

GeminiDB

API Reference

Issue 01
Date 2025-01-30



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

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

Trademarks and Permissions



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

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

Notice

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

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

Contents

1 Before You Start	1
1.1 Overview	1
1.2 API Calling	1
1.3 Endpoints	1
1.4 Constraints	1
1.5 Concepts	1
2 API Overview	3
3 Calling APIs	5
3.1 Making an API Request	5
3.2 Authentication	8
3.3 Returned Values	9
4 Quick Start	11
5 APIs v3 (Recommended)	13
5.1 API Versions	13
5.1.1 Querying API Versions	13
5.1.2 Querying Version Information of an API	15
5.2 Versions and Specifications	17
5.2.1 Querying Version Information	17
5.2.2 Querying Instance Specifications	19
5.2.3 Querying Dedicated Resources	25
5.3 Instances	28
5.3.1 Creating an Instance	28
5.3.2 Deleting an Instance	52
5.3.3 Querying Instances and Details	53
5.3.4 Scaling Up Storage Space of an Instance	69
5.3.5 Adding Nodes for an Instance	74
5.3.6 Deleting Nodes from a Specified Instance	78
5.3.7 Obtaining Sessions of a Node	81
5.3.8 Querying Session Statistics of an Instance Node	84
5.3.9 Closing Sessions of an Instance Node	86
5.3.10 Querying Instance Specifications That Can Be Changed	88
5.3.11 Changing Specifications of an Instance	93

5.3.12 Resetting the Administrator Password of an Instance.....	97
5.3.13 Editing the Name of an Instance.....	99
5.3.14 Changing the Security Group of an Instance.....	101
5.3.15 Upgrading Minor Version.....	102
5.3.16 Patching Databases in Batches.....	104
5.3.17 Creating Cold Storage.....	107
5.3.18 Scaling Up Cold Storage.....	109
5.3.19 Binding/Unbinding an EIP.....	111
5.3.20 Enabling or Disabling SSL.....	113
5.3.21 Restarting an Instance.....	115
5.3.22 Configuring an Autoscaling Policy for Storage Space.....	118
5.3.23 Changing a Database Port.....	122
5.3.24 Checking Password Strength.....	126
5.3.25 Configuring Access to a Replica Set Across CIDR Blocks.....	128
5.3.26 Deleting the Node that Fails to Be Added.....	130
5.3.27 Querying IP Addresses Required for Creating an Instance or Adding Nodes.....	132
5.3.28 Querying the Autoscaling Policy of Storage Space.....	134
5.3.29 Scaling Storage Space of an Instance.....	136
5.3.30 Querying High-Risk Commands.....	141
5.3.31 Modifying High-Risk Commands.....	143
5.3.32 Querying Hot Keys of a Redis Instance.....	145
5.3.33 Disabling Commands for a Redis Instance.....	147
5.3.34 Querying Disabled Commands for a Redis Instance.....	150
5.3.35 Deleting Disabled Commands for a Redis Instance.....	153
5.3.36 Setting the Maintenance Period of an Instance.....	156
5.3.37 Performing a Primary/Standby Switchover.....	158
5.3.38 Starting or Stopping a Node.....	160
5.3.39 Querying Big Keys of a GeminiDB Redis Instance.....	162
5.3.40 Querying the Password-Free Configuration of a GeminiDB Redis Instance.....	165
5.3.41 Modifying the Password-Free Configuration of a GeminiDB Redis Instance.....	167
5.3.42 Querying the Memory Mapping List and Details.....	169
5.3.43 Creating a Memory Acceleration Rule.....	172
5.3.44 Deleting a Memory Mapping.....	175
5.3.45 Creating a Memory Mapping.....	176
5.3.46 Modifying a Memory Acceleration Rule.....	178
5.3.47 Querying Memory Mapping Rules and Details.....	179
5.3.48 Deleting a Memory Acceleration Rule.....	183
5.3.49 Enabling or Disabling Instance Data Export.....	184
5.3.50 Enabling or Disabling Second-Level Monitoring.....	186
5.3.51 Querying Configurations of Second-Level Monitoring.....	187
5.4 Connection Management.....	189
5.4.1 Obtaining Instance Sessions.....	189

5.4.2 Closing Sessions of All Nodes on an Instance.....	192
5.5 Backups and Restorations.....	193
5.5.1 Querying Backups.....	193
5.5.2 Querying Backups (Recommended).....	201
5.5.3 Querying an Automated Backup Policy.....	209
5.5.4 Configuring an Automated Backup Policy.....	212
5.5.5 Querying Instances that Can Be Restored.....	215
5.5.6 Querying the Time Window When a Backup Can Be Restored.....	217
5.5.7 Creating a Manual Backup.....	220
5.5.8 Deleting a Manual Backup.....	223
5.5.9 Restoring Data to an Existing Instance.....	226
5.5.10 Querying the Recycling Policy.....	228
5.5.11 Modifying the Recycling Policy.....	230
5.5.12 Querying Instances in the Recycle Bin.....	231
5.5.13 Obtaining GeminiDB Cassandra Instance Database Information That Is Restored Using Tables....	235
5.5.14 Obtaining GeminiDB Cassandra Instance Table Information That Is Restored Using Tables.....	237
5.5.15 Restoring the Current Redis Instance to a Point in Time.....	239
5.5.16 Setting the Policy for Restoring Redis Data to a Specified Time Point.....	241
5.5.17 Querying the Policy for Restoring Redis Data to a Specified Time Point.....	243
5.5.18 Querying the Restoration Time Range of a Redis Instance.....	245
5.5.19 Querying the Storage Space Used for Restoring a Redis Instance to a Specified Time Point.....	248
5.5.20 Stopping a Backup.....	249
5.5.21 Deleting Manual Backups in Batches.....	251
5.6 Parameter Templates.....	252
5.6.1 Obtaining Parameter Templates.....	252
5.6.2 Creating a Parameter Template.....	256
5.6.3 Modifying Parameters in a Parameter Template.....	260
5.6.4 Applying a Parameter Template.....	262
5.6.5 Resetting a Custom Parameter Template.....	264
5.6.6 Modifying Parameters of a Specified Instance.....	265
5.6.7 Querying Instance Parameter Settings.....	267
5.6.8 Obtaining Parameters of a Specified Parameter Template.....	272
5.6.9 Deleting a Parameter Template.....	275
5.6.10 Querying Instances that a Parameter Template Can Be Applied To.....	276
5.6.11 Viewing Parameter Change History of an Instance.....	278
5.6.12 Viewing Application Records of a Parameter Template.....	283
5.6.13 Comparing Parameter Templates.....	286
5.6.14 Replicating a Parameter Template.....	288
5.6.15 Querying API that Support Parameter Templates.....	290
5.7 Managing Databases and Accounts.....	292
5.7.1 Creating a Database Account.....	292
5.7.2 Changing Permissions for a Database Account.....	294

5.7.3 Resetting the Password of a Database Account.....	297
5.7.4 Deleting a Database Account.....	299
5.7.5 Obtaining the Database Account List.....	300
5.7.6 Obtaining All Databases in an Instance.....	303
5.7.7 Performing Operations on a GeminiDB Instance Database.....	305
5.8 Tags.....	307
5.8.1 Querying an Instance by Tag.....	307
5.8.2 Adding or Deleting Resource Tags in Batches.....	312
5.8.3 Querying Tags of an Instance.....	315
5.8.4 Querying Tags of a Specified Project.....	318
5.9 Logs.....	320
5.9.1 Querying Database Slow Logs.....	320
5.9.2 Querying Slow Query Logs of a GeminiDB Redis Instance.....	325
5.9.3 Querying Slow Query Logs of a GeminiDB Influx Instance.....	332
5.9.4 Querying Slow Query Logs of a GeminiDB Cassandra Instance.....	336
5.9.5 Querying Slow Query Logs of a GeminiDB Mongo Instance.....	341
5.9.6 Querying Database Error Logs.....	345
5.9.7 Querying Error Logs of a GeminiDB Mongo Instance.....	349
5.9.8 Setting the Desensitization Status of Slow Query Logs.....	352
5.9.9 Querying the Desensitization Status of Slow Query Logs.....	354
5.9.10 Associating Instances with an LTS Log Stream.....	355
5.9.11 Disassociating Instances from an LTS Log Stream.....	358
5.9.12 Querying LTS Log Configurations.....	359
5.10 Quotas.....	363
5.10.1 Querying Quota.....	364
5.11 Disaster Recovery.....	366
5.11.1 Querying Regions Where a Dual-Active Relationship Can Be Created Between Two Instances.....	366
5.11.2 Setting the Percentage of Faulty Nodes to Be Taken Over.....	368
5.11.3 Querying the Percentage of Faulty Nodes to Be Taken Over.....	370
5.12 Task Management.....	373
5.12.1 Querying Tasks and Details.....	373
5.12.2 Querying the Maintenance Period of an Instance.....	378
5.12.3 Querying Scheduled Tasks.....	380
5.12.4 Canceling a Scheduled Task.....	384
5.13 Enterprise Projects.....	386
5.13.1 Querying Enterprise Project Quotas.....	386
5.13.2 Modifying Enterprise Project Quotas.....	389
5.14 Instance Load Balancing Management.....	391
5.14.1 Configure IP Address Blacklist/Whitelist for Instance Load Balancing.....	391
5.14.2 Querying IP Address Blacklist/Whitelist for Instance Load Balancing.....	393
6 API v3 (Unavailable Soon).....	396
6.1 Instance Specifications.....	396

6.2 Backup and Restoration.....	400
6.2.1 Querying an Automated Backup Policy.....	400
6.3 Parameter Templates.....	402
6.3.1 Obtaining Parameter Templates.....	402
6.3.2 Modifying Parameters of a Specified Instance.....	405
6.3.3 Applying a Parameter Template.....	407
6.4 Tags.....	409
6.4.1 Querying an Instance by Tag.....	409
7 Permission Policies and Supported Actions.....	416
7.1 Introduction.....	416
7.2 GeminiDB Actions.....	417
8 Appendixes.....	429
8.1 Abnormal Request Results.....	429
8.2 Status Codes.....	429
8.3 Error Codes.....	431
8.4 Obtaining a Project ID.....	465
8.5 Metrics.....	466
8.6 Events Supported by Event Monitoring.....	564

1 Before You Start

1.1 Overview

Welcome to GeminiDB API Reference. This document describes how to use application programming interfaces (APIs) to perform operations on GeminiDB, such as creating, deleting, and querying DB instances. For details about all supported operations, see [API Overview](#).

If you want to access GeminiDB using an API, ensure that you are familiar with GeminiDB concepts. For details, see [Service Overview](#).

1.2 API Calling

GeminiDB supports Representational State Transfer (REST) APIs, allowing you to call APIs using HTTPS. For details about API calling, see [Calling APIs](#).

1.3 Endpoints

An endpoint is the **request address** for calling an API. Endpoints vary depending on services and regions. For the endpoints of all services, see [Regions and Endpoints](#).

1.4 Constraints

For more constraints, see the description of each API.

1.5 Concepts


- Account

A domain is created upon successful registration. The domain has full access permissions for all of its cloud services and resources. It can be used to reset user passwords and grant user permissions. The domain is a payment entity

and should not be used directly to perform routine management. For security purposes, create users and grant them permissions for routine management.

- **IAM User**
An IAM user is created using an account to use cloud services. Each IAM user has its own identity credentials (password and access keys).
API authentication requires information such as the account name, username, and password.
- **Region**
A region is a geographic area in which cloud resources are deployed. Availability zones (AZs) in the same region can communicate with each other over an intranet, while AZs in different regions are isolated from each other. Deploying cloud resources in different regions can better suit certain user requirements or comply with local laws or regulations.
- **AZ**
An AZ contains one or more physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proof, and electricity facilities. Within an AZ, computing, network, storage, and other resources are logically divided into multiple clusters. AZs within a region are connected using high-speed optical fibers to support cross-AZ high-availability systems.
- **Project**
A project corresponds to a region. Projects group and isolate resources (including compute, storage, and network resources) across physical regions. Users can be granted permissions in a default project to access all resources in the region associated with the project. If you need more refined access control, create subprojects under a default project and purchase resources in subprojects. Then you can assign users the permissions required to access only the resources in specific subprojects.

Figure 1-1 Project isolating model



- **Enterprise Project**
Enterprise projects group and manage resources across regions. Resources in enterprise projects are logically isolated. An enterprise project can contain resources of multiple regions, and resources can be added to or removed from the enterprise project.

2 API Overview

GeminiDB provides extended APIs. These APIs enable you to use some functions of GeminiDB.

Table 2-1 API description

Type	Description
API Version Queries	<ul style="list-style-type: none"> Query API versions. Query version information of a specific API.
DB Version Queries	Query version information of a specified type of instances.
Instance Specifications Queries	Query all instance specifications under a specified condition.
Instance Management	Create, delete, and query instances, scale up or down the instance storage space, add or delete cluster instance nodes, change instance specifications, and change the instance administrator password, instance name, and instance security group.
Backup and Restoration	Query and set automated backup policies.
Parameter Management	Obtain parameter templates, create a parameter template, modify parameters in a parameter template, apply a parameter template, modify or obtain parameters of a specified instance, obtain parameters of a specified parameter template, and delete a parameter template.
Tag Management	Query resources by tag, batch add and delete tags, and query resource tags.
Log Management	Query database slow query logs.
Quota Management	Query quotas.

Type	Description
Disaster Recovery (DR) Management	Checking whether a DR relationship can be created for one instance with a specified one or deleted, creating a DR relationship between one instance and a specified instance, and deleting a DR relationship between two them.

3 Calling APIs

3.1 Making an API Request

This section describes the structure of a REST API, and uses the IAM API for [obtaining a user token](#) as an example to describe how to call an API. The obtained token is used to authenticate other APIs.

Request URI

A request URI consists of the following:

{URI-scheme}://{Endpoint}/{resource-path}?{query-string}

Although a request URI is included in the request header, most programming languages or frameworks require the request URI to be separately transmitted, rather than being conveyed in a request message separately.

Table 3-1 URI parameter description

Parameter	Description
URI-scheme	Protocol used to transmit requests. All APIs use HTTPS.
Endpoint	Domain name or IP address of a server where a REST service endpoint is hosted. The endpoint varies depending on services and regions. It can be obtained from Regions and Endpoints . For example, the endpoint of IAM in the CN-Hong Kong region is iam.ap-southeast-1.myhuaweicloud.com .
resource-path	Access path of an API for performing a specified operation. Obtain the path from the URI of the API. For example, the resource-path of the API for obtaining a user token is /v3/auth/tokens .

Parameter	Description
query-string	Query parameter, which is optional. Ensure that a question mark (?) is included before each query parameter that is in the format of "Parameter name=Parameter value". For example, ?limit=10 indicates that up to 10 data records will be displayed.

Request Methods

The HTTP protocol defines the following request methods that can be used to send a request to the server:

Table 3-2 HTTP methods

Method	Description
GET	Requests a server to return specified resources.
PUT	Requests a server to update specified resources.
POST	Requests a server to add a resource or perform a special operation.
DELETE	Requests a server to delete a specified resource (for example, an object).

For example, in the URI for [obtaining a user token](#), the request method is POST. The request is as follows:

```
POST https://iam.ap-southeast-1.myhuaweicloud.com/v3/auth/tokens
```

Request Header

You can also add additional header fields to a request, such as the fields required by a specified URI or HTTP method. For example, add **Content-Type** that defines a request body type to request for authentication information.

[Table 3-3](#) lists common request header fields.

Table 3-3 Common request headers

Parameter	Description	Mandatory	Example Value
Content-Type	MIME type of the request body. Use the default value application/json . For APIs used to upload objects or images, the value varies depending on the flow type.	Yes	application/json
Content-Length	Length of the request body. The unit is byte.	This field is optional for POST requests, but must be left blank for GET requests.	3495
X-Project-Id	Project ID. To obtain the project ID, see Obtaining a Project ID .	No	e9993fc787d94b6c886cb aa340f9c0f4
X-Auth-Token	User token. After a request is processed, the value of X-Subject-Token in the header is the token value.	Yes	The following is part of an example token: MIIPAgYJKoZlhvcNAQc- Co...ggg1BBIINPXsidG9rZ

The API used to [obtain a user token](#) does not require authentication. Therefore, this API only requires adding the **Content-Type** field. The following is an example request:

```
POST https://iam.ap-southeast-1.myhuaweicloud.com/v3/auth/tokens
Content-Type: application/json
```

(Optional) Request Body

This part is optional. The request body is often sent in a structured format (for example, JSON or XML) as specified in the **Content-Type** header field. If the request body contains full-width characters, these characters must be coded in UTF-8.

Request bodies vary depending on APIs. Some APIs do not require a request body, such as the APIs requested using the GET and DELETE methods.

For the API of [obtaining a user token](#), request parameters and parameter description can be obtained from the API request. The following is an example

request with a body included. Replace *username*, *domainname*, ******* (login password), and *xxxxxxxxxxxxxxxxxxxx* (project namesuch as cn-north-1) with required values. You can obtain the values from [Regions and Endpoints](#).

 **NOTE**

The **scope** parameter specifies where a token takes effect. You can set **scope** to an account or a project under an account. In the following example, the token takes effect only for the resources in a specified project. For details, see [Obtaining a User Token](#).

```
POST https://iam.ap-southeast-1.myhuaweicloud.com/v3/auth/tokens
Content-Type: application/json

{
  "auth": {
    "identity": {
      "methods": [
        "password"
      ],
      "password": {
        "user": {
          "name": "username",
          "password": "*****",
          "domain": {
            "name": "domainname"
          }
        }
      }
    },
    "scope": {
      "project": {
        "name": "xxxxxxxxxxxxxxxxxxxx"
      }
    }
  }
}
```

If all data required for the API request is available, you can send a request to call an API through [curl](#), [Postman](#), or coding. For the API of obtaining a user token, **x-subject-token** in the response header is the required user token. Then, this token can be used to authenticate the calling of other APIs.

3.2 Authentication

GeminiDB supports token-based authentication.

 **NOTE**

The validity period of a token is 24 hours. If a token is required, the system caches the token to avoid frequent calling.

A token specifies temporary permissions in a computer system. During API authentication using a token, the token is added to a request to get permissions for calling the API.

If you want to use a token for authentication, you need to obtain the user's token and add **X-Auth-Token** to the request header of the service API to make an API call.

When you [call an API to obtain a user token](#), set **auth.scope** in the request body to **project**.

```
{
  "auth": {
    "identity": {
      "methods": [
        "password"
      ],
      "password": {
        "user": {
          "name": "username",
          "password": "password",
          "domain": {
            "name": "domainname"
          }
        }
      }
    },
    "scope": {
      "project": {
        "name": "xxxxxxx"
      }
    }
  }
}
```

After a token is obtained, add field **X-Auth-Token** to the request header to specify the token when other APIs are called. For example, if the token is **ABCDEFJ....**, add **X-Auth-Token: ABCDEFJ....** to a request header as follows:

```
POST https://iam.ap-southeast-1.myhuaweicloud.com/v3/auth/projects
Content-Type: application/json
X-Auth-Token: ABCDEFJ....
```

3.3 Returned Values

Status Code

After you send a request, you will receive a response, including a status code, response header, and response body.

A status code is a group of digits ranging from 1xx to 5xx. It indicates the status of a response. For more information, see [Status Codes](#).

If status code 201 is returned for calling the API used to [obtain a user token](#), the request is successful.

Response Header

Similar to a request, a response also has a header, for example, **Content-Type**.

[Figure 3-1](#) shows the response header for the API used to [obtain a user token](#), in which **x-subject-token** is the required user token. Then, this token can be used to authenticate the calling of other APIs.

Figure 3-1 Response header for the API used to obtain a user token

```

connection → keep-alive

content-type → application/json

date → Tue, 12 Feb 2019 06:52:13 GMT

server → Web Server

strict-transport-security → max-age=31536000; includeSubdomains;

transfer-encoding → chunked

via → proxy A

x-content-type-options → nosniff

x-download-options → noopen

x-frame-options → SAMEORIGIN

x-iam-trace-id → 218d45ab-d674-4995-af3a-2d0255ba41b5

x-subject-token
→ MIIYXQYJKoZIhvcNAQcCoIIVTjCCGEOCAQExDTALBgIghkgBZQMEAgEwgharBgkqhkiG9w0BBwGgghacBIIWmHsidG9rZW4iOnsiZXhwaXJlc19hdCI6IjwMTktMDItMTNUMC
fj3KJs6YgKnpVNRbW2eZ5eb78SZ0kqjACgkqO1wi4JIGzrpd1.8LGXK5bdfq4lqHCYb8P4NaYONYeJcAgz/VeFYtLWT1GSO0zxKZmlQHq82HBqHdgIZO9fuEbL5dMhdavj+33wEI
xHRCE9I87o+k9-
j+CMZSEB7bUGd5Uj6eRASXl1jipPEGA270g1FruooL6jggIFkNPQuFSOU8+uSsttVwRtnfsC+qTp22Rkd5MCqFGQ8LcuUxC3a+9CMBnOintWW7oeRUUVhVpxk8pxiX1wTEboX-
RzT6MUbvpvGw-oPNFYxJECKnoH3HRozv0vN--n5d6Nbxg==

x-xss-protection → 1; mode=block;

```

(Optional) Response Body

This part is optional. A response body is generally returned in a structured format (for example, JSON or XML), corresponding to **Content-Type** in the response header, and is used to transfer content other than the response header.

