

- 10. Log in to Cassandra-1.
- 11. During the scaling, run the following command every 30 seconds:

nodetool status

If the status of **Cassandra-4** is **UJ**, data is being migrated. The migration is complete when the status changes to **UN**.

Figure 5-7 Node statuses

----End

Testing GeminiDB Cassandra

Step 1 Purchase a GeminiDB Cassandra cluster.

- 1. Log in to the management console.
- 2. Choose Databases > GeminiDB.
- 3. Click **Buy DB Instance** in the upper right corner of the page and set required parameters as follows:

Region: CN-Hong Kong

- Compatible API: Cassandra

- Specifications: 4 vCPUs | 16 GB

Storage Space: 200 GB

- **Nodes**: Enter **3**.

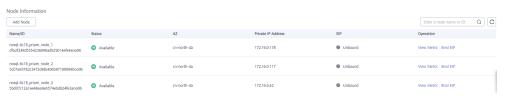
- **VPC**: The same as that of the purchased ECS.

- **Security Group**: The same as that of the purchased ECS.

Step 2 Add nodes to the GeminiDB Cassandra cluster.

- 1. Log in to the management console.
- 2. Choose Databases > GeminiDB.
- 3. Select an existing GeminiDB Cassandra instance.
- 4. Click the instance name to enter the **Basic Information** page.
- 5. In the **Node Information** area on the **Basic Information** page, click **Add Node**.

Figure 5-8 Basic information



6. On the displayed page, click + on the right of field Add Nodes.

Figure 5-9 Adding nodes



- 7. Wait until the nodes are added.
- 8. View the change of QPS during the scale-out process.

Figure 5-10 QPS changes



During the scale-out process, the QPS of the GeminiDB Cassandra instance decreases slightly for about 10 seconds, which almost has no effect on services. The whole scaling process takes about 10 minutes.

After the scale-out is complete, you can analyze test data.

----End

Test Results

Performance results

Table 5-6 Performance data

qps_av g Statisti cs	Nod e Clas s	Conc urren t Threa ds of the Clien t	Data Volu me to Be Prep ared	_read95 _update 5	_update5 0_read50	_read65_ update2 5_insert 10	_insert 90_read 10
Open- source Cassan dra	4 vCP Us 6 GB	32	50	2884	5068	8484	10694
cluster	8 vCP Us 32 GB	64	100	2796	2904	5180	7854
	16 vCP Us 64 GB	128	200	5896	14776	14304	15707

qps_av g Statisti cs	Nod e Clas s	Conc urren t Threa ds of the Clien t	Data Volu me to Be Prep ared	_read95 _update 5	_update5 0_read50	_read65_ update2 5_insert 10	_insert 90_read 10
	32 vCP Us 128 GB	256	400	8964	22284	19592	22344
Gemini DB Cassan dra	4 vCP Us 6 GB	32	50	8439	10565	9468	23830
cluster perfor mance data	8 vCP Us 32 GB	64	100	24090	24970	21716	44548
	16 vCP Us 64 GB	128	200	48985	51335	43557	67290
	32 vCP Us 128 GB	256	400	91280	85748	74313	111540
Perfor mance compar ison	4 vCP Us 6 GB	32	50	2.93	2.08	1.12	2.23
betwee n Gemini DB Cassan dra and open- source Cassan dra	8 vCP Us 32 GB	64	100	8.62	8.60	4.19	5.67
	16 vCP Us 64 GB	128	200	8.31	3.47	3.05	4.28

qps_av g Statisti cs	Nod e Clas s	Conc urren t Threa ds of the Clien t	Data Volu me to Be Prep ared	_read95 _update 5	_update5 0_read50	_read65_ update2 5_insert 10	_insert 90_read 10
	32 vCP Us 128 GB	256	400	10.18	3.85	3.79	4.99

Test Conclusion

- a. The GeminiDB Cassandra cluster performs ten times better than the open-source Cassandra cluster in terms of read latency.
- b. GeminiDB Cassandra cluster gives you basically the same write performance as the open-source cluster.
- c. Adding nodes slightly affects both the GeminiDB Cassandra and opensource clusters.
 - The scale-out for GeminiDB Cassandra is fast and only affects services briefly (10s). You do not need to change parameters, and the scale-out process takes 10 minutes.
 - For an open-source Cassandra cluster, the time needed for adding nodes depends on the data volume and parameter settings, and the impact on performance varies. In this test, the scale-out took more than 30 minutes when the preset data size was 50 GB.
 - Calculation formula: Highest migration speed = (nodetool setstreamthroughput 32 value, 200 Mbit/s by default) x Original nodes

In this test, the highest migration speed = 32 Mbit/s x 3 = 12 MB/s = 720 MB/min = 0.703 GB/min. So, the time needed for migrating 50 GB of data in this scenario was 71.1 minutes (50/0.703).