If the following information is returned for calling the API used to **obtain a user token**, the request is successful. The following describes part of the request body.

```

{
  "token": {
    "expires_at": "2019-02-13T06:52:13.855000Z",
    "methods": [
      "password"
    ],
    "catalog": [
      {
        "endpoints": [
          {
            "region_id": "aaa",
            .....

```

If an error occurs during API calling, an error code and error message will be displayed. The following is an error response body:

```

{
  "error_msg": "Parameter error",
  "error_code": "DBS.200001"
}

```

In the response, **error_code** indicates an error code, and **error_msg** describes the error.

4 Quick Start

This section describes how to create a GeminiDB Cassandra instance by calling APIs.

NOTE

The token obtained from IAM is valid for only 24 hours. If you want to use a token for authentication, you can cache it to avoid frequently obtaining the token.

Involved APIs

If you use a token for authentication, you must obtain the user's token and add **X-Auth-Token** to the request message header of the service API when making an API call.

- API for obtaining tokens from IAM
- API for creating a GeminiDB Cassandra instance

Procedure

1. Use a token for authentication by referring to [Authentication](#).
2. Send **POST** `https://{Endpoint}/v3/{project_id}/instances`.
3. Add **X-Auth-Token** to the request header.
4. Transfer the following parameters in the request body:

NOTE

Values of **region** and **availability_zone** are only for reference.

For details about the API for creating instances, see [Creating an Instance](#).

```
{
  "name": "test-cassandra-01", //Instance name
  "datastore": {
    "type": "cassandra", //Database type
    "version": "3.11", //DB engine version
    "storage_engine": "rocksDB" //Storage engine
  },
  "region": "aaa", //Region
  "availability_zone": "bbb", //AZ
  "vpc_id": "674e9b42-cd8d-4d25-a2e6-5abcc565b961", //VPC ID
  "subnet_id": "f1df08c5-71d1-406a-aff0-de435a51007b", //Subnet ID
}
```

```
"security_group_id": "7aa51dbf-5b63-40db-9724-dad3c4828b58", //Security group ID
"password": "xxxx", //Administrator password
"mode": "Cluster", //Instance type
"flavor": [
  {
    "num": 3, //Nodes
    "size": 500, //Storage space
    "storage": "ULTRAHIGH", //Disk type
    "spec_code": "geminidb.cassandra.4xlarge.4" //Resource specification code
  }
],
"backup_strategy": {
  "start_time": "08:00-09:00", //Backup time window
  "keep_days": "8" //Retention period of backup files
},
"enterprise_project_id": "0" //Enterprise project ID
}
```

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

```
{
  "id": "39b6a1a278844ac48119d86512e0000bin06",
  "name": "test-cassandra-01",
  "datastore": {
    "type": "cassandra",
    "version": "3.11",
    "storage_engine": "rocksDB"
  },
  "created": "2019-10-28 14:10:54",
  "status": "creating",
  "region": "aaa",
  "availability_zone": "bbb",
  "vpc_id": "674e9b42-cd8d-4d25-a2e6-5abcc565b961",
  "subnet_id": "f1df08c5-71d1-406a-aff0-de435a51007b",
  "security_group_id": "7aa51dbf-5b63-40db-9724-dad3c4828b58",
  "mode": "Cluster",
  "flavor": [
    {
      "num": 3,
      "size": 500,
      "storage": "ULTRAHIGH",
      "spec_code": "geminidb.cassandra.4xlarge.4"
    }
  ],
  "backup_strategy": {
    "start_time": "08:00-09:00",
    "keep_days": "8"
  },
  "job_id": "c010abd0-48cf-4fa8-8cbc-090f093eaa2f",
  "enterprise_project_id": "0"
}
```

If the request fails, an error code and error information are returned. For details, see [Error Codes](#).

5 APIs v3 (Recommended)

5.1 API Versions

5.1.1 Querying API Versions

Function

This API is used to query the supported API versions.

URI

GET https://{Endpoint}/

Request Parameters

None

Response Parameters

Status code: 200

Table 5-1 Response body parameters

Parameter	Type	Description
versions	Array of ApiVersionResponse objects	API version information

Table 5-2 ApiVersionResponse

Parameter	Type	Description
id	String	API version number
links	Array of Links objects	API link information NOTE If the version is v3, the value is [].
status	String	Version status
version	String	Subversion information of the API version
min_version	String	Minimum API version number
updated	String	Version update time The format is yyyy-mm-dd Thh:mm:ssZ. T is the separator between the calendar and the hourly notation of time. Z indicates the UTC.

Table 5-3 Links

Parameter	Type	Description
href	String	API URL. The value is "".
rel	String	The value is self , indicating that URL is a local link.

Example Requests

URI example

GET <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/>

Example Response

Status code: 200

Success

```
{
  "versions": [{
    "id": "v3",
    "links": [],
    "status": "CURRENT",
    "version": "",
    "min_version": "",
    "updated": "2019-10-30T17:34:02Z"
  }]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.1.2 Querying Version Information of an API

Function

This API is used to query version information of a specified API.

URI

GET https://{Endpoint}/{versionId}

Table 5-4 Path parameters

Parameter	Mandatory	Type	Description
versionId	Yes	String	API version

Request Parameters

None

Response Parameters

Status code: 200

Table 5-5 Response body parameters

Parameter	Type	Description
version	ApiVersionResponse object	API version information

Table 5-6 ApiVersionResponse

Parameter	Type	Description
id	String	API version number
links	Array of Links objects	API link information NOTE If the version is v3, the value is [].
status	String	Version status

Parameter	Type	Description
version	String	Subversion information of the API version
min_version	String	Minimum API version number
updated	String	Version update time The format is yyyy-mm-dd Thh:mm:ssZ. T is the separator between the calendar and the hourly notation of time. Z indicates the UTC.

Table 5-7 Links

Parameter	Type	Description
href	String	API URL. The value is "".
rel	String	The value is self , indicating that URL is a local link.

Example Requests

URI example

GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3

Example Response

Status code: 200

Success

```
{
  "version": {
    "id": "v3",
    "links": [ ],
    "status": "CURRENT",
    "version": "",
    "min_version": "",
    "updated": "2019-10-30T17:34:02Z"
  }
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.2 Versions and Specifications

5.2.1 Querying Version Information

Function

This API is used to query version information of a specified type of instances.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

URI

GET `https://{Endpoint}/v3/{project_id}/datastores/{datastore_name}/versions`

Table 5-8 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Parameter	Mandatory	Type	Description
datastore_name	Yes	String	<p>Explanation: Database type</p> <p>Constraints: N/A</p> <p>Values:</p> <ul style="list-style-type: none"> • cassandra: GeminiDB Cassandra instance • mongodb: GeminiDB Mongo instance • influxdb: GeminiDB Influx instance • redis: GeminiDB Redis instance <p>Default value: N/A</p>

Request Parameters

Table 5-9 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Response Parameters

Status code: 200

Table 5-10 Response body parameters

Parameter	Type	Description
versions	Array of strings	<p>Explanation: Database version</p> <p>Values:</p> <ul style="list-style-type: none"> • GeminiDB Cassandra instance 3.11 • GeminiDB Mongo replica set 4.0 • GeminiDB Influx instance 1.8 • GeminiDB Redis instance 5.0

Example Requests

URI example

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/
datastores/cassandra/versions
```

Example Responses

Status code: 200

Success

```
{
  "versions": [ "3.11" ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.2.2 Querying Instance Specifications

Function

This API is used to query all instance specifications under a specified condition.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

URI

GET https://{Endpoint}/v3.1/{project_id}/flavors

Table 5-11 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Table 5-12 Query parameters

Parameter	Mandatory	Type	Description
engine_name	No	String	<p>Explanation: Database type</p> <p>Constraints: N/A</p> <p>Values:</p> <ul style="list-style-type: none"> • cassandra, indicating that the instances are of the GeminiDB Cassandra type. • mongodb, indicating that the instances are of the GeminiDB Mongo type. • influxdb, indicating that the instances are of the GeminiDB Influx type. • redis, indicating that the instances are of the GeminiDB Redis type. <p>Default value: If this parameter is not transferred, the default value is cassandra.</p>

Parameter	Mandatory	Type	Description
offset	No	Integer	<p>Explanation: Index offset.</p> <p>Constraints: N/A</p> <p>Values: The value must be a positive integer.</p> <p>Default value: If offset is set to N, the resource query starts from the N+1 piece of data. If action is set to filter, offset is 0 by default, the query starts from the first record.</p>
limit	No	Integer	<p>Explanation: Maximum of specifications that can be queried</p> <p>Constraints: N/A</p> <p>Values: 1 to 100</p> <p>Default value: If this parameter is not transferred, the first 100 pieces of specification information is queried by default.</p>

Request Parameters

Table 5-13 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Response Parameters

Status code: 200

Table 5-14 Response body parameters

Parameter	Type	Description
total_count	Integer	<p>Explanation: Total number of records</p> <p>Values: N/A</p>
flavors	Array of Flavors objects	<p>Explanation: Instance specifications</p>

Table 5-15 Flavors

Parameter	Type	Description
engine_name	String	<p>Explanation: API name</p> <p>Values:</p> <ul style="list-style-type: none"> • cassandra: specifications supported by a GeminiDB Cassandra instance • mongodb: specifications supported by a GeminiDB Mongo instance • influxdb: specifications supported by a GeminiDB Influx instance • redis: specifications supported by a GeminiDB Redis instance
engine_version	String	<p>Explanation: API version.</p> <p>Values: N/A</p> <p>Default value: N/A</p>
vcpus	String	<p>Explanation: Number of CPU cores</p> <p>Values: N/A</p> <p>Default value: N/A</p>
ram	String	<p>Explanation: Memory size, in MB</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Parameter	Type	Description
spec_code	String	<p>Explanation: Resource specification code.</p> <p>Example: geminidb.cassandra.8xlarge.4</p> <p>NOTE</p> <ul style="list-style-type: none"> • geminidb.cassandra indicates the instance is a GeminiDB Cassandra instance. • 8xlarge.4 indicates node specifications. <p>Values: N/A</p>
availability_zone	Array of strings	<p>Explanation: ID of the AZ that supports the specifications</p> <p>NOTE This parameter has been discarded. Do not use it.</p> <p>Values: N/A</p>
az_status	Object	<p>Explanation: Status of specifications in an AZ</p> <p>Values:</p> <ul style="list-style-type: none"> • normal: The product is on sale. • unsupported: The specifications are not supported. • sellout: The product is sold out.

Example Requests

URI example

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3.1/375d8d8fad1f43039e23d3b6c0f60a19/flavors?engine_name=cassandra&offset=0&limit=10
```

Example Responses

Status code: 200

Success

```
{
  "total_count": 4,
  "flavors": [ {
    "engine_name": "cassandra",
    "engine_version": "3.11",
    "vcpus": "4",
    "ram": "16",
    "spec_code": "geminidb.cassandra.xlarge.4",
```

```
"availability_zone": [ "az1", "az2" ],
"az_status": {
  "az1": "normal",
  "az2": "unsupported"
}
}, {
  "engine_name": "cassandra",
  "engine_version": "3.11",
  "vcpus": "8",
  "ram": "32",
  "spec_code": "geminidb.cassandra.2xlarge.4",
  "availability_zone": [ "az1", "az2" ],
  "az_status": {
    "az1": "unsupported",
    "az2": "normal"
  }
}, {
  "engine_name": "cassandra",
  "engine_version": "3.11",
  "vcpus": "16",
  "ram": "64",
  "spec_code": "geminidb.cassandra.4xlarge.4",
  "availability_zone": [ "az1", "az2" ],
  "az_status": {
    "az1": "normal",
    "az2": "sellout"
  }
}, {
  "engine_name": "cassandra",
  "engine_version": "3.11",
  "vcpus": "32",
  "ram": "128",
  "spec_code": "geminidb.cassandra.8xlarge.4",
  "availability_zone": [ "az1", "az2" ],
  "az_status": {
    "az1": "normal",
    "az2": "normal"
  }
}
}]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.2.3 Querying Dedicated Resources

Function

This API is used to query dedicated resources created by a user.

Constraints

This API supports GeminiDB Cassandra instances.

URI

GET https://{Endpoint}/v3/{project_id}/dedicated-resources

Table 5-16 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Table 5-17 Query parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	<p>Index position. The query starts from the dedicated resource created after the dedicated resource indexed by this parameter in a specified project. If offset is set to N, the resource query starts from the N+1 piece of data.</p> <ul style="list-style-type: none"> The value must be no less than 0. If this parameter is not transferred, offset is set to 0 by default, indicating that the query starts from the latest created dedicated resource.
limit	No	Integer	<p>Maximum of dedicated resources to be queried.</p> <ul style="list-style-type: none"> The value ranges from 1 to 100. If this parameter is not transferred, the first 100 instances are queried by default.

Request Parameters

None

Response Parameters

Status code: 200

Table 5-18 Response body parameters

Parameter	Type	Description
total_count	Integer	Total number of records
resources	Array of DedicatedResource objects	Dedicated resource information

Table 5-19 DedicatedResource

Parameter	Type	Description
id	String	Dedicated resource ID
resource_name	String	Dedicated resource name
engine_name	String	API name.
availability_zone	String	AZ information
architecture	String	Type of dedicated compute hosts. The value can be X86 or ARM .
capacity	DedicatedResourceCapacity object	Capacity of the dedicated resource
status	String	Status of the dedicated resource. The value can be: <ul style="list-style-type: none"> • NORMAL: indicating that the dedicated resource is available. • BUILDING, indicating that the dedicated resource is being built.

Table 5-20 DedicatedResourceCapacity

Parameter	Type	Description
vcpus	Integer	Number of vCPUs
ram	Integer	Memory size in GB
volume	Integer	Storage size in GB

Example Requests

URI example

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/dedicated-resources?offset=0&limit=10
```

Example Responses

Status code: 200

Success

```
{
  "total_count": 2,
  "resources": [ {
    "id": "300acc71-eca9-4c6b-9a4e-0d25e20fe54c",
    "resource_name": "nosql-dedicated-resource-1",
    "engine_name": "cassandra",
    "availability_zone": "az1,az2,az3",
    "architecture": "X86",
    "capacity": {
      "vcpus": 192,
      "ram": 1536,
      "volume": 300000
    },
    "status": "NORMAL"
  }, {
    "id": "4c423cd8-551f-4b74-934f-8534e4a9bf2b",
    "resource_name": "nosql-dedicated-resource-2",
    "engine_name": "cassandra",
    "availability_zone": "az1,az2,az3",
    "architecture": "X86",
    "capacity": {
      "vcpus": 192,
      "ram": 1536,
      "volume": 300000
    },
    "status": "BUILDING"
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3 Instances

5.3.1 Creating an Instance

Function

- This API can be used to create an instance.
- The API can be used to create an instance when you restore data using a specific backup.
- The API can also be used to create an instance when you restore data of a specific instance to a specified point in time.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis
- GeminiDB HBase
- GeminiDB DynamoDB

This API supports both yearly/monthly and pay-per-use instances.

Only GeminiDB Cassandra and GeminiDB Influx allow you to restore data of a specified time point from a specified cluster instance to a new one.

URI

POST https://{Endpoint}/v3/{project_id}/instances

Table 5-21 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-22 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-23 Request body parameters

Parameter	Mandatory	Type	Description
name	Yes	String	Instance name, which can be the same as an existing instance name. The name must start with a letter and can include 4 to 64 characters. It is case-sensitive and can contain only letters, digits, hyphens (-), and underscores (_).
datastore	Yes	Datastore object	Database information.
region	Yes	String	Region ID. The value cannot be empty. For value details, see Regions and Endpoints .
availability_zone	Yes	String	AZ ID. For details about the value, see az_status returned in Querying Instance Specifications . If an instance can be created across three AZs, separate multiple AZ IDs by commas (,).
vpc_id	Yes	String	VPC ID. You can obtain the value with either of the following methods: <ul style="list-style-type: none"> Method 1: Log in to the VPC console and view the VPC ID on the VPC details page. Method 2: Query the VPC ID using the VPC API. For details, see Querying VPCs.

Parameter	Mandatory	Type	Description
subnet_id	Yes	String	<p>Subnet ID. You can obtain the subnet ID with either of the following methods:</p> <ul style="list-style-type: none"> Method 1: Log in to the VPC console and click the target subnet on the Subnets page. You can view the network ID on the displayed page. Method 2: Query the subnet ID using the VPC API. For details, see Querying Subnets.
security_group_id	Yes	String	<p>Security group ID. You can obtain the security group ID with either of the following methods:</p> <ul style="list-style-type: none"> Method 1: Log in to the VPC console. Choose Access Control > Security Groups in the navigation pane on the left. On the displayed page, click the target security group. You can view the security group ID on the displayed page. Method 2: Query the security group using the VPC API. For details, see Querying Security Groups.
password	Yes	String	<p>Database password.</p> <p>The password can include 8 to 32 characters and contain uppercase letters, lowercase letters, digits, and the following special characters: ~!@#%^*-_+=? The password of GeminiDB Redis instances can contain at least two types of the following characters: uppercase letters, lowercase letters, digits, and special characters (~!@#%^*&*()-_+=?).</p> <p>Enter a strong password against security risks such as brute force cracking.</p>

Parameter	Mandatory	Type	Description
mode	Yes	String	<p>Instance type. The value can be:</p> <ul style="list-style-type: none"> • Cluster, indicating that GeminiDB Cassandra supports the cluster type. • ReplicaSet, indicating that GeminiDB Mongo 4.0 supports the replica set type. • Cluster, indicating that GeminiDB Influx supports the cluster type. • InfluxdbSingle, indicating that GeminiDB Influx supports the single-node instance type. • Cluster: A GeminiDB Redis instance supports a proxy cluster. • A GeminiDB Redis instance supports Redis Cluster. The value is RedisCluster. For details about Redis Cluster instances, see Creating an Instance. • Replication, indicating that GeminiDB Redis supports the primary/standby type.
flavor	Yes	Array of Flavor objects	<p>Instance specifications. For details about the specifications, see parameter values under flavors in Querying Instance Specifications.</p>
configuration_id	No	String	Parameter template ID.
backup_strategy	No	BackupStrategy object	Advanced backup policy.

Parameter	Mandatory	Type	Description
enterprise_project_id	No	String	Enterprise project ID. <ul style="list-style-type: none"> Do not transfer this parameter if EPS is not enabled. If EPS is enabled but this parameter is not transferred, the default enterprise project is used. For the enterprise project ID, see the id value in the enterprise_project field data structure table in section "Querying the Enterprise Project List" of the <i>Enterprise Project Management Service API Reference</i>.
ssl_option	No	String	Whether SSL is enabled. The value can be: <ul style="list-style-type: none"> 0, indicating that SSL is disabled by default. 1, indicating that SSL is enabled by default. If this parameter is not transferred, SSL is disabled by default.
charge_info	No	ChargeInfo object	Billing mode, which includes yearly/monthly and pay-per-use. The default billing mode is pay-per-use.
dedicated_resource_id	No	String	Dedicated resource ID. This parameter can be delivered only after the dedicated resource pool is enabled.
restore_info	No	RestoreInfo object	Backup information. You can restore data from a specific backup or instance to a specific point in time during the backup retention period. Only GeminiDB Cassandra and GeminiDB Influx allow you to restore data from a specific cluster instance to a specific point in time.

Parameter	Mandatory	Type	Description
port	No	String	<p>Port number for accessing the instance.</p> <p>You can specify a port number for accessing GeminiDB Redis instances. The port number ranges from 1024 to 65535, excluding 2180, 2887, 3887, 6377, 6378, 6380, 8018, 8079, 8091, 8479, 8484, 8999, 9864, 9866, 9867, 12017, 12333, and 50069.</p> <p>If you do not specify a port number, port 6379 is used by default when you create a GeminiDB Redis instance.</p> <p>If you want to use this instance for dual-active DR, set the port to 8635.</p>
availability_zone_detail	No	object	<p>Multi-AZ details of the active/standby instance</p> <p>Currently, only GeminiDB Redis instances are supported.</p> <p>The system ignores this parameter if single-AZ deployment is selected. For details, see Table 5-29.</p>

Table 5-24 Datastore

Parameter	Mandatory	Type	Description
type	Yes	String	<p>Database type</p> <ul style="list-style-type: none"> • GeminiDB Cassandra, GeminiDB Mongo, GeminiDB Influx, and GeminiDB Redis instances are supported. • If you set this parameter to cassandra, GeminiDB Cassandra instances will be created. • If you set this parameter to mongodb, GeminiDB Mongo instances will be created. • If you set this parameter to influxdb, GeminiDB Influx instances will be created. • If you set this parameter to redis, GeminiDB Redis instances will be created. • If you set this parameter to dynamodb, GeminiDB (DynamoDB API compatible) instances will be created. • If you set this parameter to hbase, GeminiDB HBase instances will be created.
version	Yes	String	<p>Database version</p> <ul style="list-style-type: none"> • 3.11: GeminiDB Cassandra instance 3.11 • 4.0: GeminiDB Mongo instance 4.0 • 1.8: GeminiDB Influx instance 1.8 • 5.0: GeminiDB Redis instance 5.0 • For a GeminiDB DynamoDB instance, the value is null. • For a GeminiDB HBase instance, the value is null.

Parameter	Mandatory	Type	Description
storage_engine	Yes	String	Storage engine. <ul style="list-style-type: none"> • rocksDB: A GeminiDB Cassandra instance supports RocksDB. • rocksDB: A GeminiDB Mongo instance supports RocksDB. • rocksDB: A GeminiDB Influx instance supports RocksDB. • rocksDB: A GeminiDB Redis instance supports RocksDB.

Table 5-25 Flavor

Parameter	Mandatory	Type	Description
num	Yes	String	Number of nodes. <ul style="list-style-type: none"> • Each GeminiDB Cassandra instance can contain 3 to 60 nodes. • Each GeminiDB Mongo replica set 4.0 can contain 3 nodes. • Each GeminiDB Influx cluster instance can contain 3 to 16 nodes. • Each GeminiDB Influx single-node instance can contain 1 node. • Each GeminiDB Redis instance can contain 3 to 12 nodes.

Parameter	Mandatory	Type	Description
size	Yes	String	<p>Storage space. It must be an integer, in GB.</p> <p>For GeminiDB Cassandra, GeminiDB Mongo, and GeminiDB Influx instances, the minimum storage space is 100 GB, and the maximum limit depends on instance specifications. The maximum and minimum storage space of a GeminiDB Redis instance depends on node quantity and specifications of the instance.</p> <ul style="list-style-type: none"> • For details about GeminiDB Cassandra instances, see Instance Specifications. • For details about GeminiDB Mongo instances, see Instance Specifications. • For details about GeminiDB Influx instances, see Instance Specifications. • For details about GeminiDB Redis instances, see Instance Specifications.
storage	Yes	String	<p>Disk type.</p> <p>If you set this parameter to ULTRAHIGH, SSD disks are used.</p>
spec_code	Yes	String	<p>Resource specification code.</p> <p>For the code, see the value of response parameter spec_code in Querying Instance Specifications.</p>

Table 5-26 BackupStrategy

Parameter	Mandatory	Type	Description
start_time	Yes	String	<p>Backup time window. Automated backup will be triggered during the backup time window.</p> <p>The value cannot be empty. It must be the UTC time in the hh:mm-HH:MM format.</p> <ul style="list-style-type: none"> • The HH value must be 1 greater than the hh value. • The values of mm and MM must be the same and must be set to 00, 15, 30, or 45. • If this parameter is not transferred, the default backup time window is from 00:00 to 01:00. • Example value: 23:00-00:00
keep_days	No	String	<p>Backup retention days. The value ranges from 0 to 35.</p> <ul style="list-style-type: none"> • If this parameter is set to 0, the automated backup policy is not set. • If this parameter is not transferred, the automated backup policy is enabled by default. Backup files are stored for 7 days by default.

Table 5-27 ChargeInfo

Parameter	Mandatory	Type	Description
charge_mode	Yes	String	<p>Billing mode.</p> <p>Values:</p> <ul style="list-style-type: none"> • prePaid: indicates that the billing mode is yearly/monthly. • postPaid: indicates that the billing mode is pay-per-use.

Parameter	Mandatory	Type	Description
period_type	No	String	<p>Subscription period type.</p> <p>Values:</p> <ul style="list-style-type: none"> • month: indicates that the subscription unit is month. • year: indicates that the subscription unit is year. <p>NOTE This parameter is valid and mandatory only when charge_mode is set to prePaid.</p>
period_num	No	String	<p>Subscription time period. This parameter is available and mandatory only when charge_mode is set to prePaid.</p> <p>Values:</p> <ul style="list-style-type: none"> • If period_type is set to month, the parameter value ranges from 1 to 9. • If period_type is set to year, the parameter value ranges from 1 to 3.
is_auto_renew	No	String	<p>Whether automatic renewal is enabled for yearly/monthly instances. The renewal period is the same as the original period, and the order will be automatically paid during the renewal. The value can be:</p> <ul style="list-style-type: none"> • true, indicating that the subscription is automatically renewed. • false, indicating that the subscription is not automatically renewed. The default value is false.

Parameter	Mandatory	Type	Description
is_auto_pay	No	String	<p>Payment method. When you create a yearly/monthly instance, you can specify whether the order is automatically paid from your account. This parameter does not affect the payment mode of automatic renewal. The value can be:</p> <ul style="list-style-type: none"> • true, indicating that the order is automatically paid from your account. • false, indicating that the order needs to be manually paid from your account. This payment method is used by default.

Table 5-28 RestoreInfo

Parameter	Mandatory	Type	Description
backup_id	No	String	<p>Full backup file ID.</p> <p>This parameter cannot be left blank when you create an instance to restore data using a specific backup.</p>
source_instance_id	No	String	<p>ID of the specified instance that backup data is restored to.</p> <p>This parameter cannot be left blank when you restore data at a specific point in time from a specific instance to a new instance.</p>
restore_time	No	Long	<p>Time point that backup data is restored to.</p> <p>This parameter cannot be left blank when you restore data at a specific point in time from a specific instance to a new instance. The value is a 13-digit number (in milliseconds, UTC time). You can query the value by referring to Querying the Time Window When a Backup Can Be Restored.</p>

Table 5-29 AvailabilityZoneDetail

Parameter	Mandatory	Type	Description
primary_availability_zone	Yes	String	The primary AZ must be a single AZ and be different from the standby AZ.
secondary_availability_zone	Yes	String	The standby AZ must be a single AZ and be different from the primary AZ.

Response Parameters

Status code: 202

Table 5-30 Response body parameters

Parameter	Type	Description
id	String	Instance ID.
name	String	Instance name. This parameter is the same as the corresponding request parameter.
datastore	Datastore object	Database information. This parameter is the same as the corresponding request parameter.
created	String	Creation time, which is in the <i>yyyy-mm-dd hh:mm:ss</i> format
status	String	Instance status. The value is creating .
region	String	Region ID. This parameter is the same as the corresponding request parameter.
availability_zone	String	AZ ID. This parameter is the same as the corresponding request parameter.
vpc_id	String	VPC ID. This parameter is the same as the corresponding request parameter.
subnet_id	String	Subnet ID. This parameter is the same as the corresponding request parameter.
security_group_id	String	Security group ID. This parameter is the same as the corresponding request parameter.

Parameter	Type	Description
mode	String	Instance type. This parameter is the same as the corresponding request parameter.
flavor	Array of Flavor objects	Instance specifications. This parameter is the same as the corresponding request parameter.
backup_strategy	BackupStrategy object	Advanced backup policy. This parameter is the same as the corresponding request parameter.
enterprise_project_id	String	Enterprise project ID. If you set this parameter to 0 , the resource belongs to the default enterprise project.
ssl_option	String	Whether SSL is enabled. This parameter has the same effect as the corresponding request parameter.
job_id	String	ID of the workflow for creating an instance. This parameter is returned only when a pay-per-use instance is created.
order_id	String	ID of the order for creating an instance. This parameter is returned only when you create a yearly/monthly instance.
charge_info	ChargeInfo object	Billing mode, which includes yearly/monthly and pay-per-use. The default billing mode is pay-per-use.
dedicated_resource_id	String	Dedicated resource ID. This parameter is returned only when the DB instance belongs to a dedicated resource pool.

Table 5-31 Datastore

Parameter	Type	Description
type	String	<p>Database type</p> <ul style="list-style-type: none"> • GeminiDB Cassandra, GeminiDB Mongo, GeminiDB Influx, and GeminiDB Redis instances are supported. • If you set this parameter to cassandra, GeminiDB Cassandra instances will be created. • If you set this parameter to mongodb, GeminiDB Mongo instances will be created. • If you set this parameter to influxdb, GeminiDB Influx instances will be created. • If you set this parameter to redis, GeminiDB Redis instances will be created. • If you set this parameter to dynamodb, GeminiDB (DynamoDB API compatible) instances will be created. • If you set this parameter to hbase, GeminiDB HBase instances will be created.
version	String	<p>Database version</p> <ul style="list-style-type: none"> • 3.11, indicating that GeminiDB Cassandra 3.11 is supported. • 4.0, indicating that GeminiDB Mongo 4.0 is supported. • 1.8: GeminiDB Influx instance 1.8 • 5.0, indicating that GeminiDB Redis 5.0 is supported. • For a GeminiDB DynamoDB instance, the value is null. • For a GeminiDB HBase instance, the value is null.

Parameter	Type	Description
storage_engine	String	Storage engine. <ul style="list-style-type: none"> ● rocksDB: A GeminiDB Cassandra instance supports RocksDB. ● rocksDB: A GeminiDB Mongo instance supports RocksDB. ● rocksDB: A GeminiDB Influx instance supports RocksDB. ● rocksDB: A GeminiDB Redis instance supports RocksDB.

Table 5-32 Flavor

Parameter	Type	Description
num	String	Number of nodes. <ul style="list-style-type: none"> ● Each GeminiDB Cassandra instance can contain 3 to 60 nodes. ● Each GeminiDB Mongo replica set 4.0 can contain 3 nodes. ● Each GeminiDB Influx instance can contain 3 to 16 nodes. ● Each GeminiDB Redis instance can contain 3 to 12 nodes.

Parameter	Type	Description
size	String	<p>Storage space. It must be an integer, in GB.</p> <p>For GeminiDB Cassandra, GeminiDB Mongo, and GeminiDB Influx instances, the minimum storage space is 100 GB, and the maximum limit depends on instance specifications. The maximum and minimum storage space of a GeminiDB Redis instance depends on node quantity and specifications of the instance.</p> <ul style="list-style-type: none"> • For details about GeminiDB Cassandra instances, see Instance Specifications. • For details about GeminiDB Mongo instances, see Instance Specifications. • For details about GeminiDB Influx instances, see Instance Specifications. • For details about GeminiDB Redis instances, see Instance Specifications.
storage	String	<p>Disk type.</p> <p>If you set this parameter to ULTRAHIGH, SSD disks are used.</p>
spec_code	String	<p>Resource specification code.</p> <p>For the code, see the value of response parameter spec_code in Querying Instance Specifications.</p>

Table 5-33 BackupStrategy

Parameter	Type	Description
start_time	String	<p>Backup time window. Automated backup will be triggered during the backup time window.</p> <p>The value cannot be empty. It must be the UTC time in the hh:mm-HH:MM format.</p> <ul style="list-style-type: none"> The HH value must be 1 greater than the hh value. The values of mm and MM must be the same and must be set to 00, 15, 30, or 45. If this parameter is not transferred, the default backup time window is from 00:00 to 01:00. Example value: 23:00-00:00
keep_days	String	<p>Backup retention days.</p> <p>The value ranges from 0 to 35.</p> <ul style="list-style-type: none"> If this parameter is set to 0, the automated backup policy is not set. If this parameter is not transferred, the automated backup policy is enabled by default. Backup files are stored for 7 days by default.

Table 5-34 ChargeInfo

Parameter	Type	Description
charge_mode	String	<p>Billing mode.</p> <p>Values:</p> <ul style="list-style-type: none"> prePaid: indicates that the billing mode is yearly/monthly. postPaid: indicates that the billing mode is pay-per-use.

Parameter	Type	Description
period_type	String	<p>Subscription period type.</p> <p>Values:</p> <ul style="list-style-type: none"> • month: indicates that the subscription unit is month. • year: indicates that the subscription unit is year. <p>NOTE This parameter is valid and mandatory only when charge_mode is set to prePaid.</p>
period_num	String	<p>Subscription time period. This parameter is available and mandatory only when charge_mode is set to prePaid.</p> <p>Values:</p> <ul style="list-style-type: none"> • If period_type is set to month, the parameter value ranges from 1 to 9. • If period_type is set to year, the parameter value ranges from 1 to 3.
is_auto_renew	String	<p>Whether automatic renewal is enabled for yearly/monthly instances. The renewal period is the same as the original period, and the order will be automatically paid during the renewal. The value can be:</p> <ul style="list-style-type: none"> • true, indicating that the subscription is automatically renewed. • false, indicating that the subscription is not automatically renewed. The default value is false.

Parameter	Type	Description
is_auto_pay	String	<p>Payment method. When you create a yearly/monthly instance, you can specify whether the order is automatically paid from your account. This parameter does not affect the payment mode of automatic renewal. The value can be:</p> <ul style="list-style-type: none"> • true, indicating that the order is automatically paid from your account. • false, indicating that the order needs to be manually paid from your account. This payment method is used by default.

Example Request

- URI example
POST `https://{Endpoint}/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances`
- Creating a pay-per-use 3-node GeminiDB Cassandra instance with 16 vCPUs and 64 GB of memory

NOTE

Values of **region** and **availability_zone** in the request body are only examples. Set them based on service requirements.

```
{
  "name": "test-cassandra-01",
  "datastore": {
    "type": "cassandra",
    "version": "3.11",
    "storage_engine": "rocksDB"
  },
  "region": "aaa",
  "availability_zone": "bbb",
  "vpc_id": "674e9b42-cd8d-4d25-a2e6-5abcc565b961",
  "subnet_id": "f1df08c5-71d1-406a-aff0-de435a51007",
  "security_group_id": "7aa51dbf-5b63-40db-9724-dad3c4828b58",
  "password": "*****",
  "mode": "Cluster",
  "flavor": [ {
    "num": 3,
    "storage": "ULTRAHIGH",
    "size": 500,
    "spec_code": "geminidb.cassandra.4xlarge.4"
  } ],
  "backup_strategy": {
    "start_time": "08:15-09:15",
    "keep_days": 8
  },
  "ssl_option": 1
}
```

- Creating a yearly/monthly 3-node GeminiDB Cassandra instance with 16 vCPUs and 64 GB of memory

 NOTE

Values of **region** and **availability_zone** in the request body are only examples. Set them based on service requirements.

```
{
  "name": "test-cassandra-01",
  "datastore": {
    "type": "cassandra",
    "version": "3.11",
    "storage_engine": "rocksDB"
  },
  "region": "aaa",
  "availability_zone": "bbb",
  "vpc_id": "674e9b42-cd8d-4d25-a2e6-5abcc565b961",
  "subnet_id": "f1df08c5-71d1-406a-aff0-de435a51007",
  "security_group_id": "7aa51dbf-5b63-40db-9724-dad3c4828b58",
  "password": "*****",
  "mode": "Cluster",
  "flavor": [ {
    "num": 3,
    "storage": "ULTRAHIGH",
    "size": 500,
    "spec_code": "geminidb.cassandra.4xlarge.4"
  } ],
  "backup_strategy": {
    "start_time": "08:15-09:15",
    "keep_days": 8
  },
  "ssl_option": 1,
  "charge_info": {
    "charge_mode": "prePaid",
    "period_type": "year",
    "period_num": 3,
    "is_auto_renew": true,
    "is_auto_pay": true
  }
}
```

- Creating a pay-per-use 3-node GeminiDB Cassandra instance with 16 vCPUs and 64 GB of memory based on data restored using a specific backup

 NOTE

Values of **region** and **availability_zone** in the request body are only examples. Set them based on service requirements.

```
{
  "name": "test-cassandra-01",
  "datastore": {
    "type": "cassandra",
    "version": "3.11",
    "storage_engine": "rocksDB"
  },
  "region": "aaa",
  "availability_zone": "bbb",
  "vpc_id": "674e9b42-cd8d-4d25-a2e6-5abcc565b961",
  "subnet_id": "f1df08c5-71d1-406a-aff0-de435a51007",
  "security_group_id": "7aa51dbf-5b63-40db-9724-dad3c4828b58",
  "password": "*****",
  "mode": "Cluster",
  "flavor": [ {
    "num": 3,
    "storage": "ULTRAHIGH",
    "size": 500,
    "spec_code": "geminidb.cassandra.4xlarge.4"
  } ],
  "backup_strategy": {
    "start_time": "08:15-09:15",
    "keep_days": 8
  }
}
```



```

    },
    "ssl_option" : 1,
    "restore_info" : {
      "backup_id" : "2f4ddb93b9014b0893d81d2e472f30fe"
    }
  }
}

```

- Creating a yearly/monthly 3-node GeminiDB Cassandra instance with 16 vCPUs and 64 GB of memory based on the data of a specified instance at a specified point in time

 **NOTE**

Values of **region** and **availability_zone** in the request body are only examples. Set them based on service requirements.

```

{
  "name" : "test-cassandra-01",
  "datastore" : {
    "type" : "cassandra",
    "version" : "3.11",
    "storage_engine" : "rocksDB"
  },
  "region" : "aaa",
  "availability_zone" : "bbb",
  "vpc_id" : "674e9b42-cd8d-4d25-a2e6-5abcc565b961",
  "subnet_id" : "f1df08c5-71d1-406a-aff0-de435a51007",
  "security_group_id" : "7aa51dbf-5b63-40db-9724-dad3c4828b58",
  "password" : "*****",
  "mode" : "Cluster",
  "flavor" : [ {
    "num" : 3,
    "storage" : "ULTRAHIGH",
    "size" : 500,
    "spec_code" : "geminidb.cassandra.4xlarge.4"
  } ],
  "backup_strategy" : {
    "start_time" : "08:15-09:15",
    "keep_days" : 8
  },
  "ssl_option" : 1,
  "charge_info" : {
    "charge_mode" : "prePaid",
    "period_type" : "year",
    "period_num" : 3,
    "is_auto_renew" : true,
    "is_auto_pay" : true
  },
  "restore_info" : {
    "restore_time" : 1607731200000,
    "source_instance_id" : "054e292c9880d4992f02c0196d3ein12"
  }
}

```

Example Response

Status code: 202

Accepted

Creating a pay-per-use instance:

```

{
  "id" : "39b6a1a278844ac48119d86512e0000bin06",
  "name" : "test-cassandra-01",
  "datastore" : {
    "type" : "cassandra",
    "version" : "3.11",

```

```
"storage_engine": "rocksDB"
},
"created": "2019-10-28 14:10:54",
"status": "creating",
"region": "aaa",
"availability_zone": "bbb,ccc,ddd",
"vpc_id": "490a4a08-ef4b-44c5-94be-3051ef9e4fce",
"subnet_id": "0e2eda62-1d42-4d64-a9d1-4e9aa9cd994f",
"security_group_id": "2a1f7fc8-3307-42a7-aa6f-42c8b9b8f8c5",
"mode": "Cluster",
"flavor": [ {
  "num": 3,
  "size": 500,
  "storage": "ULTRAHIGH",
  "spec_code": "geminidb.cassandra.4xlarge.4"
} ],
"backup_strategy": {
  "start_time": "08:15-09:15",
  "keep_days": "8"
},
"ssl_option": "1",
"job_id": "c010abd0-48cf-4fa8-8cbc-090f093eaa2f"
}
```

Creating a yearly/monthly instance:

```
{
  "id": "39b6a1a278844ac48119d86512e0000bin06",
  "name": "test-cassandra-01",
  "datastore": {
    "type": "cassandra",
    "version": "3.11",
    "storage_engine": "rocksDB"
  },
  "created": "2019-10-28 14:10:54",
  "status": "creating",
  "region": "aaa",
  "availability_zone": "bbb,ccc,ddd",
  "vpc_id": "490a4a08-ef4b-44c5-94be-3051ef9e4fce",
  "subnet_id": "0e2eda62-1d42-4d64-a9d1-4e9aa9cd994f",
  "security_group_id": "2a1f7fc8-3307-42a7-aa6f-42c8b9b8f8c5",
  "mode": "Cluster",
  "flavor": [ {
    "num": 3,
    "size": 500,
    "storage": "ULTRAHIGH",
    "spec_code": "geminidb.cassandra.4xlarge.4"
  } ],
  "backup_strategy": {
    "start_time": "08:15-09:15",
    "keep_days": "8"
  },
  "enterprise_project_id": "0",
  "ssl_option": "1",
  "charge_info": {
    "charge_mode": "prePaid",
    "period_type": "year",
    "period_num": 3,
    "is_auto_renew": true,
    "is_auto_pay": true
  }
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.2 Deleting an Instance

Function

This API is used to delete an instance. Only pay-per-use instances can be deleted. Yearly/Monthly instances need to be unsubscribed from if they are no longer needed.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

URI

DELETE https://{Endpoint}/v3/{project_id}/instances/{instance_id}

Table 5-35 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-36 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 202

Table 5-37 Response body parameters

Parameter	Type	Description
job_id	String	Task ID.

Example Requests

URI example