5.2 Buying and Connecting to a GeminiDB Cassandra Instance

This section describes how to buy a GeminiDB Cassandra instance and uses a Linux ECS as an example to describe how to connect to the instance over a private network.

- Step 1: Buy a GeminiDB Cassandra Instance
- Step 2: Buy an Instance
- Step 3: Connect to the GeminiDB Cassandra Instance

Basic Syntax

Step 1: Buy a GeminiDB Cassandra Instance

- Step 1 Log in to the management console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB Cassandra API**.
- Step 3 On the Instances page, click Buy DB Instance.
- **Step 4** Click **Buy DB Instance**, select a billing mode, and configure instance parameters. Then, click **Next** and complete subsequent operations.

Figure 5-11 Basic information

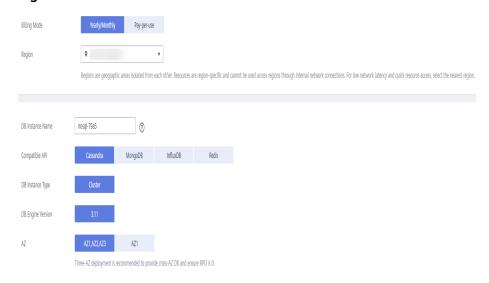
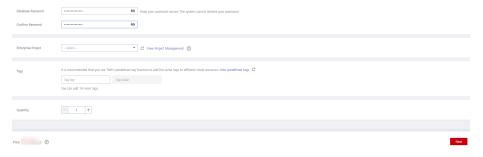


Figure 5-12 Setting a password



Step 5 View the purchased GeminiDB Cassandra instance.

Figure 5-13 Successful purchase



----End

Step 2: Buy an Instance

- Step 1 Log in to the management console.
- **Step 2** In the service list, choose **Compute** > **Elastic Cloud Server**. On the Elastic Cloud Server console, click **Buy ECS**.

Figure 5-14 Logging in to the ECS console



Step 3 Configure basic settings and click **Next: Configure Network**. Make sure that the ECS is in the same region, AZ, VPC, and security group as the GeminiDB Cassandra instance you created.

Figure 5-15 Basic settings

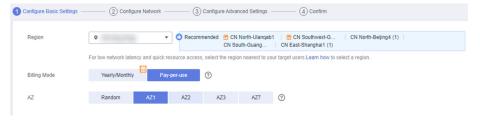


Figure 5-16 Selecting specifications

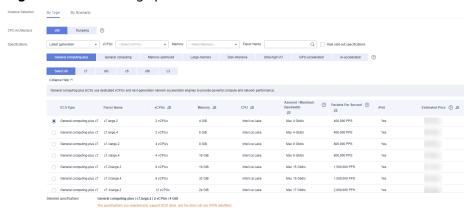


Figure 5-17 Selecting an image



- **Step 4** Configure the ECS network and click **Next: Configure Advanced Settings**. Make sure that the ECS is in the same VPC and security group as the GeminiDB Cassandra instance.
 - If security group rules allow access from the ECS, you can connect to the instance using the ECS.
 - If the security group rules do not allow access from the ECS, add an inbound rule to the security group.

Figure 5-18 Network settings

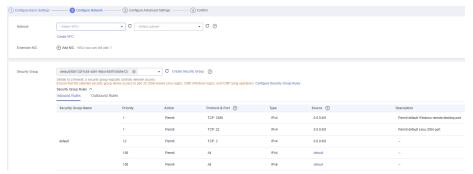
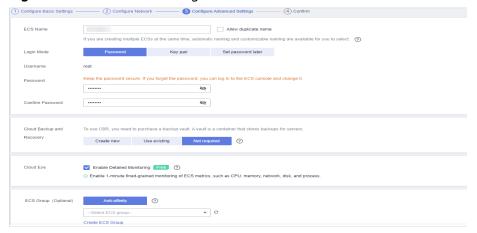


Figure 5-19 Selecting an EIP



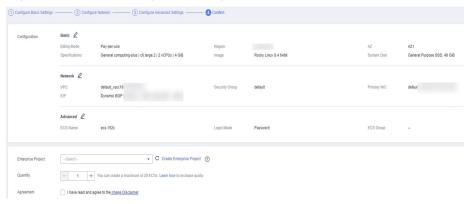
Step 5 Configure a password for the ECS and click **Next: Confirm**.