```
DELETE https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/9136fd2a9fcd405ea4674276ce36dae8in02
```

Example Responses

Status code: 202

Accepted

```
{  "job_id" : "04efe8e2-9255-44ae-a98b-d87cae411890"}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.3 Querying Instances and Details

Function

This API is used to query instances and details based on specified conditions.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

URI

```
GET https://{Endpoint}/v3/{project_id}/instances
```

Table 5-38 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Table 5-39 Query parameters

Parameter	Mandatory	Type	Description
id	No	String	<p>Explanation: Instance ID.</p> <p>Constraints: If you enter an instance ID starting with an asterisk (*), fuzzy search results are returned. If you enter a valid instance ID, an exact result is returned.</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Parameter	Mandatory	Type	Description
name	No	String	<p>Explanation: Instance name.</p> <p>Constraints: If you enter an instance name starting with an asterisk (*), fuzzy search results are returned. If you enter a valid instance name, an exact result is returned.</p> <p>Values: N/A</p> <p>Default value: N/A</p>
datastore_type	No	String	<p>Explanation: Database type</p> <p>Constraints: N/A</p> <p>Values: cassandra, indicating that GeminiDB Cassandra instances are queried. mongodb, indicating that GeminiDB Mongo instances are queried. influxdb, indicating that GeminiDB Influx instances are queried. redis, indicating that GeminiDB Redis instances are queried.</p> <p>Default value: If this parameter is not transferred, all types of instances are queried.</p>

Parameter	Mandatory	Type	Description
mode	No	String	<p>Explanation: Instance type. The value can be:</p> <p>Constraints: N/A</p> <p>Values:</p> <ul style="list-style-type: none"> • Cluster: GeminiDB Cassandra, GeminiDB Influx, or GeminiDB Redis instance in a proxy cluster • RedisCluster: GeminiDB Redis instance in a Redis cluster • InfluxdbSingle: single-node GeminiDB Influx instance • ReplicaSet: GeminiDB Mongo instance in a replica set • The system ignores this parameter if datastore_type is not transferred. <p>Default value: N/A</p>
vpc_id	No	String	<p>Explanation: VPC ID. You can obtain the value with either of the following methods:</p> <ul style="list-style-type: none"> • Method 1: Log in to the VPC console and view the VPC ID on the VPC details page. • Method 2: Query the VPC ID using the VPC API. For details, see Querying VPCs. <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Parameter	Mandatory	Type	Description
subnet_id	No	String	<p>Explanation: Subnet ID. You can obtain the subnet ID with either of the following methods:</p> <ul style="list-style-type: none"> • Method 1: Log in to the VPC console and click the target subnet on the Subnets page. You can view the network ID on the displayed page. • Method 2: Query the subnet ID using the VPC API. For details, see Querying Subnets. <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>
offset	No	Integer	<p>Explanation: Index position. The query starts from the next instance creation time indexed by this parameter under a specified project. If offset is set to N, the resource query starts from the N+1 piece of data.</p> <p>Constraints: N/A</p> <p>Values: The value must be no less than 0.</p> <p>Default value: If this parameter is not transferred, the index offset is 0 by default, indicating that the query starts from the latest created instance.</p>

Parameter	Mandatory	Type	Description
limit	No	Integer	<p>Explanation: Maximum number of instances that can be queried.</p> <p>Constraints: N/A</p> <p>Values: 1 to 100</p> <p>Default value: If this parameter is not transferred, the first 100 instances are queried by default.</p>

Request Parameters

Table 5-40 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Response Parameters

Status code: 200

Table 5-41 Response body parameters

Parameter	Type	Description
instances	Array of ListInstancesResult objects	Explanation: Instance information.
total_count	Integer	Explanation: Total number of records. Values: N/A

Table 5-42 ListInstancesResult

Parameter	Type	Description
id	String	Explanation: Instance ID. Values: N/A
name	String	Explanation: Instance name. Values: N/A
status	String	Explanation: Instance status. Values: <ul style="list-style-type: none"> • normal, indicating that the instance is running normally. • abnormal, indicating that the instance is abnormal. • creating, indicating that the instance is being created. • frozen, indicating that the instance is frozen. • data_disk_full, indicating that the instance disk is full. • createfail, indicating that the instance failed to be created. • enlargefail, indicating that nodes failed to be added to the instance.

Parameter	Type	Description
port	String	Explanation: Database port. Values: N/A
mode	String	Explanation: Instance type Values: This parameter is the same as the corresponding request parameter.
region	String	Explanation: Region where the instance is deployed. Values: N/A
datastore	ListInstancesDatastoreResult object	Explanation: Database information.
engine	String	Explanation: Storage engine. Value: rocksDB
created	String	Explanation: Instance creation time. Values: N/A
updated	String	Explanation: Time when an instance is updated. Values: N/A
db_user_name	String	Explanation: Default username Value: rwuser
vpc_id	String	Explanation: VPC ID. Values: N/A

Parameter	Type	Description
subnet_id	String	Explanation: Subnet ID. Values: One GeminiDB Cassandra instance may use multiple subnets. For the subnet ID, see Table 5-47 .
security_group_id	String	Explanation: Security group ID. Values: N/A
backup_strategy	ListInstancesBackupStrategyResult object	Explanation: Backup policy.
pay_mode	String	Explanation: Billing mode. Options: Values: <ul style="list-style-type: none"> • 0: indicates the instance is billed on a pay-per-use basis. • 1: indicates the instance is billed on a yearly/monthly basis.
maintenance_window	String	Explanation: Maintenance time window. Values: N/A
groups	Array of ListInstancesGroupResult objects	Explanation: Group information.
enterprise_project_id	String	Explanation: Enterprise project ID. Value: 0 : default enterprise project
time_zone	String	Explanation: Time zone. Values: N/A

Parameter	Type	Description
actions	Array of strings	Explanation: Operation that is executed on the instance. Example values: <ul style="list-style-type: none"> • CREATE • REBOOT • RESTORE
dedicated_resource_id	String	Explanation: Dedicated resource ID. This parameter is returned only when the instance belongs to a dedicated resource pool. Values: N/A
lb_ip_address	String	Explanation: Load balancer IP address Constraints: This parameter is returned only when a load balancer IP address is assigned. Values: N/A
lb_port	String	Explanation: Load balancer port Constraints: This parameter is returned only when a load balancer IP address is assigned. Values: N/A
availability_zone	String	Explanation: AZ Values: N/A

Table 5-43 ListInstancesDatastoreResult

Parameter	Type	Description
type	String	Explanation: DB API. Values: N/A

Parameter	Type	Description
version	String	Explanation: DB version number. Values: N/A
patch_available	Boolean	Explanation: Whether the current instance can be patched Values: <ul style="list-style-type: none"> • true: A database can be upgraded through a patching API. • false: A database cannot be upgraded through a patching API.
whole_version	String	Explanation: Complete database version number Constraints: Currently, this parameter is available only for GeminiDB Cassandra API. Values: N/A

Table 5-44 ListInstancesBackupStrategyResult

Parameter	Type	Description
start_time	String	Explanation: Backup time window. Automated backup will be triggered during the backup time window. The current time is the UTC time. Values: N/A
keep_days	Integer	Explanation: Backup retention days Values: 0-35

Table 5-45 ListInstancesGroupResult

Parameter	Type	Description
id	String	Explanation: Group ID. Values: N/A
status	String	Explanation: Group status. Values: <ul style="list-style-type: none"> • normal, indicating that the group is normal. • abnormal, indicating that the group is abnormal. • creating, indicating that the group is being created. • createfail, indicating that the group failed to be created. • deleted: indicating that the group has been deleted. • resizefailed: indicating that the group specifications failed to be changed. • enlargefail: indicating the group failed to be scaled out.
volume	Volume object	Explanation: Volume information.
nodes	Array of ListInstancesNodeResult objects	Explanation: Node information.

Table 5-46 Volume

Parameter	Type	Description
size	String	Explanation: Storage space in GB. Values: N/A

Parameter	Type	Description
used	String	Explanation: Used storage space, in GB. Values: N/A

Table 5-47 ListInstancesNodeResult

Parameter	Type	Description
id	String	Explanation: Node ID. Values: N/A
name	String	Explanation: Node name. Values: N/A
status	String	Explanation: Node status. Values: <ul style="list-style-type: none"> • normal, indicating that the node is running normally. • abnormal, indicating that the node is abnormal. • creating, indicating that the node is being created. • createfail, indicating that the node failed to be created. • deleted, indicating that the node has been deleted. • resizefailed: indicating that the node specifications failed to be changed. • enlargefail: indicating nodes failed to be added.

Parameter	Type	Description
role	String	<p>Explanation: Node role.</p> <p>Constraints: This parameter is available only for GeminiDB MongoAPI replica set instances.</p> <p>Values: N/A</p>
subnet_id	String	<p>Explanation: ID of the subnet where the instance node is deployed.</p> <p>Values: N/A</p>
private_ip	String	<p>Explanation: Private IP address of a node</p> <p>Constraints: This parameter value is available after an ECS is created. Otherwise, the value is "".</p> <p>Values: N/A</p>
public_ip	String	<p>Explanation: Bound EIP</p> <p>Constraints: This parameter is valid only for nodes bound with EIPs.</p> <p>Values: N/A</p>
spec_code	String	<p>Explanation: Resource specification code. For the code, see the value of parameter flavors.spec_code in Querying Instance Specifications.</p> <p>Values: N/A</p>
availability_zone	String	<p>Explanation: AZ.</p> <p>Values: N/A</p>

Parameter	Type	Description
support_reduce	Boolean	<p>Explanation: Whether instance nodes can be deleted. The value can be:</p> <p>Values:</p> <ul style="list-style-type: none"> • true, indicating that instance nodes can be deleted. • false, indicating that instance nodes cannot be deleted.

Example Requests

- URI example

Querying all instances and details

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/0483b6b16e954cb88930a360d2c4e663/instances
```

- URI example

Querying instances and details based on specified conditions

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/0483b6b16e954cb88930a360d2c4e663/instances?offset=0&limit=10&id=ed7cc6166ec24360a5ed5c5c9c2ed726in06&name=hy&mode=Cluster&datastore_type=cassandra&vpc_id=19e5d45d-70fd-4a91-87e9-b27e71c9891f&subnet_id=bd51fb45-2dcb-4296-8783-8623bfe89bb7
```

Example Responses

Status code: 200

Success

```
{
  "instances": [ {
    "id": "8436a91546294036b75931e879882200in06",
    "name": "nosql-efa6",
    "status": "normal",
    "port": "8635",
    "mode": "Cluster",
    "region": "aaa",
    "datastore": {
      "type": "Cassandra",
      "version": "3.11",
      "whole_version": "3.11.3.11204",
      "patch_available": false
    },
    "engine": "rocksDB",
    "created": "2019-01-17T07:05:52",
    "updated": "2019-01-17T07:05:47",
    "db_user_name": "rwuser",
    "vpc_id": "674e9b42-cd8d-4d25-a2e6-5abcc565b961",
    "subnet_id": "f1df08c5-71d1-406a-aff0-de435a51007b",
    "security_group_id": "7aa51dbf-5b63-40db-9724-dad3c4828b58",
    "backup_strategy": {
      "start_time": "16:00-17:00",
      "keep_days": 7
    },
    "pay_mode": 0,
  }
]
```

```
"maintenance_window" : "02:00-06:00",
"groups" : [ {
  "id" : "0b0ff12541794e1084f6827e424be2d6gr06",
  "status" : "creating",
  "volume" : {
    "size" : 10,
    "used" : 0.33
  },
  "nodes" : [ {
    "id" : "233eaac9c6f245c0bb9c2d21eea12d1bno06",
    "name" : "nosql-efa6_priam_node_1",
    "status" : "normal",
    "subnet_id" : "f1df08c5-71d1-406a-aff0-de435a51007b",
    "private_ip" : "192.168.0.174",
    "spec_code" : "geminidb.redis.xlarge.4",
    "availability_zone" : "bbb"
  }, {
    "id" : "d57d76d6320a4a7b86db82c317550c4ano06",
    "name" : "nosql-efa6_priam_node_2",
    "status" : "normal",
    "subnet_id" : "f1df08c5-71d1-406a-aff0-de435a51007b",
    "private_ip" : "192.168.0.175",
    "spec_code" : "geminidb.redis.xlarge.4",
    "availability_zone" : "bbb"
  }, {
    "id" : "f46b0a1cf4d9400e9fd7af17f8742d37no06",
    "name" : "nosql-efa6_priam_node_3",
    "status" : "normal",
    "subnet_id" : "f1df08c5-71d1-406a-aff0-de435a51007b",
    "private_ip" : "192.168.0.176",
    "spec_code" : "geminidb.redis.xlarge.4",
    "availability_zone" : "bbb"
  }
]
}, {
  "enterprise_project_id" : "0",
  "availability_zone" : "cn-north-4a,cn-north-4b,cn-north-4c",
  "time_zone" : "",
  "actions" : [ "CREATE" ],
  "lb_ip_address" : "192.168.11.145",
  "lb_port" : "8635"
}, {
  "id" : "1236a91546294036b75931e879882200in02",
  "name" : "nosql-efa7",
  "status" : "normal",
  "port" : "8635",
  "mode" : "ReplicaSet",
  "region" : "aaa",
  "datastore" : {
    "type" : "ReplicaSet",
    "version" : "4.0",
    "patch_available" : false
  },
  "engine" : "rocksDB",
  "created" : "2019-01-17T07:05:52",
  "updated" : "2019-01-17T07:05:47",
  "db_user_name" : "rwuser",
  "vpc_id" : "674e9b42-cd8d-4d25-a2e6-5abcc565b961",
  "subnet_id" : "f1df08c5-71d1-406a-aff0-de435a51007b",
  "security_group_id" : "7aa51dbf-5b63-40db-9724-dad3c4828b58",
  "backup_strategy" : {
    "start_time" : "16:00-17:00",
    "keep_days" : 7
  },
  "pay_mode" : 0,
  "maintenance_window" : "02:00-06:00",
  "groups" : [ {
    "id" : "0b0ff12541794e1084f6827e424be2d1gr02",
    "status" : "normal",
    "volume" : {
```

```
"size" : 100,
"used" : 0.003
},
"nodes" : [ {
  "id" : "233eaac9c6f245c0bb9c2d21eea12d1bno02",
  "name" : "nosql-efa7_replica_node_2",
  "status" : "normal",
  "role" : "Primary",
  "subnet_id" : "f1df08c5-71d1-406a-aff0-de435a51007b",
  "private_ip" : "192.168.0.174",
  "public_ip" : "10.154.217.134",
  "spec_code" : "geminidb.mongodb.xlarge.4",
  "availability_zone" : "bbb"
}, {
  "id" : "d57d76d6320a4a7b86db82c317550c4ano02",
  "name" : "nosql-efa7_replica_node_1",
  "status" : "normal",
  "role" : "Secondary",
  "subnet_id" : "f1df08c5-71d1-406a-aff0-de435a51007b",
  "private_ip" : "192.168.0.39",
  "spec_code" : "geminidb.mongodb.xlarge.4",
  "availability_zone" : "bbb"
}, {
  "id" : "f46b0a1cf4d9400e9fd7af17f8742d37no02",
  "name" : "nosql-efa7_replica_node_3",
  "status" : "normal",
  "role" : "Secondary",
  "subnet_id" : "f1df08c5-71d1-406a-aff0-de435a51007b",
  "private_ip" : "192.168.0.176",
  "spec_code" : "geminidb.mongodb.xlarge.4",
  "availability_zone" : "bbb"
} ]
}],
"enterprise_project_id" : "0",
"availability_zone" : "cn-north-4a,cn-north-4b,cn-north-4c",
"time_zone" : "",
"actions" : [ ]
}],
"total_count" : 2
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.4 Scaling Up Storage Space of an Instance

Function

This API is used to scale up storage space of an instance.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

- GeminiDB Redis (proxy cluster)
- GeminiDB Redis (primary/standby)

This API supports both yearly/monthly and pay-per-use instances.

URI

POST https://{Endpoint}/v3/{project_id}/instances/{instance_id}/extend-volume

Table 5-48 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>
instance_id	Yes	String	<p>Explanation: Instance ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Request Parameters

Table 5-49 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token. You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Table 5-50 Request body parameters

Parameter	Mandatory	Type	Description
size	Yes	Integer	<p>Explanation: Requested storage space</p> <p>Constraints: It must be an integer greater than the current storage space.</p> <p>Values: The maximum storage space depends on the API type and specifications.</p> <ul style="list-style-type: none"> • For details about GeminiDB Cassandra instances, see Instance Specifications. • For details about GeminiDB Mongo instances, see Instance Specifications. • For details about GeminiDB Influx instances, see Instance Specifications. • For details about GeminiDB Redis instances, see Instance Specifications. <p>Default value: N/A</p>
is_auto_pay	No	String	<p>Explanation: Whether the order will be automatically paid after a yearly/monthly instance is created. This parameter does not affect the payment mode of automatic renewal.</p> <p>Constraints: N/A</p> <p>Values:</p> <ul style="list-style-type: none"> • true: An order is automatically paid from user's account. • false: An order is manually paid by a user. The default value is false. <p>Default value: false</p>

Response Parameters

Status code: 202

Table 5-51 Response body parameters

Parameter	Type	Description
job_id	String	<p>Explanation: Task ID. This parameter is returned only for pay-per-use instances.</p> <p>Values: N/A</p>
order_id	String	<p>Explanation: Order ID. This parameter is returned only when a yearly/monthly instance is created.</p> <p>Values: N/A</p>

Example Request

- URI example
POST `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/9136fd2a9fcd405ea4674276ce36dae8in06/extend-volume`
- Scaling up storage space of an instance to 550 GB

```
{
  "size" : 550
}
```

Example Response

Status code: 202

Accepted

Example response for a pay-per-use instance:

```
{
  "job_id" : "04efe8e2-9255-44ae-a98b-d87cae411890"
}
```

Example response for a yearly/monthly instance:

```
{
  "order_id" : "CS20070721568OVO9"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.5 Adding Nodes for an Instance

Function

This API is used to add nodes for a specified instance.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

The maximum and minimum storage space of each GeminiDB Redis instance depends on its node quantity and specifications. If the new storage space cannot meet the requirements for running the instance, nodes cannot be added.

This API supports both yearly/monthly and pay-per-use instances.

GeminiDB Influx does not allow you to add nodes to single-node instances.

GeminiDB Mongo allows you to add only read-only nodes for instances.

This API can be used to add nodes for two GeminiDB Redis instances between which there is a dual-active DR relationship.

URI

POST https://{Endpoint}/v3/{project_id}/instances/{instance_id}/enlarge-node

Table 5-52 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Parameter	Mandatory	Type	Description
instance_id	Yes	String	Explanation: Instance ID. Constraints: N/A Values: N/A Default value: N/A

Request Parameters

Table 5-53 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication . Constraints: N/A Values: N/A Default value: N/A

Table 5-54 Request body parameters

Parameter	Mandatory	Type	Description
num	Yes	Integer	<p>Explanation: Number of new nodes.</p> <p>Constraints: The value must be a positive integer.</p> <p>Values: N/A</p> <p>Default value: N/A</p>
subnet_id	No	String	<p>Explanation: ID of the subnet where the new node is deployed.</p> <p>Constraints:</p> <ul style="list-style-type: none"> • This parameter can be transferred when nodes are added to the GeminiDB Cassandra database instances. • The transferred subnet ID must belong to the VPC where the current instance is deployed. • If this parameter is not transferred, the system will allocate a subnet with sufficient IP addresses for the new node. <p>Values: N/A</p> <p>Default value: N/A</p>

Parameter	Mandatory	Type	Description
is_auto_pay	No	String	<p>Explanation: Whether the order will be automatically paid after a yearly/monthly instance is created. This parameter does not affect the payment mode of automatic renewal.</p> <p>Constraints: N/A</p> <p>Values:</p> <ul style="list-style-type: none"> • true: An order is automatically paid from user's account. • false: An order is manually paid by a user. The default value is false. <p>Default value: false</p>

Response Parameters

Status code: 202

Table 5-55 Response body parameters

Parameter	Type	Description
job_id	String	<p>Explanation: Task ID. This parameter is returned only for pay-per-use instances.</p> <p>Values: N/A</p>
order_id	String	<p>Explanation: Order ID. This parameter is returned only for yearly/monthly instances.</p> <p>Values: N/A</p>

Example Request

- URI example
POST `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/9136fd2a9fcd405ea4674276ce36dae8in06/enlarge-node`

- Adding a node

```
{  
  "num" : 1  
}
```

Example Response

Status code: 202

Accepted

Example response for a pay-per-use instance:

```
{  
  "job_id" : "3711e2ad-5787-49bc-a47f-3f0b066af9f5"  
}
```

Example response for a yearly/monthly instance:

```
{  
  "order_id" : "CS20070721568OVO9"  
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.6 Deleting Nodes from a Specified Instance

Function

This API is used to delete nodes from a specified instance.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra yearly/monthly and pay-per-use instances
- GeminiDB Redis yearly/monthly and pay-per-use instances
- You can change the billing mode of a GeminiDB Mongo instance from yearly/monthly to pay-per-use.

The maximum and minimum storage space of each GeminiDB Redis instance depends on its node quantity and specifications. If the new storage space cannot meet the requirements for running the instance, nodes cannot be removed.

GeminiDB Mongo allows you to delete only read-only nodes from instances.

This API can be used to delete nodes from two GeminiDB Redis instances between which there is a dual-active DR relationship.

URI

POST https://{Endpoint}/v3/{project_id}/instances/{instance_id}/reduce-node

Table 5-56 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-57 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-58 Request body parameters

Parameter	Mandatory	Type	Description
num	No	Integer	<p>Number of nodes to be deleted randomly.</p> <p>For GeminiDB Cassandra instances, the value ranges from 1 to 10.</p> <p>For GeminiDB Redis instances, the value is 1.</p> <p>NOTE If users connect to nodes using the client, do not choose to delete node randomly.</p>

Parameter	Mandatory	Type	Description
node_list	No	Array of strings	<p>ID of the node to be deleted. Make sure that the node can be deleted. If this parameter is not transferred, the number of nodes to be deleted is based on the internal policy of the system.</p> <p>NOTE</p> <ul style="list-style-type: none"> • Either num or node_list must be set. • If node_list is transferred, its value can contain 1 to 10 characters for GeminiDB Cassandra and contain 1 character for GeminiDB Redis. • If num and node_list are both transferred, the value of node_list takes effect. • If node_list is empty, instance nodes are deleted randomly. If node_list is not empty, only the node whose ID is specified is deleted. • Before a node is deleted, do not connect to the node directly to avoid service interruptions.

Response Parameters

Status code: 202

Table 5-59 Response body parameters

Parameter	Type	Description
job_id	String	Task ID. This parameter is returned only for pay-per-use instances.
order_id	String	Order ID. This parameter is returned only for yearly/monthly instances.

Example Requests

- URI example
 POST <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/9136fd2a9fcd405ea4674276ce36dae8in06/reduce-node>

- Deleting a node

```
{
  "num" : 1,
  "node_list" : [ "116ba14da34a42d28ecd83a38c218907no12" ]
}
```

Example Responses

Status code: 202

Accepted

Example response for a pay-per-use instance:

```
{
  "job_id" : "04efe8e2-9255-44ae-a98b-d87cae411890"
}
```

Example response for a yearly/monthly instance:

```
{
  "order_id" : "CS20070721568OVO9"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.7 Obtaining Sessions of a Node

Function

This API is used to obtain all sessions of a node.