Figure 5-20 Advanced settings



Step 6 Confirm the configurations and click **Submit**.

Figure 5-21 Confirming the configurations



Step 7 View the purchased ECS.

----End

Step 3: Connect to the GeminiDB Cassandra Instance

Step 1 On the ECS console, log in to the ECS using the remote login option.

Figure 5-22 Remote login



Step 2 Enter the username and password of the ECS.

Figure 5-23 Entering the username and password

```
Rocky Linux 8.4 (Green Obsidian)
Kernel 4.18.0-37

Hint: Num Lock on
ecs-fd82 login: root
Password:
Last failed login: Tue May 30 13:53:07 CST 2023 from 114.116.222.88 on ssh:notty
There were 10 failed login attempts since the last successful login.

Welcome to Huawei Cloud Service

[root@ecs-fd82~]#_
```

Step 3 Download the Cassandra installation package and upload it to the ECS.

Method 1:

wget https://dbs-download.obs.cn-north-1.myhuaweicloud.com/nosql/ Cassandra_cqlsh_x86_64.zip

Method 2:

Download the **Cassandra client** installation package using your browser and upload it to the ECS.

Step 4 Decompress the client package.

unzip Cassandra_cqlsh_x86_64.zip

Step 5 Make the files executable:

chmod +x *

Step 6 Connect to the GeminiDB Cassandra instance in the directory where the cqlsh is located.

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

Example:

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

Table 5-7 Required description

Parameter	Description
<db_host></db_host>	The private IP address of the instance to be accessed. To obtain this IP address, go to the Instances page, locate the instance, and click its name. The IP address can be found in the Private IP Address field under Node Information on the Basic Information page.
	If the GeminiDB Cassandra instance you purchased has multiple nodes, select the private IP address of any node.

Parameter	Description
<db_port></db_port>	The port used to access the instance. The default port number is 8635. Set this parameter based on service requirements.
	Click the instance name to go to the Basic Information page and obtain the port number in the Network Information area.
<db_user></db_user>	Username of the instance administrator. The default value is rwuser .

Step 7 If information similar to the following is displayed, the connection was successful. rwuser@cqlsh>

----End

Basic Syntax

- Keyspace syntax
 - Create a keyspace.

Example:

CREATE KEYSPACE IF NOT EXISTS nosql WITH replication = {'class': 'SimpleStrategy', 'replication_factor': '3'};

In this example, **keyspace_name** is set to **nosql**, **class** to **SimpleStrategy**, and **replication_factor** to **3**. GeminiDB Cassandra provides strong consistency and 3 data copies for each GeminiDB Cassandra instance no matter how many copies you set.

Run DESC <keyspace_name> to verify the creation results.

Figure 5-24 Verifying the creation results

```
rwuser@cqlsh> CREATE KEYSPACE IF NOT EXISTS nosql WITH replication = {'class': 'SimpleStrategy', 'replication factor': '3'};
rwuser@cqlsh> DESC nosql;
CREATE KEYSPACE nosql WITH replication = {'class': 'SimpleStrategy', 'replication_factor': '3'} AND d
urable writes = true;
```

Run use <keyspace_name> to switch to the keyspace you created.

Figure 5-25 Switching the keyspace

```
rwuser@cqlsh> use nosql;
rwuser@cqlsh:nosql> |
```

Run DROP KEYSPACE < keyspace_name > to delete the keyspace you created.

Figure 5-26 Deleting the keyspace

```
rwuser@cqlsh> use nosql;
rwuser@cqlsh:nosql> DROP KEYSPACE nosql;
rwuser@cqlsh:nosql> DESC nosql;
Keyspace 'nosql' not found.
rwuser@cqlsh:nosql> ■
```

- Table syntax
 - Create a table.

Example:

CREATE TABLE nosql_table(user_id int, age int, user_name text, PRIMARY KEY(user_id));

nosql_table is a table name defined by the following three columns: **user_id**, **age**, and **user_name**. **user_id** is the user ID of the INT type. **age** is the user age of the INT type. **user_name** is the username of the TEXT type. The primary key is **user_id**.

- Run **DESC** < table_name > to verify the creation results.