Constraints

This API supports GeminiDB Redis instances.

URI

GET https://{Endpoint}/v3/{project_id}/redis/nodes/{node_id}/sessions

Table 5-60 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
node_id	Yes	String	Node ID.

Table 5-61 Query parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Index offset. The value must be no less than 0 . If this parameter is not transferred, the index offset is 0 by default, indicating that the query starts from the latest created session to the instance node.
limit	No	Integer	Number of pages in a pagination query. If this parameter is not transferred, 50 sessions are displayed on each page by default. A maximum of 100 sessions can be displayed.
addr_prefix	No	String	Prefix of the address on the user side. It is a character string consisting of an IP address and port number. If this parameter is not transferred, all prefixes of addresses on the user side are queried by default.

Request Parameters

Table 5-62 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.
X-Language	No	String	Language

Response Parameters

Status code: 200

Table 5-63 Response body parameters

Parameter	Type	Description
sessions	Array of objects	Instance sessions. For details, see Table 5-64 .

Parameter	Type	Description
total_count	Integer	Total sessions that meet search criteria.

Table 5-64 InstanceSession

Parameter	Type	Description
addr	String	IP address and port number of the client.
id	String	Client ID.
name	String	Client name, which is specified by running CLIENT SETNAME .
cmd	String	Last executed command.
age	String	Setup duration of the client connection, in seconds.
idle	String	Idle duration of the client connection, in seconds.
db	String	ID of the currently accessed database.
fd	String	File descriptor for sockets.
sub	String	Number of subscribed channels (Pub/Sub).
psub	String	Number of subscribed channels (Pub/Sub) in batches.
multi	String	Number of commands contained in a MULTI or EXEC transaction.

Example Requests

```
GET https://{Endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/redis/nodes/784b3fb7bac14bc490659950dd4f022fno12/sessions?offset=0&limit=20&addr_prefix=192.0.0.1:80
```

Example Responses

Status code: 200

Success

```
{
  "total_count": 100,
  "sessions": [ {
    "addr": "127.0.0.1:8080",
    "id": "254487",
    "name": "cli",
    "cmd": "get",
    "age": "8888581",
    "idle": "8888581",
    "db": "0",
    "fd": "1311",
```

```
"sub": "0",
"psub": "0",
"multi": "-1"
}]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.8 Querying Session Statistics of an Instance Node

Function

This API is used to query session statistics of an instance node.

Constraints

This API supports GeminiDB Redis instances.

URI

GET https://{Endpoint}/v3/{project_id}/redis/nodes/{node_id}/session-statistics

Table 5-65 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
node_id	Yes	String	Node ID.

Request Parameters

Table 5-66 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-67 Response body parameters

Parameter	Type	Description
total_connection_count	Integer	Total client connections.
active_connection_count	Integer	Number of active client connections.
top_source_ips	Array of objects	Top ten clients with the most connections. Total connected clients and their IP addresses are displayed and ordered based on sessions. For details, see Table 5-68 .
top_dbs	Array of objects	Top ten databases with the most connections. Total connected clients and their IP addresses are displayed and ordered based on sessions. For details, see Table 5-69 .

Table 5-68 SourceTopConnection

Parameter	Type	Description
client_ip	String	Client IP address.
connection_count	Integer	Number of client connections.

Table 5-69 TopDbConnection

Parameter	Type	Description
db	String	GeminiDB Redis database ID.
connection_count	Integer	Number of client connections.

Example Requests

```
GET https://{Endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/redis/nodes/784b3fb7bac14bc490659950dd4f022fno12/session-statistics
```

Example Responses

Status code: 200

Success

```
{
  "total_connection_count": 10,
  "active_connection_count": 5,
  "top_source_ips": [ {
```

```

"client_ip" : "127.0.0.1",
"connection_count" : 10
}, {
"client_ip" : "192.10.14.1",
"connection_count" : 9
}],
"top_dbs" : [ {
"db" : "12",
"connection_count" : 10
}, {
"db" : "14",
"connection_count" : 8
} ]
}

```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.9 Closing Sessions of an Instance Node

Function

This API is used to close sessions of an instance node.

Constraints

This API supports GeminiDB Redis instances and can be used to close inactive sessions of an instance node.

URI

DELETE https://{Endpoint}/v3/{project_id}/redis/nodes/{node_id}/sessions

Table 5-70 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
node_id	Yes	String	Node ID.

Request Parameters

Table 5-71 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-72 Request body parameters

Parameter	Mandatory	Type	Description
is_all	Yes	Boolean	Whether all sessions are closed.
session_ids	No	Array of strings	ID of the session to be closed. When the value of is_all is false , this parameter cannot be left empty.

Response Parameters

Status code: 200

No response parameters

Example Requests

- URI example
DELETE https://{Endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/redis/nodes/784b3fb7bac14bc490659950dd4f022fno12/sessions
- Deleting sessions **1131** and **2323**

```
{
  "is_all" : false,
  "session_ids" : [ "1131", "2323" ]
}
```

Example Responses

Status code: 200

Success

```
{}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.10 Querying Instance Specifications That Can Be Changed

Function

This API is used to query instance specifications that can be changed.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/available-flavors

Table 5-73 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>
instance_id	Yes	String	<p>Explanation: Instance ID</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Table 5-74 Query parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	<p>Explanation: Index position. The query starts from the next piece of data indexed by this parameter.</p> <p>Constraints: N/A</p> <p>Values: The value must be greater than or equal to 0.</p> <p>Default value: 0</p>
limit	No	Integer	<p>Explanation: Maximum records to be queried</p> <p>Constraints: N/A</p> <p>Values:</p> <ul style="list-style-type: none"> • The value ranges from 1 to 100. • If this parameter is not transferred, the first 100 records are queried by default. <p>Default value: 100</p>

Request Parameters

Table 5-75 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Response Parameters

Status code: 200

Table 5-76 Response body parameters

Parameter	Type	Description
instance_id	String	<p>Explanation: Instance ID</p> <p>Values: N/A</p>
instance_name	String	<p>Explanation: Instance name</p> <p>Values: N/A</p>
current_flavor	ComputeFlavor object	<p>Explanation: Instance specifications.</p>
optional_flavors	OptionalFlavorsInfo object	<p>Explanation: Available specification options that the instance specifications can be changed to.</p>

Table 5-77 ComputeFlavor

Parameter	Type	Description
vcpus	String	Explanation: Number of CPUs Values: N/A
ram	String	Explanation: Memory size in GB Values: N/A
spec_code	String	Explanation: Specification code Values: N/A
az_status	Map<String,String>	Explanation: AZ status Values: N/A
region_status	String	Explanation: Region status Values: N/A

Table 5-78 OptionalFlavorsInfo

Parameter	Type	Description
list	Array of objects in ComputeFlavor	Explanation: Available specification options
total_count	Integer	Explanation: Total records Values: N/A

Example Request

- URI example

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/0549b4a43100d4f32f51c01c2fe4acdb/instances/094424666ef04f79a2dfbe9f5b8b31a5in06/available-flavors
```

- Example request body
None

Example Response

Status code: 200

Success

```
{
  "instance_id": "094424666ef04f79a2dfbe9f5b8b31a5in06",
  "instance_name": "geminidb_instance_noreuse_0_ZKv2FSkxgoc3F8bGzsaxNg",
  "current_flavor": {
    "vcpus": "4",
    "ram": "16",
    "spec_code": "geminidb.cassandra.xlarge.4",
    "az_status": {
      "az2***": "unknown",
      "az1***": "normal",
      "az3***": "unknown"
    },
    "region_status": null
  },
  "optional_flavors": {
    "list": [ {
      "vcpus": "2",
      "ram": "8",
      "spec_code": "geminidb.cassandra.large.4",
      "az_status": {
        "az2***": "unknown",
        "az1***": "normal",
        "az4***": "normal",
        "az3***": "unknown"
      },
      "region_status": "normal"
    }, {
      "vcpus": "8",
      "ram": "32",
      "spec_code": "geminidb.cassandra.2xlarge.4",
      "az_status": {
        "az2***": "unknown",
        "az1***": "normal",
        "az3***": "unknown"
      },
      "region_status": "normal"
    }, {
      "vcpus": "16",
      "ram": "64",
      "spec_code": "geminidb.cassandra.4xlarge.4",
      "az_status": {
        "az2***": "unknown",
        "az1***": "normal",
        "az3***": "unknown"
      },
      "region_status": "normal"
    }, {
      "vcpus": "32",
      "ram": "128",
      "spec_code": "geminidb.cassandra.8xlarge.4",
      "az_status": {
        "az2***": "unknown",
        "az1***": "normal",
        "az3***": "unknown"
      },
      "region_status": "normal"
    }
  ]
}
```

```
  } ],  
  "total_count": 4  
 }  
 }
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.11 Changing Specifications of an Instance

Function

This API is used to change specifications of an instance.

NOTE

Services will be interrupted for 5 to 10 minutes when you change specifications of an instance. Exercise caution when performing this operation.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Redis

This API supports both yearly/monthly and pay-per-use instances.

This API can be used to scale up or down specifications of an instance.

The new specifications cannot be the same as the original specifications.

Specifications can be modified only when the instance status is **normal**.

If specifications cannot meet the requirements for running the instance, the specifications cannot be changed.

This API can be used to change specifications for two GeminiDB Redis instances between which there is a dual-active DR relationship.

URI

PUT https://{Endpoint}/v3/{project_id}/instances/{instance_id}/resize

Table 5-79 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>
instance_id	Yes	String	<p>Explanation: Instance ID</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Request Parameters

Table 5-80 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token</p> <p>You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Table 5-81 Request body parameters

Parameter	Mandatory	Type	Description
resize	Yes	ResizeInstanceOption object	Explanation: Target specification information.
is_auto_pay	No	String	Explanation: Whether the order will be automatically paid after a yearly/monthly instance is created. This parameter does not affect the payment mode of automatic renewal. Constraints: This parameter is valid only when specifications need to be expanded. Values: <ul style="list-style-type: none"> • true: An order is automatically paid from user's account. • false: An order is manually paid by a user. The default value is false. Default value: false

Table 5-82 ResizeInstanceOption

Parameter	Mandatory	Type	Description
target_spec_code	Yes	String	<p>Explanation:</p> <ul style="list-style-type: none"> Target resource specification code To obtain this value, see response values of flavors.spec_code in Querying Instance Specifications. <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Response Parameters

Status code: 202

Table 5-83 Response body parameters

Parameter	Type	Description
job_id	String	<p>Explanation: Task ID. This parameter is returned only for pay-per-use instances.</p> <p>Values: N/A</p>
order_id	String	<p>Explanation: Order ID. This parameter is returned only for yearly/monthly instances.</p> <p>Values: N/A</p>

Example Request

- URI example
PUT `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/9136fd2a9fcd405ea4674276ce36dae8in06/resize`
- Changing instance specifications to 16 vCPUs | 64 GB

```
{
  "resize": {
```

```
"target_spec_code" : "geminidb.cassandra.4xlarge.4"  
}  
}
```

Example Response

Status code: 202

Accepted

Example response for a pay-per-use instance:

```
{  
  "job_id" : "3711e2ad-5787-49bc-a47f-3f0b066af9f5"  
}
```

Example response for a yearly/monthly instance:

```
{  
  "order_id" : "CS20070721568OVO9"  
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.12 Resetting the Administrator Password of an Instance

Function

This API is used to reset the administrator password of an instance.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

This operation cannot be performed on frozen or abnormal instances.

Abnormal instances do not support this operation.

Only the password of user **rwuser** can be reset.

URI

PUT https://{Endpoint}/v3/{project_id}/instances/{instance_id}/password

Table 5-84 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-85 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-86 Request body parameters

Parameter	Mandatory	Type	Description
password	Yes	String	Database password. The password can include 8 to 32 characters and contain uppercase letters, lowercase letters, digits, and a combination of any two of the following special characters: ~!@#%^*-_+=? The password of GeminiDB Redis instances can contain at least two types of the following characters: uppercase letters, lowercase letters, digits, and special characters (~!@#%\$^&*()-_+=?). Enter a strong password against security risks such as brute force cracking.

Response Parameters

Status code: 204

No response parameters

Example Requests

- URI example
PUT https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/9136fd2a9fcd405ea4674276ce36dae8in06/password
- Resetting the administrator password of an instance to *****

```
{
  "password" : "*****"
}
```

Example Responses

Status code: 204

No Content

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.13 Editing the Name of an Instance

Function

This API is used to edit the name of an instance.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

The name of the instance that is being created or fails to be created cannot be edited.

URI

PUT https://{Endpoint}/v3/{project_id}/instances/{instance_id}/name

Table 5-87 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Parameter	Mandatory	Type	Description
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-88 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-89 Request body parameters

Parameter	Mandatory	Type	Description
name	Yes	String	New instance name. The name: Must start with a letter and can include 4 to 64 characters. It is case-sensitive and can contain only letters, digits, hyphens (-), and underscores (_).

Response Parameters

Status code: 204

No response parameters

Example Requests

- **URI example**
PUT `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/9136fd2a9fcd405ea4674276ce36dae8in06/name`
- **Changing the instance name to `myNewName`**

```
{
  "name": "myNewName"
}
```

Example Responses

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.14 Changing the Security Group of an Instance

Function

This API is used to change the security group associated with an instance.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

Abnormal instances do not support this operation.

Please confirm the modified security group rule. This policy may affect connections to the current instance, interrupting services.

URI

PUT https://{Endpoint}/v3/{project_id}/instances/{instance_id}/security-group

Table 5-90 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-91 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-92 Request body parameters

Parameter	Mandatory	Type	Description
security_group_id	Yes	String	ID of the new security group.

Response Parameters

Status code: 202

Table 5-93 Response body parameters

Parameter	Type	Description
job_id	String	Task ID.

Example Requests

- **URI example**
PUT `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/9136fd2a9fcd405ea4674276ce36dae8in02/security-group`
- **Example request body**

```
{
  "security_group_id": "73bed21a-708b-4985-b697-a96d0e0d2b39"
}
```

Example Responses

Status code: 202

No Content

```
{
  "job_id": "3711e2ad-5787-49bc-a47f-3f0b066af9f5"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.15 Upgrading Minor Version

Function

This API is used to upgrade the minor version of an instance.

Constraints

- This API supports the following types of instances:
 - GeminiDB Cassandra
 - GeminiDB Influx
 - GeminiDB Redis
- This API is not available to frozen or abnormal instances.
- This API is not available if there are abnormal instance nodes.
- View field **patch_available** in the result returned by the API for querying instance details and check whether a minor version upgrade is supported.
- Perform an upgrade during off-peak hours.

URI

POST https://{Endpoint}/v3/{project_id}/instances/{instance_id}/db-upgrade

Table 5-94 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-95 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 202

Table 5-96 Response body parameters

Parameter	Type	Description
job_id	String	Task ID.

Example Requests

URI example

POST https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/e73893ef73754465a8bd2e0857bbf13ein02/db-upgrade

Example Responses

Status code: 202

Accepted

```
{
  "job_id" : "3711e2ad-5787-49bc-a47f-3f0b066af9f5"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.16 Patching Databases in Batches

Function

This API is used to patch databases in batches.

Constraints

- This API supports the following types of instances:
 - GeminiDB Cassandra
 - GeminiDB Influx
 - GeminiDB Redis
- Upgrade is triggered immediately, so patch databases during off-peak hours.

URI

POST https://{Endpoint}/v3/{project_id}/instances/db-upgrade

Table 5-97 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-98 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-99 Request body parameters

Parameter	Mandatory	Type	Description
instance_ids	Yes	Array of strings	IDs of instances where databases need to be patched. Up to 10 IDs of an instance of the same GeminiDB engine type can be specified at a time.

Response Parameters

Status code: 202

Table 5-100 Response body parameters

Parameter	Type	Description
upgrade_results	Array of UpgradeResult objects	Batch upgrade result

Table 5-101 UpgradeResult

Parameter	Type	Description
job_id	String	Task ID returned only when a patch installation task is successfully submitted
instance_id	String	Instance ID
error_code	String	Error code returned only when a patch installation task fails to be submitted
error_message	String	Failure cause returned only when a patch installation task fails to be submitted

Status code: 400

Table 5-102 Response body parameters

Parameter	Type	Description
error_code	String	Error code
error_msg	String	Error message

Status code: 500

Table 5-103 Response body parameters

Parameter	Type	Description
error_code	String	Error code
error_msg	String	Error message

Example Request

URI example

POST <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/db-upgrade>

Example Request

Patching databases in batches

```
{
  "instance_ids" : [ "046287aae57843b1a7bc61b7a8812f41in13",
                    "3d1e04f49efa473a8c7eaf07ed7ff870in13" ]
}
```

Example Response

Status code: 202

Accepted

```
{
  "upgrade_results" : [ {
    "instance_id" : "046287aae57843b1a7bc61b7a8812f41in13",
    "job_id" : "e4616470-733d-41de-a9b0-a260709293d3"
  }, {
    "instance_id" : "3d1e04f49efa473a8c7eaf07ed7ff870in13",
    "error_code" : "DBS.200011",
    "error_message" : "The status of DB instance does not allow the operation."
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.17 Creating Cold Storage

Function

This API is used to create cold storage for an instance.

Constraints

This API supports only GeminiDB Influx cluster and GeminiDB Influx single-node instances.

This API supports both yearly/monthly and pay-per-use instances.

URI

POST `https://{Endpoint}/v3/{project_id}/instances/{instance_id}/cold-volume`

Table 5-104 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-105 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-106 Request body parameters

Parameter	Mandatory	Type	Description
size	Yes	Integer	Size in GB of the cold storage to be created. The minimum cold storage is 500 GB, and the maximum is 100,000 GB.

Parameter	Mandatory	Type	Description
is_auto_pay	No	String	Whether the order is paid automatically from your account when you buy cold storage for a yearly/monthly instance. This parameter does not affect the payment mode of automatic renewal. The value can be: <ul style="list-style-type: none"> • true, indicating that the order is automatically paid from the account. • false, indicating that the order is manually paid from the account. The default value is false.

Response Parameters

Status code: 202

Table 5-107 Response body parameters

Parameter	Type	Description
job_id	String	Task ID.
order_id	String	Order ID. This parameter is returned only when cold storage is created for yearly/monthly instances.