Figure 5-27 Verifying the creation results

```
rwuser@cqlsh:nosql> CREATE TABLE nosql_table(user_id int, age int, user_name text, PRIMARY KEY(user_id
));
rwuser@cqlsh:nosql> DESC nosql_table;

CREATE TABLE nosql.nosql_table (
    user_id int PRIMARY KEY,
    age int,
    user_name text
) WITH bloom_filter_fp_chance = 0.01
    AND caching = {'keys': 'ALL', 'rows_per_partition': 'NONE'}
    AND comment = ''
    AND comment = ''
    AND compaction = {'class': 'org.apache.cassandra.db.compaction.SizeTieredCompactionStrategy', 'max
    threshold': '32', 'man threshold': '4'}
    AND compression = {'chunk_length_in_kb': '64', 'class': 'org.apache.cassandra.io.compress.LZ4Compr
essor')
    AND crc_check_chance = 1.0
    AND dcfault_time_to_live = 0
    AND gc_grace_seconds = 864000
    AND gc_grace_seconds = 864000
    AND max_index_interval = 2048
    AND memtable_flush_period_in_ms = 0
    AND min_index_interval = 128
    AND mead_repair_chance = 0.0
    AND speculative_retry = '99PERCENTILE';
```

- Insert data into the table, for example,

INSERT INTO nosql_table (user_id, age, user_name) VALUES (1, 10, 'user1');

INSERT INTO nosql_table (user_id, age, user_name) VALUES (2, 20, 'user2');

INSERT INTO nosql_table (user_id, age, user_name) VALUES (3, 30, 'user3');

Run SELECT * FROM to query table data.

Figure 5-28 Querying table data

```
rwuser@cqlsh:nosql> INSERT INTO nosql_table (user_id, age, user_name) VALUES (1, 10, 'user1');INSERT I
NTO nosql_table (user_id, age, user_name) VALUES (2, 20, 'user2');INSERT INTO nosql_table (user_id, age, user_name) VALUES (3, 30, 'user3');
rwuser@cqlsh:nosql> SELECT * FROM nosql_table;

user_id | age | user_name

1 | 10 | user1
2 | 20 | user2
3 | 30 | user3

(3 rows)
rwuser@cqlsh:nosql> ■
```

Add a column to the table, for example,

ALTER TABLE nosql_table ADD gender text;

- Update data in a table of a keyspace, for example,

```
UPDATE nosql.nosql_table SET gender = 'male' WHERE user_id = 1;
UPDATE nosql.nosql_table SET gender = 'male' WHERE user_id = 2;
UPDATE nosql.nosql_table SET gender = 'female' WHERE user_id = 3;
```

Figure 5-29 Viewing the update results

- Delete data from a table in a keyspace, for example,

Delete age data of the user whose ID is 1.

DELETE age FROM nosql.nosql_table WHERE user_id=1;

Figure 5-30 Deleting a data record

Delete the entire record of the user whose ID is 2.

DELETE FROM nosql.nosql_table WHERE user_id=2;

Figure 5-31 Deleting the entire record

- Delete an entire table, for example,

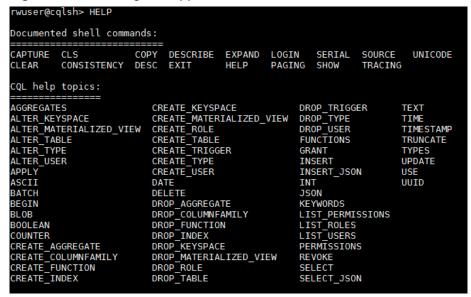
DROP TABLE nosql.nosql table;

Figure 5-32 Deleting an entire table

```
rwuser@cqlsh:nosql> DROP TABLE nosql.nosql_table;
rwuser@cqlsh:nosql> DESC nosql_table;
'nosql_table' not found in keyspace 'nosql'
rwuser@cqlsh:nosql>
```

- HELP command
 - Run the HELP command to view all supported commands.

Figure 5-33 Viewing all supported commands



- **HELP** *<COMMAND>* is used to query the usage of a command.

Example: **HELP DESC**

6 Performance White Paper

6.1 Test Method

This section describes performance testing of GeminiDB Cassandra instances, including the test environment, procedure, and results.

Test Environment

- Region: AP-Singapore
- AZ: AZ1, AZ2, and AZ3 (three-AZ deployment)
- Elastic Cloud Server (ECS): h3.4xlarge.2 with 16 vCPUs, 32 GB of memory, and CentOS 7.5 64-bit image
- Nodes per instance: 3
- Instance specifications: All specifications described in Table 6-1

Table 6-1 Instance specifications

No.	Specifications
Cluster 1	4 vCPUs 16 GB
Cluster 2	8 vCPUs 32 GB
Cluster 3	16 vCPUs 64 GB
Cluster 4	32 vCPUs 128 GB

Test Tool

YCSB is an open-source tool for testing performance of databases. In this test, YCSB 0.15.0 is used. RoundRobinPolicy is used for load balancing.

For details on how to use this tool, see YCSB.

Test Metrics

Operations per Second (OPS): operations executed by a database per second

Test Procedure

1. Configure the **workload** file.

Set values for fields **readproportion**, **insertproportion**, **updateproportion**, **scanproportion**, and **readmodifywriteproportion** in the file by referring to **Table 6-2**.

Set a value for field **recordcount** in the file by referring to **Table 6-3**.

2. Use workload-insert-mostly as an example. Run the following command to prepare test data:

sh bin/ycsb.sh load cassandra-cql -P workloads/workload-insert-mostly -p "hosts=\${ContactPoints}" -p "port=\${port}" -p "cassandra.username=\$ {username}" -p "cassandra.password=\${password}" -p operationcount=400000000 -p recordcount=400000000 -p exportfile=./ data load.exp -threads \${threadNum} -s > data load.log 2>&1 &

3. Use workload-insert-mostly as an example. Run the following command to test performance:

sh bin/ycsb.sh run cassandra-cql -P workloads/workload-insert-mostly -p "hosts=\${ContactPoints}" -p "port=\${port}" -p "cassandra.username=\$ {username}" -p "cassandra.password=\${password}" -p operationcount=9000000 -p recordcount=9000000 -p maxexecutiontime=3600 -p exportfile=./workload-insert-mostly.exp - threads \${threadNum} -s > workload-insert-mostly_run.log 2>&1 &

Test Models

Workload model

Table 6-2 Workload models

Workload Model	Description
workload-read-mostly	95% read, 5% update
workload-read-write- combination	50% update, 50% read
workload-read-modify-write	50% read, 50% read-modify-write
workload-mixed-operational- analytical	65% read, 25% update, 10% insert
workload-insert-mostly	90% insert, 10% read

 Data model fieldlength = 100, fieldcount = 10

Preset data volume

Different preset data volumes were used to test performance of instances of each type of specifications.

The following table describes the preset data volumes.

Table 6-3 Preset data volumes

No.	Specifications	Preset Data Volume
Cluster 1	4 vCPUs 16 GB	50 GB
Cluster 2	8 vCPUs 32 GB	100 GB
Cluster 3	16 vCPUs 64 GB	200 GB
Cluster 4	32 vCPUs 128 GB	400 GB

6.2 Test Data

The OPS of instances of different specifications can be tested using different service models with the same preset data volume. For details, see the numbers in bold in **Table 6-4**.

Table 6-4 Test data

Node Specifications		4 vCPUs 16 GB	8 vCPUs 32 GB	16 vCPUs 64 GB	32 vCPUs 128 GB
Concurrent Client Requests		32	64	128	256
Preset D	ata Volume	50 GB	100 GB	200 GB	400 GB
Test Model	workload-read- mostly	15627	44612	90713	169037
	workload-read- write-combination	19565	46240	95065	158793
	workload-read- modify-write	11768	29488	59332	96964
	workload-mixed- operational- analytical	17534	40214	80661	137616
	workload-insert- mostly	23830	44548	67290	111540

- Operations per Second (OPS): operations executed by a database per second
- Test Model No.: indicates the test model sequence number. For details, see Table 6-2.

7 FAQs

7.1 Product Consulting

7.1.1 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.

7.1.2 What Is GeminiDB Cassandra Instance Availability?

The formula for calculating the instance availability is as follows:

DB instance availability = (1 - Failure duration/Total service duration) × 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.

7.2 Billing

7.2.1 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.

7.2.2 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.

7.3 Database Usage

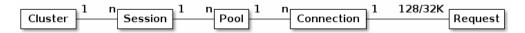
7.3.1 Why Does the Overall Instance Performance Deteriorate When QPS Increases After the Batch Size Is Decreased?

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.