Example Requests

- URI example
POST <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/4d77eb5d9b9d407d88bbeba254b81aa0in13/cold-volume>
- Creating 500 GB of cold storage

```
{
  "size" : 500
}
```

Example Responses

Status code: 202

Accepted

```
{
  "job_id" : "f85104b5-4a9c-4e0f-9505-fc5409d8f7ae"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.18 Scaling Up Cold Storage

Function

This API is used to scale up cold storage of an instance.

Constraints

This API supports only GeminiDB Influx cluster and GeminiDB Influx single-node instances.

This API supports both yearly/monthly and pay-per-use instances.

URI

PUT `https://{Endpoint}/v3/{project_id}/instances/{instance_id}/cold-volume`

Table 5-108 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a user in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-109 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-110 Request body parameters

Parameter	Mandatory	Type	Description
size	Yes	Integer	<p>Size in GB of the new cold storage.</p> <ul style="list-style-type: none"> You can increase cold storage in increments of 1 GB and ensure that the new cold storage is an integer. The maximum cold storage can be up to 100,000 GB.
is_auto_pay	No	String	<p>Whether the order can be paid automatically from your account when you scale up cold storage of your instance. This parameter does not affect the payment mode of automatic renewal. The parameter value can be:</p> <ul style="list-style-type: none"> true, indicating that the order is automatically paid from the account. false, indicating that the order is manually paid from the account. The default value is false.

Response Parameters

Status code: 202

Table 5-111 Response body parameters

Parameter	Type	Description
job_id	String	Task ID.
order_id	String	Order ID. This parameter is returned only when storage space of a yearly/monthly instance is scaled up.

Example Requests

- URI example

```
PUT https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/944bdc55da6c4b33b260b34185ac86bein13/cold-volume
```

- Scaling up cold data storage space of an instance to 1,000 GB

```
{
  "size" : 1000
}
```

Example Responses

Status code: 202

Accepted

```
{
  "job_id" : "f85104b5-4a9c-4e0f-9505-fc5409d8f7ae"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.19 Binding/Unbinding an EIP

Function

This API is used to bind an EIP to a node of an instance or unbind an EIP from the node.

Constraints

- This API supports the following types of instances:
 - GeminiDB Cassandra
 - GeminiDB Mongo
 - GeminiDB Influx
 - GeminiDB Redis
- This API is not available to frozen or abnormal instances.
- Multiple EIPs cannot be bound to the same node.
- Unbinding an EIP from a node of a frozen instance is not supported.
- Unbinding an EIP from a node with no EIPs bound is not supported.

URI

POST https://{Endpoint}/v3/{project_id}/instances/{instance_id}/nodes/{node_id}/public-ip

Table 5-112 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.
node_id	Yes	String	Instance node ID.

Request Parameters

Table 5-113 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-114 Request body parameters

Parameter	Mandatory	Type	Description
action	Yes	String	Operation identifier. The value can be: <ul style="list-style-type: none"> BIND, indicating that an EIP is bound. UNBIND, indicating that an EIP is unbound.
public_ip	No	String	EIP. This parameter is mandatory when an EIP needs to be bound.
public_ip_id	No	String	EIP ID. This parameter is mandatory when an EIP needs to be bound.

Response Parameters

Status code: 202

Table 5-115 Response body parameters

Parameter	Type	Description
job_id	String	Task ID.

Example Requests

- Binding an EIP

- URI example

```
POST https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/61f554065802400c9c33f87e8114f081in10/nodes/1345bb07d1834f8fb0b4acbc26e989aano10/public-ip
```

- Example request body

```
{  "action": "BIND",  "public_ip": "10.154.218.161",  "public_ip_id": "45da4782-e0c8-4aa4-a290-b8740014f710"}
```

- Unbinding an EIP

- URI example

```
POST https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/61f554065802400c9c33f87e8114f081in10/nodes/1345bb07d1834f8fb0b4acbc26e989aano10/public-ip
```

- Example request body

```
{  "action": "UNBIND",  "public_ip": "10.154.218.161"}
```

Example Responses

Status code: 202

Accepted

```
{  "job_id": "f85104b5-4a9c-4e0f-9505-fc5409d8f7ae"}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.20 Enabling or Disabling SSL

Function

This API is used to enable or disable SSL.

Constraints

- This API supports the following types of instances:
 - GeminiDB Mongo
 - GeminiDB Influx
 - GeminiDB Redis

- Enabling or disabling SSL is not supported for frozen or abnormal instances.
- The instance will be restarted after SSL is enabled or disabled on it. Exercise caution when you enable or disable SSL.

URI

POST https://{Endpoint}/v3/{project_id}/instances/{instance_id}/ssl-option

Table 5-116 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a user in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-117 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-118 Request body parameters

Parameter	Mandatory	Type	Description
ssl_option	Yes	String	Whether SSL is enabled. The value can be: <ul style="list-style-type: none"> • on, indicating that SSL is enabled by default. • off, indicating that SSL is not enabled by default.

Response Parameters

Status code: 202

Table 5-119 Response body parameters

Parameter	Type	Description
job_id	String	Task ID.

Example Requests

- Enabling SSL

- URI example

```
POST https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/944bdc55da6c4b33b260b34185ac86bein13/ssl-option
```

- Enabling SSL

```
{  "ssl_option" : "on"}
```

- Disabling SSL

- URI example

```
POST https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/944bdc55da6c4b33b260b34185ac86bein13/ssl-option
```

- Disabling SSL

```
{  "ssl_option" : "off"}
```

Example Responses

Status code: 202

Accepted

```
{  "job_id" : "f85104b5-4a9c-4e0f-9505-fc5409d8f7ae"}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.21 Restarting an Instance

Function

This API is used to restart an instance.

Constraints

- This API supports the following types of instances:
 - GeminiDB Cassandra
 - GeminiDB Mongo
 - GeminiDB Influx
 - GeminiDB Redis
- If the instance status is not normal, the instance cannot be restarted.

URI

POST https://{Endpoint}/v3/{project_id}/instances/{instance_id}/restart

Table 5-120 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>
instance_id	Yes	String	<p>Explanation: Instance ID</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Request Parameters

Table 5-121 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Response Parameters

Status code: 202

Table 5-122 Response body parameters

Parameter	Type	Description
job_id	String	<p>Explanation: Task ID.</p> <p>Values: N/A</p>

Example Request

- URI example
POST `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/944bdc55da6c4b33b260b34185ac86bein13/restart`
- Example request body
{ }

Example Response

Status code: 202

Accepted

```
{
  "job_id" : "f85104b5-4a9c-4e0f-9505-fc5409d8f7ae"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.22 Configuring an Autoscaling Policy for Storage Space

Function

This API is used to configure an autoscaling policy for storage space.

Constraints

- This API supports GeminiDB Cassandra and GeminiDB Redis instances.
- This API supports both pay-per-use and yearly/monthly instances.
- For yearly/monthly instances, the system bills new storage space automatically by default.
- Autoscaling is available only when your account balance is sufficient.
- If the instance status is not normal, autoscaling of storage space cannot be configured.

URI

PUT https://{Endpoint}/v3/{project_id}/instances/disk-auto-expansion

Table 5-123 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-124 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-125 Request body parameters

Parameter	Mandatory	Type	Description
instance_ids	Yes	Array of strings	IDs of the instances where autoscaling is enabled for storage space. Up to 50 instances are supported.
switch_option	No	String	Whether autoscaling is enabled. The value can be: <ul style="list-style-type: none"> • on: Autoscaling is enabled for storage space. • off: Autoscaling is disabled for storage space. The default value is on .
policy	Yes	Array of diskAutoExpansionPolicy	Autoscaling policies for storage space.

Table 5-126 diskAutoExpansionPolicy

Parameter	Mandatory	Type	Description
threshold	No	Integer	Threshold for triggering autoscaling. <ul style="list-style-type: none"> • GeminiDB Cassandra instances <ul style="list-style-type: none"> - The value can be 80, 85, or 90. - The default threshold is 90, indicating that autoscaling is enabled when the used storage space exceeds 90% of total storage space or the available storage space is less than 10 GB. • GeminiDB Redis instances <ul style="list-style-type: none"> - The value can be 60, 65, 70, 75, 80, 85, and 90. - The default threshold is 80, indicating that autoscaling is enabled when the used storage space exceeds 80% of total storage space.

Parameter	Mandatory	Type	Description
step	No	Integer	<p>Autoscaling step (s%).</p> <ul style="list-style-type: none"> • GeminiDB Cassandra instances <ul style="list-style-type: none"> - The value can be 10, 15, or 20, and the default value is 10. - After autoscaling is enabled, storage space will increase by s% automatically. • GeminiDB Redis instances <ul style="list-style-type: none"> - The value can be 10, 15, or 20, and the default value is 20. - When the storage usage is greater than 98%: If the total storage is less than 600 GB, the storage usage after autoscaling (used storage space/total storage space) will be less than 85%. If the total storage is greater than or equal to 600 GB, the system automatically scales up the storage space by over 90 GB. <p>NOTE</p> <ul style="list-style-type: none"> • GeminiDB Cassandra instances <ul style="list-style-type: none"> • If the autoscaling step is not a multiple of 10, round it up. • The value after the decimal point is rounded. The minimum step is 100 GB by default. • GeminiDB Redis instances <ul style="list-style-type: none"> • The value after the decimal point is rounded. The minimum step is 1 GB by default. • If there is insufficient balance in your account, storage space may fail to be scaled up for yearly/monthly instances.

Parameter	Mandatory	Type	Description
size	No	Integer	<p>Storage limit in GB that autoscaling can increase storage space to.</p> <ul style="list-style-type: none"> • GeminiDB Cassandra instances <ul style="list-style-type: none"> - Maximum amount that the system can automatically scale up an instance's storage space to must be greater than or equal to 100 GB of the current DB instance storage space. The value must be no less than the total storage of the instance and cannot exceed its maximum storage. - 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. • GeminiDB Redis instances <ul style="list-style-type: none"> - 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.

Response Parameters

Status code: 204

No response parameters

Example Request

- Enabling the autoscaling policy of storage space
 - URI example
PUT <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/disk-auto-expansion>
 - Enabling autoscaling for storage space (Set **threshold** to **90**, **step** to **10**, and **size** to **600**.)


```
{
  "instance_ids" : [ "93e4b3eda14349b1b870f72829bc3b9bin06" ],
  "policy" : {
    "threshold" : 90,
    "step" : 10,
    "size" : 600
  }
}
```

- Disabling the autoscaling policy of storage space

- URI example

```
PUT https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/disk-auto-expansion
```

- Disabling the autoscaling policy of storage space

```
{
  "instance_ids": [
    "93e4b3eda14349b1b870f72829bc3b9bin06"
  ],
  "switch_option": "off",
  "policy": {
    "threshold": 90,
    "step": 10,
    "size": 600
  }
}
```

Example Response

Status code: 204

No Content

```
{}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.23 Changing a Database Port

Function

This API is used to change the database port of an instance.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Redis

This operation is not supported for frozen or abnormal instances.

GeminiDB Redis and GeminiDB Cassandra do not allow you to change the database port of a DR instance. To change the database port, you have to change the port of the primary instance first. The load balancer port of a GeminiDB Redis instance is changed accordingly.

The instance must be restarted to make changes take effect. Exercise caution when modifying the database port.

URI

PUT https://{Endpoint}/v3/{project_id}/instances/{instance_id}/port

Table 5-127 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>
instance_id	Yes	String	<p>Explanation: Instance ID, which can be obtained by calling the API described in Querying Instances and Details. If there are no instances available, call the API described in Creating an Instance to create one.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Request Parameters

Table 5-128 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Table 5-129 Request body parameters

Parameter	Mandatory	Type	Description
port	Yes	Integer	<p>Explanation: New database port number.</p> <p>Constraints: N/A</p> <p>Values: The database port number of GeminiDB Mongo replica sets 4.0 ranges from 2100 to 9500, except 8636, 8637, and 8638. The database port number of GeminiDB Cassandra instances ranges from 2100 to 9500, except 7000, 7001, 7199, 8636, 8479, 8484, 8999, 8018, 2180, 2887, 3887, 8079, 8091, and 8092. The port number of the GeminiDB Redis instance ranges from 1024 to 65535, except 2180, 2887, 3887, 6377, 6378, 6380, 8018, 8079, 8091, 8479, 8484, 8999, 9864, 9866, 9867, 12017, 12333 and 50069.</p> <p>Default value: N/A</p>

Response Parameters

Status code: 202

Table 5-130 Response body parameters

Parameter	Type	Description
job_id	String	<p>Explanation: Task ID</p> <p>Values: N/A</p>

Example Request

- **URI example**
PUT `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/c865f921f3dd45198f209a607533a779in06/port`
- **Changing the database port number to 8888**

```
{
  "port": 8888
}
```

Example Response

Status code: 202

Accepted

```
{
  "job_id": "89638f5e-0780-497c-b3c0-4d0968383e19"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.24 Checking Password Strength

Function

This API is used to check whether the password is weak.

URI

POST `https://{Endpoint}/v3/{project_id}/weak-password-verification`

Table 5-131 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-132 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-133 Request body parameters

Parameter	Mandatory	Type	Description
password	Yes	String	Database password.

Response Parameters

Status code: 200

Table 5-134 Response body parameters

Parameter	Type	Description
weak	Boolean	Whether the password is a weak password. <ul style="list-style-type: none"> true: It is a weak password. false: It is not a weak password.

Example Requests

- **URI example**
POST `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/weak-password-verification`
- **Checking Password Strength**

```
{
  "password" : "xxxx"
}
```

Example Responses

Status code: 200

Success

```
{
  "weak" : false
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.25 Configuring Access to a Replica Set Across CIDR Blocks

Function

This API is used to configure access to a replica set across CIDR blocks.

Constraints

This API supports the following types of instances:

- GeminiDB Mongo

This operation is not supported for frozen or abnormal instances.

URI

POST `https://{Endpoint}/v3/{project_id}/instances/{instance_id}/client-network`

Table 5-135 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID, which can be obtained by calling the API described in Querying Instances and Details . If there are no instances available, call the API described in Creating an Instance to create one.

Request Parameters

Table 5-136 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-137 Request body parameters

Parameter	Mandatory	Type	Description
client_network_ranges	Yes	Array of strings	<p>CIDR block where the client is located.</p> <ul style="list-style-type: none"> • Access across CIDR blocks is required only when the CIDR blocks of the client and the replica set instance are different. For example, if the client CIDR block is 192.168.0.0/16 and the replica set instance's CIDR block is 172.16.0.0/24, add the CIDR block 192.168.0.0/16 so that the client can access the replica set instance. • For example, if the source network segment is 192.168.0.0/xx, the value of xx must range from 8 to 32. • To ensure the ECS and the instance can communicate with each other, configure the connection by referring to VPC Peering Connection Overview.

Response Parameters

Status code: 202

Table 5-138 Response body parameters

Parameter	Type	Description
job_id	String	Task ID.

Example Requests

- URI example
POST `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/9136fd2a9fcd405ea4674276ce36dae8in10/client-network`
- Changing the CIDR block where the client is located to 192.168.0.0/16

```
{
  "client_network_ranges": [ "192.168.0.0/16" ]
}
```


Example Responses

Status code: 202

Accepted

```
{
  "job_id" : "89638f5e-0780-497c-b3c0-4d0968383e19"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.26 Deleting the Node that Fails to Be Added

Function

This API is used to delete the node that fails to be added to an instance.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Redis

URI

DELETE https://{Endpoint}/v3/{project_id}/instances/{instance_id}/enlarge-failed-nodes

Table 5-139 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID, which can be obtained by calling the API described in Querying Instances and Details . If there are no instances available, call the API described in Creating an Instance to create one.

Request Parameters

Table 5-140 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-141 Request body parameters

Parameter	Mandatory	Type	Description
node_id	Yes	String	Node ID.

Response Parameters

Status code: 202

Table 5-142 Response body parameters

Parameter	Type	Description
job_id	String	Task ID.

Example Requests

- URI example
DELETE https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/c865f921f3dd45198f209a607533a779in06/enlarge-failed-nodes
- Example request body

```
{
  "node_id" : "b60f00f19cd044fc8d7b52908978f629no06"
}
```

Example Responses

Status code: 202

Accepted

```
{
  "job_id" : "89638f5e-0780-497c-b3c0-4d0968383e19"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.27 Querying IP Addresses Required for Creating an Instance or Adding Nodes

Function

This API is used to query IP addresses required for creating an instance or adding nodes to an instance.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

URI

GET https://{Endpoint}/v3/{project_id}/ip-num-requirement

Table 5-143 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Table 5-144 Query parameters

Parameter	Mandatory	Type	Description
node_num	Yes	Integer	Nodes required for creating or scaling out an instance. The maximum value is 200 .

Parameter	Mandatory	Type	Description
engine_name	No	String	DB API name. If no instance ID is transferred, this parameter is mandatory. The value can be: <ul style="list-style-type: none"> • cassandra: GeminiDB Cassandra API • mongodb: GeminiDB Mongo API • influxdb: GeminiDB Influx API • redis: GeminiDB Redis API
instance_mode	No	String	Instance type. The value can be: If no instance ID is transferred, this parameter is mandatory. The value can be: <ul style="list-style-type: none"> • Cluster: GeminiDB Cassandra, GeminiDB Influx, or GeminiDB Redis instance in a proxy cluster • ReplicaSet: GeminiDB Mongo instance in a replica set
instance_id	No	String	Instance ID, which can be obtained by calling the API described in Querying Instances and Details . If there are no instances available, call the API described in Creating an Instance to create one.

Request Parameters

Table 5-145 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-146 Response body parameters

Parameter	Type	Description
count	Integer	Number of IP addresses used.

Example Requests

- URI example
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/ip-num-requirement?node_num=3&engine_name=cassandra&instance_mode=Cluster
- Example request body
None

Example Responses

Status code: 200

Success

```
{
  "count" : 3
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.28 Querying the Autoscaling Policy of Storage Space

Function

This API is used to query the autoscaling policy of storage space.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/disk-auto-expansion

Table 5-147 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID, which can be obtained by calling the API described in Querying Instances and Details . If there are no instances available, call the API described in Creating an Instance to create one.

Request Parameters

Table 5-148 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-149 Response body parameters

Parameter	Type	Description
policy	object	Autoscaling policy for storage space. No information is returned if the autoscaling policy is disabled. For details, see Table 5-150 .

Table 5-150 AutoEnlargePolicy

Parameter	Type	Description
threshold	Integer	Threshold for triggering autoscaling.

Parameter	Type	Description
step	Integer	Percentage increase (step%). When autoscaling is triggered, the database system automatically scales up the current storage space of your instance by step%. If the increased storage space is not a multiple of 10 GB, the system rounds it up to the nearest multiple of 10 GB. The default minimum increment is 100 GB. NOTE If there is insufficient balance in your account, storage space may fail to be scaled up for yearly/monthly instances.
size	Integer	Storage limit in GB that autoscaling can increase storage space to.

Example Request

- URI example

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/93e4b3eda14349b1b870f72829bc3b9bin06/disk-auto-expansion
```

- Example request body

None

Example Response

Status code: 200

Success

```
{
  "policy": {
    "threshold": 90,
    "step": 10,
    "size": 600
  }
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.29 Scaling Storage Space of an Instance

Function

This API is used to scale storage space of an instance.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Redis (proxy cluster)
- GeminiDB Redis (primary/standby)

This API supports both yearly/monthly and pay-per-use instances.

URI

PUT `https://{Endpoint}/v3/{project_id}/instances/{instance_id}/volume`

Table 5-151 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>
instance_id	Yes	String	<p>Explanation: Instance ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Request Parameters

Table 5-152 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication . Constraints: N/A Values: N/A Default value: N/A

Table 5-153 Request body parameters

Parameter	Mandatory	Type	Description
size	Yes	Integer	<p>Explanation: Requested storage space in GB. The value is an integer.</p> <p>Constraints: To scale up storage, ensure the new storage space is greater than the current one. To scale down storage, ensure the new storage space is at least 1.25 times more than the used space and rounded up.</p> <p>Values: The maximum and minimum storage space depends on the API type and specifications.</p> <ul style="list-style-type: none"> • For details about GeminiDB Cassandra instances, see Instance Specifications. • For details about GeminiDB Redis instances, see Instance Specification. <p>Default value: N/A</p>
is_auto_pay	No	Boolean	<p>Explanation: Whether the order is automatically paid from your account when you scale up the storage of a yearly/ monthly instance. This parameter does not affect the payment mode of automatic renewal.</p> <p>Constraints: N/A</p> <p>Values:</p> <ul style="list-style-type: none"> • true: An order is automatically paid from user's account. • false: An order is manually paid by a user. The default value is false. <p>Default value: false</p>

Response Parameters

Status code: 200

Table 5-154 Response body parameters

Parameter	Type	Description
job_id	String	<p>Explanation: Task ID. This parameter is returned only for pay-per-use instances.</p> <p>Values: N/A</p>
order_id	String	<p>Explanation: Order ID. This parameter is returned only when storage space of a yearly/monthly instance is scaled.</p> <p>Values: N/A</p>

Example Request

- URI example
PUT `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/9136fd2a9fcd405ea4674276ce36dae8in02/volume`
- Changing storage space of an instance to 550 GB

```
{
  "size" : 550
}
```

Example Response

Status code: 200

Success

Example response for a pay-per-use instance:

```
{
  "job_id" : "04efe8e2-9255-44ae-a98b-d87cae411890"
}
```

Example response for a yearly/monthly instance:

```
{
  "order_id" : "CS20070721568OVO9"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.30 Querying High-Risk Commands

Function

This API is used to query high-risk commands supported by a GeminiDB Redis instance.