7.3.2 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

7.3.3 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 = $16/2 \times (3/2 + 1) = 16 \text{ GB}$
 - 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 = $16/2 \times (4/2) = 16 \text{ GB}$
 - Maximum number of created tables = 16 x 1024 MB/768 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 x 1024 MB/768 MB = 32

For details about the mapping between the number of nodes (4 vCPUs, 16 GB) and the number of tables, see **Table 7-1**.

Table 7-	1 L	Jpper	limit	on	the	num	ber	of	tab	les
----------	-----	-------	-------	----	-----	-----	-----	----	-----	-----

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

- 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) x 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 7-2**.

Table 7-2 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

- 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) x 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 7-3**.

Table 7-3 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) x 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 7-4**.

Table 7-4 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

- 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) x 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: **Z00_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
 Z00_THROUGHPUT = 'small';
 - Medium throughput
 CREATE TABLE test2 (k int,p int,s int static,v int,PRIMARY KEY (k, p)) WITH
 Z00_THROUGHPUT = 'medium';
 - High throughput
 CREATE TABLE test3 (k int,p int,s int static,v int,PRIMARY KEY (k, p)) WITH
 Z00_THROUGHPUT = 'biq';
- 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 Z01_TABLE_TOKENS = 24;
- 3. Table parameters: Z00_BUFFER_SIZE and Z00_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.

- Z00_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 Z00_BUFFER_SIZE = {'default': 16};
- Z00_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 Z00_BUFFER_NUMBER = {'default': 3};

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 Z00 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 Z00_THROUGHPUT = 'small';

7.3.4 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 provided by Huawei 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 7-1 BigKey alarm



The alarm is described as follows:

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,
    accessstring text,
    moviename text,
    access_time timestamp,
    PRIMARY KEY (movieid, appid, uid, accessstring, 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.

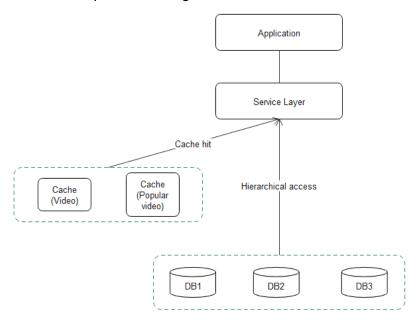
Create a table to store the short video information. Only public short video information is retained, and user information is not included. This ensures that the table does not generate large partition keys. Write the short video information to the table.

```
CREATE TABLE hotmovieaccess (
    movieid text,
    appid int,
    accessstring 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:

- 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 HotKey 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 7-2 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.

7.3.5 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 7-3 Example value

calculate Office Note in 1770 VISH and AS STRICT vd., cst, pd2 FROM to MedEE v2 > 2 MO vd IS NOT NULL MO pdc IS NOT NULL AND pd2 IS NOT NULL AND p

• 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 7-4 Querying the result

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 7-5 Mapping between views and base tables

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 base table and materialized view.

Figure 7-6 Querying the result

4. Delete the materialized view.

DROP MATERIALIZED VIEW person_adder;

Figure 7-7 Deleting a view

```
cqlsh:ks> DROP MATERIALIZED VIEW person_addr ; cqlsh:ks>
```

7.3.6 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 7-8 Querying the result

```
id | song_id | album | artist | song_order | title
```

If an index is created for the artist field, the hidden table structure is as follows.

Figure 7-9 Querying the result

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 7-10 Querying the result

2. Query records by indexed column.

Figure 7-11 Querying the result

Each table can have multiple indexes, but the write performance may be affected.

7.3.7 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.

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.

7.4 Database Connection

7.4.1 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.VehicleHistoryIoknob.main(VehicleHistoryToknob.scala:12)
at co.mega.tetris.analyzer.history.VehicleHistoryToknob.main(VehicleHistoryToknob.scala)

Caused by: java.uttil.MoSuchElementException: No value present
at java.uttil.Optional.get(Optional.java:135)
at com.datastax.spark.connector.rdd.partitioner.CassandraPartitionGenerator.SanonfunSwithSessionDoS1(CassandraConnector.scala:112)
at com.datastax.spark.connector.cql.CassandraConnector.SanonfunSwithSessionDoS1(CassandraConnector.scala:112)
at com.datastax.spark.connector.cql.CassandraConnector.scala:desconnector.scala:113)
at com.datastax.spark.connector.cql.CassandraConnector.withSessionDoC(cassandraConnector.scala:113)
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.CassandraToBleScanRDD.getPartitions(CassandraToBleScanRDD.scala:273)
at org.apache.spark.rdd.RDD.SanonfunSpartitionsS2(RDD.scala:276)
at scala.Option.getOrtIse(Option.scala:189)
at org.apache.spark.rdd.RDD.partitions(RDD.scala:272)
```

Configuration Details

The following shows the components and account details.