Constraints

This API supports only proxy-based general-purpose GeminiDB Redis instances.

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/high-risk-commands

Table 5-155 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-156 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-157 Response body parameters

Parameter	Type	Description
commands	Array of objects	High-risk commands and their corresponding renamed commands. For details, see Table 5-158 .

Table 5-158 CommandInfo

Parameter	Type	Description
origin_name	String	Original high-risk commands. The following commands are supported: KEYS, FLUSHDB, FLUSHALL, HGETALL, HKEYS, HVALS, and SMEMBERS.
name	String	Name of the command that takes effect. If this parameter is left blank, the command is disabled. The name can include 0 to 30 characters, including digits, uppercase letters, lowercase letters, and underscores (_). Minimum characters: 0 Maximum characters: 30

Example Requests

```
GET https://{Endpoint}/v3/054e292c9880d4992f02c0196d3ea468/instances/41409e5a49ee467287d738b9b6d0adcbin12/high-risk-commands
```

Example Responses

Status code: 200

Success

```
{
  "commands" : [ {
    "origin_name" : "keys",
    "name" : "newKeys"
  }, {
    "origin_name" : "flushdb",
    "name" : "newFlushdb"
  }, {
    "origin_name" : "flushall",
    "name" : "newFlushall"
  }, {
    "origin_name" : "hgetall",
    "name" : "newHgetall"
  }, {
    "origin_name" : "hkeys",
    "name" : "newHkeys"
  }, {
    "origin_name" : "hvals",
    "name" : "newHvals"
  }, {
    "origin_name" : "smembers",
    "name" : "newSmembers"
  }
]
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.31 Modifying High-Risk Commands

Function

This API is used to modify high-risk commands supported by a GeminiDB Redis instance in batches.

Constraints

This API supports only proxy-based general-purpose GeminiDB Redis instances.

URI

PUT https://{Endpoint}/v3/{project_id}/instances/{instance_id}/high-risk-commands

Table 5-159 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-160 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-161 Request body parameters

Parameter	Mandatory	Type	Description
commands	Yes	Array of objects	High-risk commands and their corresponding renamed commands. For details, see Table 5-162 .

Table 5-162 CommandInfo

Parameter	Mandatory	Type	Description
origin_name	Yes	String	Original high-risk commands. The following commands are supported: KEYS, FLUSHDB, FLUSHALL, HGETALL, HKEYS, HVALS, and SMEMBERS.
name	Yes	String	Name of the command that takes effect. If this parameter is left blank, the command is disabled. The value can include 0 to 30 characters, including digits, uppercase letters, lowercase letters, and underscores (_). Minimum characters: 0 Maximum characters: 30

Response Parameters

Status code: 200

No response parameters

Example Requests

- URI example
PUT
`https://{Endpoint}/v3/054e292c9880d4992f02c0196d3ea468/instances/41409e5a49ee467287d738b9b6d0adcb12/high-risk-commands`

- Modifying high-risk commands

```
{
  "commands": [ {
    "origin_name": "keys",
    "name": "newKeys"
  }, {
    "origin_name": "flushdb",
    "name": "newFlushdb"
  }, {
    "origin_name": "flushall",
    "name": "newFlushall"
  }, {
    "origin_name": "hgetall",
    "name": "newHgetall"
  }, {
    "origin_name": "hkeys",
    "name": "newHkeys"
  }, {
    "origin_name": "hvals",
    "name": "newHvals"
  }, {
    "origin_name": "smembers",
    "name": "newSmembers"
  }
]
```

Example Responses

Status code: 200

Success

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.32 Querying Hot Keys of a Redis Instance

Function

This API is used to query hot keys of a Redis instance.

Constraints

This API supports only GeminiDB Redis instances.

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/hot-keys

Table 5-163 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Table 5-164 Query parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Index position. If offset is set to N, the resource query starts from the N+1 piece of data. The value must be no lower than 0. If this parameter is not transferred, the value is 0 by default.

Parameter	Mandatory	Type	Description
limit	No	Integer	Maximum records to be queried. <ul style="list-style-type: none"> The value ranges from 1 to 50. If this parameter is not transferred, the first 50 instances are queried by default.

Request Parameters

Table 5-165 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-166 Response body parameters

Parameter	Type	Description
keys	Array of objects in Table 5-167	Queried hot keys.
count	Integer	Total number.

Table 5-167 HotKeysInfoRspResponseBody

Parameter	Type	Description
name	String	Hot key name.
type	String	Hot key type.
command	String	Hot key command.
qps	Integer	Hot key QPS.
db_id	Integer	Database where a hot key is located.

Status code: 400

Table 5-168 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Status code: 500

Table 5-169 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Example Requests

Querying the first 30 keys (hot keys) in a Redis instance
 GET https://{endpoint}/v30549b4a43100d4f32f51c01c2fe4acdb/instances/
 91b5c48e11bf4913b35814f28b8f79f2in12/hot-keys

Example Responses

Status code: 200

Success.

```
{
  "keys": [ {
    "name": "address",
    "type": "String",
    "command": "set",
    "qps": 3000,
    "db_id": 0
  } ],
  "count": 1
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.33 Disabling Commands for a Redis Instance

Function

This API is used to disable commands for a Redis instance.

Constraints

This API supports only GeminiDB Redis instances.

URI

POST https://{Endpoint}/v3/{project_id}/redis/instances/{instance_id}/disabled-commands

Table 5-170 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID

Request Parameters

Table 5-171 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-172 Request body parameters

Parameter	Mandatory	Type	Description
disabled_type	Yes	String	Disabled type. Values: <ul style="list-style-type: none"> command key

Parameter	Mandatory	Type	Description
commands	No	Array of strings	This parameter is transferred when disabled_type is set to command . Values: <ul style="list-style-type: none"> • keys • hkeys • hvals • hgetall • smembers • flushdb • flushall
keys	No	Array of objects in Table 5-173	This parameter is transferred when disabled_type is set to key . A maximum of 20 keys are allowed.

Table 5-173 RedisDisabledCommandsDetail

Parameter	Mandatory	Type	Description
db_id	Yes	String	Database where a key is located.
key	Yes	String	Key name.
commands	Yes	Array of strings	Command list.

Response Parameters

Status code: 400

Table 5-174 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Status code: 500

Table 5-175 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Example Requests

Disabling commands for a Redis instance

POST https://{endpoint}/v3/0549b4a43100d4f32f51c01c2fe4acdb/redis/instances/0a9a2be86bdf493981d17626ad2f8668in12/disabled-commands

```
{
  "disabled_type": "key",
  "keys": [ {
    "db_id": "0",
    "key": "name",
    "commands": [ "get" ]
  }, {
    "db_id": "0",
    "key": "address",
    "commands": [ "get", "set" ]
  } ]
}
```

Example Responses

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.34 Querying Disabled Commands for a Redis Instance

Function

This API is used to query disabled commands for a Redis instance.

Constraints

This API supports only GeminiDB Redis instances.

URI

GET https://{Endpoint}/v3/{project_id}/redis/instances/{instance_id}/disabled-commands

Table 5-176 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Table 5-177 Query parameters

Parameter	Mandatory	Type	Description
type	Yes	String	Disabled type. Values: <ul style="list-style-type: none"> • command • key
offset	No	Integer	Index position. If offset is set to N, the resource query starts from the N+1 piece of data. The value must be no lower than 0. If this parameter is not transferred, the value is 0 by default.
limit	No	Integer	Maximum dedicated resources to be queried. <ul style="list-style-type: none"> • The value ranges from 1 to 50. If this parameter is not transferred, the first 50 instances are queried by default.

Request Parameters

Table 5-178 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-179 Response body parameters

Parameter	Type	Description
total_count	String	Total number.
disabled_type	String	Command type. Values: <ul style="list-style-type: none"> • command • key
commands	Array of strings	This parameter is displayed when disabled_type is set to command . Values: <ul style="list-style-type: none"> • keys • hkeys • hvals • hgetall • smembers • flushdb • flushall
keys	Array of objects in Table 5-180	This parameter is displayed when disabled_type is set to key . A maximum of 20 keys are supported.

Table 5-180 RedisDisabledCommandsDetail

Parameter	Type	Description
db_id	String	Database where a key is located.
key	String	Key name.
commands	Array of strings	Command list.

Status code: 400

Table 5-181 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Status code: 500

Table 5-182 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Example Requests

- Querying disabled commands for a redis instance. By default, the first 50 instances are queried.

```
GET https://{endpoint}/v3/0549b4a43100d4f32f51c01c2fe4acdb/redis/instances/0a9a2be86bdf493981d17626ad2f8668in12/disabled-commands?type=command
```
- Querying disabled commands for a Redis instance

```
GET https://{endpoint}/v3/0549b4a43100d4f32f51c01c2fe4acdb/redis/instances/0a9a2be86bdf493981d17626ad2f8668in12/disabled-commands?type=key&offset=0&limit=10
```

Example Responses

Status code: 200

Success.

```
{
  "total_count": 2,
  "disabled_type": "key",
  "keys": [ {
    "db_id": "0",
    "key": "name",
    "commands": [ "get" ]
  }, {
    "db_id": "0",
    "key": "address",
    "commands": [ "get", "set" ]
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.35 Deleting Disabled Commands for a Redis Instance

Function

This API is used to delete disabled commands for a Redis instance.

Constraints

This API supports only GeminiDB Redis instances.

URI

DELETE https://{Endpoint}/v3/{project_id}/redis/instances/{instance_id}/disabled-commands

Table 5-183 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-184 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-185 Request body parameters

Parameter	Mandatory	Type	Description
disabled_type	Yes	String	Disabled type. Values: <ul style="list-style-type: none"> • command • key
commands	No	Array of strings	This parameter is transferred when disabled_type is set to command . Values: <ul style="list-style-type: none"> • keys • hkeys • hvals • hgetall • smembers • flushdb • flushall

Parameter	Mandatory	Type	Description
keys	No	Array of objects in Table 5-186	This parameter is transferred when disabled_type is set to key . A maximum of 20 keys are allowed.

Table 5-186 RedisDisabledCommandsDetail

Parameter	Mandatory	Type	Description
db_id	Yes	String	Database where a key is located.
key	Yes	String	Key name.
commands	Yes	Array of strings	Command list.

Response Parameters

Status code: 400

Table 5-187 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Status code: 500

Table 5-188 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Example Requests

- Deleting a disabled command for a Redis instance
 DELETE `https://{endpoint}/v3/0549b4a43100d4f32f51c01c2fe4acdb/redis/instances/0a9a2be86bdf493981d17626ad2f8668in12/disabled-commands`

```
{
  "disabled_type": "command",
```

```
"commands" : [ "keys", "hkeys" ]
}
```

- **Deleting disabled commands for a Redis instance**

DELETE <https://{endpoint}/v3/0549b4a43100d4f32f51c01c2fe4acdb/redis/instances/0a9a2be86bdf493981d17626ad2f8668in12/disabled-commands>

```
{
  "disabled_type" : "key",
  "keys" : [ {
    "db_id" : "0",
    "key" : "name",
    "commands" : [ "get" ]
  }, {
    "db_id" : "0",
    "key" : "address",
    "commands" : [ "get", "set" ]
  } ]
}
```

Example Responses

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.36 Setting the Maintenance Period of an Instance

Function

This API is used to set the maintenance period of a specified instance.

Constraints

This API supports the following types of instances:

- GeminiDB Influx
- GeminiDB Redis
- GeminiDB Mongo
- GeminiDB Cassandra

URI

PUT https://{Endpoint}/v3/{project_id}/instances/{instance_id}/maintenance-window

Table 5-189 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID

Request Parameters

Table 5-190 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-191 Request body parameters

Parameter	Mandatory	Type	Description
start_time	Yes	String	Start time. The value must be a valid value in the "HH:MM" format. The current time is the UTC time and must be on the hour. By default, the interval between the end time and the start time is four hours.

Response Parameters

Status code: 204

None

Example Request

Set the maintenance window of an instance to 02:00-06:00.

```
PUT https://gaussdb-nosql.cn-north-7.myhuaweicloud.com/v3/0549b4a43100d4f32f51c01c2fe4acdb/instances/b0965c9010f44ffca9af4ee00746aa8din12/maintenance-window
```

```
{
  "start_time": "02:00"
}
```

Example Response

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.37 Performing a Primary/Standby Switchover

Function

This API is used to switch the primary and standby of a GeminiDB Redis instance.

Constraints

This API supports the following instances:

- GeminiDB Redis

URI

PUT https://{Endpoint}/v3/{project_id}/instance/{instance_id}/switchover

Table 5-192 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-193 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 202

Table 5-194 Response body parameters

Parameter	Type	Description
job_id	String	Task ID.

Status code: 400

Table 5-195 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Status code: 500

Table 5-196 Response body parameters

Parameter	Type	Description
error_code	String	Error code
error_msg	String	Error message

Example Requests

```
PUT https://{Endpoint}/v3/2900b7b8d03e4619b8db8d43bc6234ee/instance/f42ecf40870644b288884742bff6ef51in12/switchover
```

Example Responses

Status code: 202

Accepted.

```
{
  "job_id" : "528747ac-c30b-4904-a4ed-806c6943ec96"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.38 Starting or Stopping a Node

Function

This API is used to stop a node and enable other nodes to take over services after a node cannot work properly due to an underlying fault.

Constraints

This API supports the following instances:

- GeminiDB Redis

URI

PUT https://{Endpoint}/v3/{project_id}/instances/{instance_id}/nodes

Table 5-197 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-198 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-199 Request body parameters

Parameter	Mandatory	Type	Description
action	Yes	String	Operation type. shutdown indicates stopping a node and startup indicates starting a node.
node_ids	Yes	Array of strings	Node ID list. A maximum of 10 node IDs are supported.

Response Parameters

Status code: 202

Table 5-200 Response body parameters

Parameter	Type	Description
job_id	String	Task ID.

Status code: 400

Table 5-201 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Status code: 500

Table 5-202 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Example Requests

```
PUT https://{endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/
b31f6915c8be45e38958eac14e96a5d2in12/nodes
```

```
{
  "action": "shutdown",
  "node_ids": [ "d96eaf85e4b44b2aa800dcf2a5981711no12", "e76f07fb1dae420791af867fc2ac4c35no12" ]
}
```

Example Responses

Status code: 202

Success.

```
{
  "job_id": "e20eae83-1e77-4d58-8d92-c195e6372b0b"
}
```


Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.39 Querying Big Keys of a GeminiDB Redis Instance

Function

This API is used to query big keys of a GeminiDB Redis instance.

Constraints

This API supports the following instances:

- GeminiDB Redis

URI

POST https://{Endpoint}/v3/{project_id}/instances/{instance_id}/big-keys

Table 5-203 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-204 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-205 Request body parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Index position. If offset is set to N , the resource query starts from the $N+1$ piece of data. The value must be no lower than 0 . If this parameter is not transferred, offset is set to 0 by default, indicating that the query starts from the first big key.
limit	No	Integer	Maximum records to be queried. The value ranges from 1 to 100 . If this parameter is not transferred, the first 100 records are queried by default.
key_types	No	Array of strings	Big key type. It is a character string list and supports six types such as strings, hashes, zsets, sets, lists, and streams.

Response Parameters

Status code: 200

Table 5-206 Response body parameters

Parameter	Type	Description
keys	Array of Table 5-207 objects	Queried big keys.
count	Integer	Total big keys.

Table 5-207 BigKeysInfoResponseBody

Parameter	Type	Description
db_id	Integer	Database where a big key is located.
key_type	String	Type of a big key.
key_name	String	Name of a big key.
key_size	Integer	Length of a big key.

Status code: 400

Table 5-208 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Status code: 500

Table 5-209 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Example Requests

```
POST https://{endpoint}/v3/instances/91b5c48e11bf4913b35814f28b8f79f2in12/big-keys
{
  "key_types" : [ "string", "hash", "zset", "set", "list", "stream" ]
}
```

Example Responses

Status code: 200

Success.

```
{
  "keys" : [ {
    "db_id" : 0,
    "key_type" : "String",
    "key_name" : "a",
    "key_size" : 12
  } ],
  "count" : 1
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.40 Querying the Password-Free Configuration of a GeminiDB Redis Instance

Function

This API is used to query the password-free configuration of a GeminiDB Redis instance.

Constraints

This API supports the following instances:

- GeminiDB Redis

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/passwordless-config

Table 5-210 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Table 5-211 Query parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Index position. If offset is set to N , the resource query starts from the $N+1$ piece of data. The value must be no lower than 0 . If this parameter is not transferred, offset is set to 0 by default, indicating that the query starts from the first record.
limit	No	Integer	Maximum number of records to be queried. The value ranges from 1 to 100 . If this parameter is not transferred, the first 100 records are queried by default.

Request Parameters

Table 5-212 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-213 Response body parameters

Parameter	Type	Description
config_ips	Array of strings	Password-free configuration, which involves IP addresses and CIDR blocks. Only IPv4 addresses or CIDR blocks are supported.
total_count	Integer	Total number.

Status code: 400

Table 5-214 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Status code: 500

Table 5-215 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Example Requests

```
GET https://{endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/
b31f6915c8be45e38958eac14e96a5d2in12/passwordless-config
```

Example Responses

Status code: 200

Success.

```
{
  "config_ips" : [ "10.23.26.0/24" ],
  "total_count" : 1
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.41 Modifying the Password-Free Configuration of a GeminiDB Redis Instance

Function

This API is used to modify the password-free configuration of a GeminiDB Redis instance.

Constraints

This API supports the following instances:

- GeminiDB Redis

URI

PUT https://{Endpoint}/v3/{project_id}/instances/{instance_id}/passwordless-config

Table 5-216 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-217 Request header parameter

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-218 Request body parameters

Parameter	Mandatory	Type	Description
config_ips	Yes	Array of strings	List of IP addresses and CIDR blocks to be configured. If the list is empty, the password-free configuration is cleared.

Response Parameters

Status code: 400

Table 5-219 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Status code: 500

Table 5-220 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Example Requests

```
PUT https://{endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/
b31f6915c8be45e38958eac14e96a5d2in12/passwordless-config
```

```
{
  "config_ips": [ "10.23.26.0/24" ]
}
```

Example Responses

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.42 Querying the Memory Mapping List and Details

Function

This API is used to query the memory mapping list and details based on specified conditions.

Constraints

This API can be used only on primary/standby GeminiDB Redis instances.

URI

GET https://{Endpoint}/v3/{project_id}/dbcache/mappings

Table 5-221 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Table 5-222 Query parameters

Parameter	Mandatory	Type	Description
id	No	String	Mapping ID, which can be obtained by calling the API for querying the memory mapping list and details.
name	No	String	Mapping name. If you enter a name starting with an asterisk (*), fuzzy search results are returned. If you enter a valid name, an exact result is returned.

Parameter	Mandatory	Type	Description
source_instance_id	No	String	Source instance ID
source_instance_name	No	String	Source instance name. If you enter a name starting with an asterisk (*), fuzzy search results are returned. If you enter a valid name, an exact result is returned.
target_instance_id	No	String	Target instance ID
target_instance_name	No	String	Target instance name. If you enter a name starting with an asterisk (*), fuzzy search results are returned. If you enter a valid name, an exact result is returned.
offset	No	Integer	Index offset. If offset is set to N , the resource query starts from the N+1 piece of data. If action is set to filter , offset is 0 by default, the query starts from the first record. The offset value must be a positive integer.
limit	No	Integer	Maximum records to be queried. The value ranges from 1 to 100 . If this parameter is not transferred, the first 100 records are queried by default.

Request Parameters

Table 5-223 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

Status code: 200

Table 5-224 Response body parameters

Parameter	Type	Description
total_count	Integer	Total number of records
dbcache_mappings	Array of objects in Table 5-225	Memory mapping information

Table 5-225 QueryDBCacheMappingResponse

Parameter	Type	Description
id	String	Memory mapping ID
name	String	Memory mapping name
source_instance_id	String	Source instance ID
source_instance_name	String	Source instance name
target_instance_id	String	Target instance ID
target_instance_name	String	Target instance name
status	String	Memory mapping relationship <ul style="list-style-type: none"> ● normal: A memory mapping is normal. ● creating: A memory mapping is being created. ● createfail: A memory mapping failed to be created. ● deleting: A memory mapping is being deleted. ● stopped: A memory mapping is stopped. ● deleted: A memory mapping is deleted.
created	String	Time when a memory mapping was created
updated	String	Last time when a memory mapping was changed
rule_count	Integer	Number of rules in a memory mapping

Example Request

Querying the first 10 memory mappings

GET https://{endpoint}/v3/054e292c9880d4992f02c0196d3ea468/dbcache/mappings?offset=0&limit=10

Example Response

Status code: 200

Success.

```
{
  "total_count": 1,
  "dbcache_mappings": [ {
    "id": "0231e002-d278-4df8-815c-8284164168e5",
    "name": "dbcache_rule_01",
    "source_instance_id": "d8e6ca5a624745bcb546a227aa3ae1cfin01",
    "target_instance_id": "9136fd2a9fcd405ea4674276ce36dae8in12",
    "source_instance_name": "source_name",
    "target_instance_name": "target_name",
    "status": "normal",
    "rule_count": 10,
    "created": "2024-01-17T07:05:52",
    "updated": "2024-01-20T08:07:02"
  } ]
}
```

Status Codes

See [Status Codes](#).