Component configuration details

Table 7-5 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.

7.4.2 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 7-12 Log information

```
Together the control of the control has together the correct datacenter name with DCAsareRoundRobinPolicy (if this is incorrect, please provide the correct datacenter name with DCAsareRoundRobinPolicy (if this is incorrect, please provide the correct datacenter name with DCAsareRoundRobinPolicy (if this is incorrect, please provide the correct datacenter name with DCAsareRoundRobinPolicy (if this is incorrect, please provide the correct datacenter name with DCAsareRoundRobinPolicy (if this is incorrect, please provide the correct datacenter name with DCAsareRoundRobinPolicy (if this is incorrect, please provide the correct datacenter name with DCAsareRoundRobinPolicy (if this is incorrect, please provide the correct datacenter name with DCAsareRoundRobinPolicy (if this is incorrect, please provide the correct datacenter name with DCAsareRoundRobinPolicy (if this is incorrect, please provide the correct datacenter one with DCAsareRoundRobinPolicy (if this is incorrect, please provide the correct datacenter) conditions the provide datacenter of the DCAsareRoundRobinPolicy (if this is incorrect, please provide the correct datacenter) conditions the provided datacenter of the DCAsareRoundRobinPolicy (if this is incorrect, please provide) conditions the provided datacenter of the DCAsareRoundRobinPolicy (if this is incorrect, please provide) conditions the provided datacenter one datacenter of the DCAsareRoundRobinPolicy (if this is incorrect, please provided datacenter) conditions the DCAsareRoundRobinPolicy (if this is incorrect, please provided datacenter) conditions the DCAsareRoundRobinPolicy (if this is incorrect, please provided datacenter) conditions the DCAsareRoundRobinPolicy (if this is incorrect, please provided datacenter) conditions the DCAsareRoundRobinPolicy (if this is incorrect, please provided datacenter) conditions the DCAsareRoundRobinPolicy (if this is incorrect, please provided datacenter) conditions the DCAsareRoundRobinPolicy (if this is incorrect, please provided datacenter) conditions the
```

Possible Cause

Java Driver was not used correctly, as shown in **Figure 7-13**. Do not use addContactPointsWitPorts when using Java Driver and do not map each IP address.

Figure 7-13 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();
```

7.4.3 How Can I Create and Connect to an ECS?

- 1. 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.
- 2. To connect to an ECS, see "Logging in to an ECS" *Getting Started with Elastic Cloud Server User Guide*.

7.4.4 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**.

7.5 Backup and Restoration

7.5.1 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.

7.6 Regions and AZs

7.6.1 What Is AZ and How Can I Select an AZ?

AZ

An AZ is a part of a physical region with its own independent power supply and network. An AZ is generally an independent physical equipment room, ensuring independence of the AZ.

Each region contains multiple AZs. If one AZ becomes faulty, the other AZs in the same region can continue to provide services.

AZs in the same region can communicate with each other over an intranet.

Selecting an AZ

You can deploy your instances in different AZs for high availability. If one of an AZ becomes faulty, databases in other AZs will not be affected. When selecting AZs:

- If only one AZ is available in a region, there is only one AZ in the region.
- The AZ of a purchased DB instance cannot be changed.
- The AZs in one region can communicate with each other over an intranet.

For more information, see **Regions and AZs**.

7.6.2 Can Different AZs Communicate with Each Other?

An AZ is a part of a physical region with its own independent power supply and network. An AZ is generally an independent physical equipment room, ensuring independence of the AZ.

Each region contains multiple AZs. If one AZ becomes faulty, the other AZs in the same region can continue to provide services normally.

By default, different AZs in the same VPC can communicate with each other through an internal network.

For more information, see Regions and AZs.

7.6.3 Can I Change the Region of a GeminiDB Cassandra Instance?

No. After an instance is created, its region cannot be changed.

7.7 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. For details, see **Resource Suspension and Release**.

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 a Pay-per-Use Instance**.
- To delete a yearly/monthly instance, see Unsubscribing from a Yearly/ Monthly Instance.

A Change History

Release Date	Description
2023-08-30	This issue is the thirty-second official release. Added Billing.
2023-07-18	This issue is the thirty-first official release. Added Getting Started with Common Practices.
2023-06-30	This issue is the thirtieth official release. Added Event Monitoring.
2023-05-30	This issue is the twenty-ninth official release. Added Viewing Incremental Backups.
2023-04-28	This issue is the twenty-eighth official release. Added Connecting to an Instance Using Go. Added Configuring an Incremental Backup Policy.
2023-03-31	This issue is the twenty-seventh official release. Added Updating the OS of an Instance.
2023-01-05	This issue is the twenty-sixth official release. Added the description of cross-region backup policies in Setting a Cross-Region Backup Policy.
2022-08-11	This issue is the twenty-fifth official release. Supported disk encryption in section Buying an Instance. Added FAQ Instance Freezing, Release, Deletion, and Unsubscription.
2022-07-07	This issue is the twenty-fourth official release. Added the description of patch installation in Upgrading a Minor Version.

Release Date	Description
2022-06-01	This issue is the twenty-third official release. Added the description of cross-region dual-active DR in Overview.
2022-02-15	This issue is the twenty-second official release. In Deleting Nodes , added the description of removing nodes from a yearly/monthly GeminiDB Cassandra instance. In Creating a DR Instance , added the description of creating a DR instance in a different subnet from the primary instance in the same VPC.
2022-01-27	This issue is the twenty-first official release. Added description of batch operations in Renewing Instances, Changing the Billing Mode from Pay-per-Use to Yearly/Monthly, Changing the Billing Mode from Yearly/Monthly to Pay-per-Use, and Unsubscribing from a Yearly/Monthly Instance.
2021-11-20	This issue is the twentieth official release. Added description of scaling operations in Buying an Instance and Configuring Autoscaling.
2021-09-27	This issue is the nineteenth official release. Supported node deletion in Deleting Nodes .
2021-06-30	This issue is the eighteenth official release. Modified the naming rule description of the GeminiDB Cassandra instance in Buying an Instance . Supported changing security groups of GeminiDB Cassandra instances in Changing a Security Group .
2021-06-23	This issue is the seventeenth official release. Added a method for connecting to GeminiDB Cassandra instances in Connecting to an Instance Using Java.
2021-06-15	This issue is the sixteenth official release. Adjusted the document structure.
2021-04-30	This issue is the fifteenth official release. Added the minimum storage capacity supported by instances of different specifications in Instance Specifications.
2021-04-15	This issue is the fourteenth official release. Adjusted the document structure. Modified the description of performance specifications on the purchase page.

Release Date	Description
2021-01-30	This issue is the thirteenth official release. Supported changing instance class, adding nodes to an instance, and scaling up storage capacity.
2020-12-30	This issue is the twelfth official release. Supported creating DR instances for GeminiDB Cassandra.
2020-10-30	 This issue is the tenth official release. Supported up to 20 tags per instance. Allowed users to select at least 1 GB each time they scale up. Supported scaling down the specifications of GeminiDB Cassandra instances.
2020-09-30	This issue is the ninth official release. Supported instance quota management. Supported storage space usage monitoring for GeminiDB Cassandra.
2020-08-30	This issue is the eighth official release. Supported the following functions of GeminiDB Cassandra: • Supported instance name modifications. • Supported instance recycle bin. • Supported exporting log files.
2020-07-30	This issue is the seventh official release. Supported point-in-time recovery of GeminiDB Cassandra.
2020-06-30	This issue is the sixth official release. Supported the following functions of GeminiDB Cassandra: Purchasing yearly/monthly instances Renewing yearly/monthly instances Changing the billing mode of instances from payper-use to yearly/monthly Changing the billing mode of instances from yearly/monthly to pay-per-use Unsubscribing from yearly/monthly instances SSL Cross-AZ disaster recovery
2020-04-30	This issue is the fifth official release. Added monitoring metrics.

Release Date	Description
2020-03-31	This issue is the fourth official release.
	Added Creating a Parameter Template.
	Added Exporting a Parameter Template.
	Added Comparing Parameter Templates.
	Added Replicating a Parameter Template.
	Added Resetting a Parameter Template.
	Added Applying a Parameter Template.
	 Added Viewing Application Records of a Parameter Template.
	 Added Modifying a Parameter Template Description.
	Added Deleting a Parameter Template.
	 Added monitoring metrics in GeminiDB Cassandra Metrics.
2020-03-13	This issue is the third official release.
	Supported slow query logs.
2020-01-20	This issue is the second official release.
	Support modification of GeminiDB Cassandra instance specifications.
2019-11-18	This issue is the first official release.