Error Codes

See [Error Codes](#).

5.3.43 Creating a Memory Acceleration Rule

Function

This API is used to create a memory acceleration rule.

Constraints

This API can be used only on primary/standby GeminiDB Redis instances.

URI

POST https://{Endpoint}/v3/{project_id}/dbcache/rule

Table 5-226 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-227 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-228 Request body parameters

Parameter	Mandatory	Type	Description
dbcache_mapping_id	Yes	String	Memory mapping ID
name	Yes	String	Name of a memory acceleration rule, which contains no more than 256 characters and must be unique in the current mapping
source_db_schema	Yes	String	Source database
source_db_table	Yes	String	Source database table
storage_type	Yes	String	Storage type of the destination database. The value is hash .
target_database	Yes	String	Destination database. The value ranges from 0 to 999 .
key_columns	Yes	Array of strings	Columns used by mapped keys
value_columns	Yes	Array of strings	Columns used by mapped values
ttl	No	String	Lifetime of a key (unit: ms). If this parameter is not transferred, the default value 2592000000 is used, indicating 30 days.
key_separator	Yes	String	Key separator of a mapping. Only one character is allowed.
value_separator	No	String	Value separator of a mapping. Only one character is allowed.

Parameter	Mandatory	Type	Description
key_prefix	Yes	String	Key prefix. The value can contain no more than 1024 characters. key_prefix and target_database must be unique in the current mapping.

Response Parameters

Status code: 200

Table 5-229 Response body parameters

Parameter	Type	Description
id	String	ID of a memory acceleration rule

Example Request

Creating a memory acceleration rule

```
POST https://{endpoint}/v3/054e292c9880d4992f02c0196d3ea468/dbcache/rule
{
  "dbcache_mapping_id": "e6f8526c-8b9c-4811-9218-4d2d7c1ddde1",
  "name": "mysql_a1-to-redis_01",
  "source_db_schema": "test",
  "source_db_table": "students_info_t",
  "storage_type": "hash",
  "target_database": "0",
  "key_columns": [ "sno", "sname" ],
  "value_columns": [ "sno", "sname", "age" ],
  "ttl": "-1",
  "key_separator": ":",
  "key_prefix": "prefix1"
}
```

Example Response

Status code: 200

Success.

```
{
  "id": "04efe8e2-9255-44ae-a98b-d87cae411890"
}
```

Status Codes

See [Status Codes](#).

Error Codes

See [Status Codes](#).

5.3.44 Deleting a Memory Mapping

Function

This API is used to delete a memory mapping.

Constraints

This API can be used only on primary/standby GeminiDB Redis instances.

URI

DELETE https://{Endpoint}/v3/{project_id}/dbcache/mapping

Table 5-230 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-231 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-232 Request body parameters

Parameter	Mandatory	Type	Description
id	Yes	String	Memory mapping ID

Response Parameters

Status code: 202

Table 5-233 Response body parameters

Parameter	Type	Description
job_id	String	Task ID

Example Request

```
DELETE https://{endpoint}/v3/054e292c9880d4992f02c0196d3ea468/dbcache/mapping
{
  "id": "1094e5d3-c140-4095-9b13-242de267c175"
}
```

Example Response

Status code: 202

Accepted.

```
{
  "job_id": "04efe8e2-9255-44ae-a98b-d87cae411890"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.3.45 Creating a Memory Mapping

Function

This API is used to create a memory mapping.

Constraints

This API can be used only on primary/standby GeminiDB Redis instances.

URI

POST https://{Endpoint}/v3/{project_id}/dbcache/mapping

Table 5-234 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-235 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-236 Request body parameters

Parameter	Mandatory	Type	Description
source_instance_id	Yes	String	ID of a source instance for which a memory mapping is created. Currently, RDS for MySQL and GaussDB(for MySQL) instances are supported.
target_instance_id	Yes	String	ID of a destination instance for which a memory mapping is created. Currently, only primary/standby GeminiDB Redis instances are supported.

Response Parameters

Status code: 202

Table 5-237 Response body parameters

Parameter	Type	Description
job_id	String	ID of a memory mapping task to be created

Example Request

```
POST https://{endpoint}/v3/054e292c9880d4992f02c0196d3ea468/dbcache/mapping
{
  "source_instance_id": "356cf333ab77413f9d2526605227be08in01",
  "target_instance_id": "2bb234c7175c400aab91de7c62174c0ein12"
}
```

Example Response

Status code: 202

Accepted.


```
{
  "job_id" : "04efe8e2-9255-44ae-a98b-d87cae411890"
}
```

Status Codes

See [Status Codes](#).

Error Codes

See [Status Codes](#).

5.3.46 Modifying a Memory Acceleration Rule

Function

This API is used to modify a memory acceleration rule.

Constraints

This API can be used only on primary/standby GeminiDB Redis instances.

URI

PUT https://{Endpoint}/v3/{project_id}/dbcache/rule

Table 5-238 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-239 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-240 Request body parameters

Parameter	Mandatory	Type	Description
dbcache_rule_id	Yes	String	ID of a memory acceleration rule

Parameter	Mandatory	Type	Description
value_columns	Yes	Array of strings	Columns used by mapped values
ttl	No	String	Lifetime of a key (unit: ms). If this parameter is not transferred, the default value 2592000000 is used, indicating 30 days.
value_separator	No	String	Value separator of a mapping. Only one character is allowed.

Response Parameters

None

Example Request

Updating a memory acceleration rule

```
PUT https://{endpoint}/v3/054e292c9880d4992f02c0196d3ea468/dbcache/rule
{
  "dbcache_rule_id": "0231e002-d278-4df8-815c-8284164168e5",
  "value_columns": [ "sno", "sname", "age" ],
  "ttl": 3000,
  "value_separator": "_"
}
```

Example Response

None

Status Codes

See [Status Codes](#).

Error Codes

See [Status Codes](#).

5.3.47 Querying Memory Mapping Rules and Details

Function

This API is used to query memory acceleration rules and details.

Constraints

This API can be used only on primary/standby GeminiDB Redis instances.

URI

GET https://{Endpoint}/v3/{project_id}/dbcache/rules

Table 5-241 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Table 5-242 Query parameters

Parameter	Mandatory	Type	Description
dbcache_mapping_id	Yes	String	Memory mapping ID
rule_id	No	String	ID of a memory acceleration rule
rule_name	No	String	Name of a memory acceleration rule. If you enter a name starting with an asterisk (*), fuzzy search results are returned. If you enter a valid name, an exact result is returned.
source_db_schema	No	String	Source database name. If you enter a name starting with an asterisk (*), fuzzy search results are returned. If you enter a valid name, an exact result is returned.
source_db_table	No	String	Source database table name. If you enter a name starting with an asterisk (*), fuzzy search results are returned. If you enter a valid name, an exact result is returned.
offset	No	String	Index offset. If offset is set to N , the resource query starts from the N+1 piece of data. If action is set to filter , offset is 0 by default, the query starts from the first record. The offset value must be a positive integer.

Parameter	Mandatory	Type	Description
limit	No	String	Maximum records to be queried. The value ranges from 1 to 100 . If this parameter is not transferred, the first 100 records are queried by default.

Request Parameters

Table 5-243 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

Status code: 200

Table 5-244 Response body parameters

Parameter	Type	Description
total_count	Integer	Total number of records
dbcache_mapping_id	String	Memory mapping ID
rules	Array of objects in Table 5-245	Memory acceleration rule details

Table 5-245 QueryDBCacheRuleResponse

Parameter	Type	Description
id	String	ID of a memory acceleration rule
name	String	Name of a memory acceleration rule
status	String	Status of a memory acceleration rule <ul style="list-style-type: none"> • normal: A memory acceleration rule is normal. • createfail: A memory acceleration rule failed to be created.

Parameter	Type	Description
source_db_schema	String	Source database
source_db_table	String	Source database table
storage_type	String	Storage type of the destination database. The value is hash .
target_database	String	Destination database
key_columns	Array of strings	Columns used by mapped keys
value_columns	Array of strings	Columns used by mapped values
ttl	String	Lifetime of a key (unit: ms). If this parameter is not transferred, default value -1 is used, indicating permanent storage.
key_separator	String	Key separator of a mapping.
value_separator	String	Value separator of a mapping.
key_prefix	String	Key prefix.

Example Request

Querying the first 10 memory acceleration rules whose mapping ID is **1094e5d3-c140-4095-9b13-242de267c175**

```
GET https://{endpoint}/v3/054e292c9880d4992f02c0196d3ea468/dbcache/rule?offset=0&limit=10
```

```
{
  "dbcache_mapping_id" : "1094e5d3-c140-4095-9b13-242de267c175"
}
```

Example Response

Status code: 200

Success.

```
{
  "total_count" : 1,
  "dbcache_mapping_id" : "0231e002-d278-4df8-815c-8284164168e5",
  "rules" : [ {
    "id" : "0231e002-d278-4df8-815c-8284164168e5",
    "name" : "dbcache_rule_01",
    "source_db_schema" : "test",
    "source_db_table" : "student_t",
    "status" : "normal",
    "storage_type" : "hash",
    "target_database" : "0",
  }
]
```

```

"key_columns" : [ "sno", "sname", "age" ],
"value_columns" : [ "sno", "sname", "age" ],
"ttl" : "10000",
"key_separator" : ":"
} ]
}

```

Status Codes

See [Status Codes](#).

Error Codes

See [Status Codes](#).

5.3.48 Deleting a Memory Acceleration Rule

Function

This API is used to delete a memory acceleration rule.

Constraints

This API can be used only on primary/standby GeminiDB Redis instances.

URI

DELETE https://{Endpoint}/v3/{project_id}/dbcache/rule

Table 5-246 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-247 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-248 Request body parameters

Parameter	Mandatory	Type	Description
id	Yes	String	ID of a memory acceleration rule

Response Parameters

None

Example Request

DELETE https://{endpoint}/v3/054e292c9880d4992f02c0196d3ea468/dbcache/rule

Example Response

None

Status Codes

See [Status Codes](#).

Error Codes

See [Status Codes](#).

5.3.49 Enabling or Disabling Instance Data Export

Function

This API is used to enable or disable instance data export.

Currently, GeminiDB Influx instance data can be converted into a .parquet file and then uploaded to a specified OBS bucket.

Constraints

This API can be used only on performance-enhanced GeminiDB Influx instances.

URI

PUT https://{Endpoint}/v3/{project_id}/instances/{instance_id}/data-dump

Table 5-249 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Parameter	Mandatory	Type	Description
instance_id	Yes	String	Instance ID

Request Parameters

Table 5-250 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-251 Request body parameters

Parameter	Mandatory	Type	Description
bucket_name	Yes	String	OBS bucket name
action	Yes	String	Enables or disables instance data export.

Response Parameters

Status code: 400

Table 5-252 Response body parameters

Parameter	Type	Description
error_code	String	Error code
error_msg	String	Error message

Example Request

Enabling or disabling instance data export

```
PUT https://{endpoint}/v3/0549b4a43100d4f32f51c01c2fe4acdb/instances/0a9a2be86bdf493981d17626ad2f8668in13/data-dump
```

```
{
  "bucket_name": "influxdb-obs",
  "action": "open"
}
```

Example Response

None

Status Codes

See [Status Codes](#).

Error Codes

See [Status Codes](#).

5.3.50 Enabling or Disabling Second-Level Monitoring

Function

This API is used to enable or disable monitoring by 5s for an instance.

Constraints

This API can be used to only on GeminiDB Redis instances with four or more CPUs in a proxy cluster.

URI

PUT `https://{Endpoint}/v3/{project_id}/instances/{instance_id}/monitoring-by-seconds/switch`

Table 5-253 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID

Request Parameters

Table 5-254 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-255 Request body parameters

Parameter	Mandatory	Type	Description
enabled	Yes	Boolean	Whether second-level monitoring is enabled <ul style="list-style-type: none"> • true: Enabled • false: Disabled

Response Parameters

None

Example Request

- URI example
PUT https://{endpoint}/v3/054e292c9880d4992f02c0196d3ea468/instances/054e292c9880d4992f02c0196d3ein12/monitoring-by-seconds/switch
- Example request
Enabling second-level monitoring

```
{
  "enabled" : true
}
```

Example Response

None

Status Codes

See [Status Codes](#).

Error Codes

See [Status Codes](#).

5.3.51 Querying Configurations of Second-Level Monitoring

Function

This API is used to query configurations of second-level monitoring.

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/monitoring-by-seconds/switch

Table 5-256 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID

Request Parameters

Table 5-257 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

Status code: 200

Table 5-258 Response body parameters

Parameter	Type	Description
enabled	Boolean	Whether second-level monitoring is enabled <ul style="list-style-type: none"> true: Enabled false: Disabled

Example Request

URI example

```
GET https://{endpoint}/v3/054e292c9880d4992f02c0196d3ea468/instances/054e292c9880d4992f02c0196d3ein12/monitoring-by-seconds/switch
```

Example Response

Status code: 200

Success.

```
{
  "enabled" : false
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.4 Connection Management

5.4.1 Obtaining Instance Sessions

Function

This API is used to obtain instance sessions.

Constraints

This API can be used only for the following instance type:

- GeminiDB Redis

URI

GET `https://{Endpoint}/v3/{project_id}/instances/{instance_id}/sessions`

Table 5-259 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID

Table 5-260 Query parameters

Parameter	Mandatory	Type	Description
node_id	No	String	Node ID

Request Parameters

Table 5-261 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

Status code: 200

Table 5-262 Response body parameters

Parameter	Type	Description
node_sessions	Array of objects in Table 5-263	Node sessions

Table 5-263 ListNodeSessionsResult

Parameter	Type	Description
node_id	String	Node ID
total_count	Integer	Total number of sessions
sessions	Array of objects in Table 5-264	Node sessions

Table 5-264 sessions

Parameter	Type	Description
id	String	Session ID
name	String	Connection name
cmd	String	Last executed command
age	String	Connection duration, in seconds
idle	String	Idle duration, in seconds
db	String	ID of a database that is being used by a client
addr	String	IP address and port number of a client
fd	String	File descriptor for sockets
sub	String	Number of subscribed channels
psub	String	Number of subscribed modes
multi	String	Number of commands executed in a transaction

Status code: 400

Table 5-265 Response body parameters

Parameter	Type	Description
error_code	String	Error code
error_msg	String	Error message

Status code: 500

Table 5-266 Response body parameters

Parameter	Type	Description
error_code	String	Error code
error_msg	String	Error message

Example Request

- Querying sessions of all nodes on an instance

```
GET https://{endpoint}/v3/054e292c9880d4992f02c0196d3ea468/instances/e73893ef73754465a8bd2e0857bbf13ein12/sessions
```
- Querying sessions of a specified node on an instance

```
GET https://{endpoint}/v3/054e292c9880d4992f02c0196d3ea468/instances/e73893ef73754465a8bd2e0857bbf13ein12/sessions?node_id=9d3288bbbb81448e8a99ff41a5717281no12
```

Example Response

Status code: 200

Success.

```
{
  "node_sessions": [ {
    "node_id": "8436a91546294036b75931e879882200in06",
    "total_count": 1,
    "sessions": [ {
      "addr": "127.0.0.1:8080",
      "id": 54487,
      "name": "cli",
      "cmd": "get",
      "age": 8888581,
      "idle": 8888581,
      "db": 0,
      "fd": 1311,
      "sub": 0,
      "psub": 0,
      "multi": -1
    } ]
  }, {
    "node_id": "8123sa46294daw13def423fw79882200in06",
    "total_count": 1,
    "sessions": [ {
      "addr": "127.0.0.1:8080",
      "id": 54488,
      "name": "cli",
```

```

"cmd" : "get",
"age" : 3124124,
"idle" : 4414141,
"db" : 0,
"fd" : 1312,
"sub" : 0,
"psub" : 0,
"multi" : -1
}]
}]
}

```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.4.2 Closing Sessions of All Nodes on an Instance

Function

This API is used to close sessions of all nodes on an instance.

Constraints

This API can be used only for the following instance type:

- GeminiDB Redis (proxy cluster)

URI

DELETE https://{Endpoint}/v3/{project_id}/instances/{instance_id}/sessions

Table 5-267 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID

Request Parameters

Table 5-268 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

None

Example Request

```
DELETE https://{endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/  
c865f921f3dd45198f209a607533a779in12/sessions
```

Example Response

Status code: 200

Success

```
{  
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5 Backups and Restorations

5.5.1 Querying Backups

Function

This API is used to query backups based on specified conditions.

GeminiDB Cassandra only allows you to view incremental backups and differential backups and their sizes.

URI

GET https://{Endpoint}/v3/{project_id}/backups

Table 5-269 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Table 5-270 Query parameters

Parameter	Mandatory	Type	Description
offset	Yes	Integer	<p>Explanation: Page number.</p> <p>Constraints: N/A</p> <p>Values: This value is greater than or equal to 1.</p> <p>Default value: N/A</p>
limit	Yes	Integer	<p>Explanation: Number of records on each page.</p> <p>Constraints: N/A</p> <p>Values: 1 to 100</p> <p>Default value: N/A</p>

Parameter	Mandatory	Type	Description
datastore_type	No	String	<p>Explanation: API type</p> <p>Values:</p> <ul style="list-style-type: none"> ● cassandra: GeminiDB Cassandra API ● mongodb: GeminiDB Mongo API ● redis: GeminiDB Redis API ● influxdb: GeminiDB Influx API <p>Default value: If this parameter is not transferred, all APIs are queried.</p>
type	No	String	<p>Explanation: Backup policy type</p> <p>Values:</p> <ul style="list-style-type: none"> ● Instance, indicating that an instance backup is queried. ● DatabaseTable: A table-level backup is queried. This parameter is available only for GeminiDB Cassandra API. <p>Default value: Instance</p>
instance_id	No	String	<p>Explanation: Instance ID</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: If this parameter is not transferred, all backups are queried.</p>

Parameter	Mandatory	Type	Description
backup_id	No	String	<p>Explanation: Backup ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: If this parameter is not transferred, all backups are queried.</p>
backup_type	No	String	<p>Explanation: Backup type.</p> <p>Constraints: N/A</p> <p>Values:</p> <ul style="list-style-type: none"> • Auto: indicates that the backup is an automated full backup. • Manual: indicates that the backup is a manual full backup. • Incremental: indicates that the backup is an incremental backup. • Differential: indicates that the backup is a differential backup. <p>Default value: N/A</p>
begin_time	No	String	<p>Explanation: Time when a backup query starts</p> <p>Constraints: The value is in UTC format.</p> <p>Values: The format is <i>yyyy-mm-dd hh:mm:ss</i>.</p> <p>Default value: N/A</p>

Parameter	Mandatory	Type	Description
end_time	No	String	<p>Explanation: Time when a backup query ends</p> <p>Constraints: The value is in UTC format.</p> <p>Values: The format is <i>yyyy-mm-dd hh:mm:ss</i>.</p> <p>Default value: N/A</p>

Request Parameters

Table 5-271 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Response Parameters

Status code: 200

Table 5-272 Response body parameters

Parameter	Type	Description
total_count	Long	Explanation: Total number of records. Values: N/A
backups	Array of objects	Explanation: Backup list. For details, see Table 5-273 .

Table 5-273 backups

Parameter	Type	Description
id	String	Explanation: Backup ID. Values: N/A
name	String	Explanation: Backup name. Values: N/A
description	String	Explanation: Backup description. Values: N/A
begin_time	String	Explanation: Backup start time Constraints: UTC time in the <i>yyyy-mm-dd hh:mm:ss</i> format
end_time	String	Explanation: Backup end time Constraints: UTC time in the <i>yyyy-mm-dd hh:mm:ss</i> format

Parameter	Type	Description
status	String	<p>Explanation: Backup status.</p> <p>Values:</p> <ul style="list-style-type: none"> ● BUILDING: indicates that the backup is in progress. ● COMPLETED: indicates that the backup is completed. ● FAILED: indicates that the backup failed.
size	Double	<p>Explanation: Backup size, in KB.</p> <p>Values: N/A</p>
type	String	<p>Explanation: Backup type</p> <p>Values:</p> <ul style="list-style-type: none"> ● Auto: automated full backup ● Manual: manual full backup ● Incremental: incremental backup ● Differential: differential backup
instance_id	String	<p>Explanation: Instance ID</p> <p>Values: N/A</p>
instance_name	String	<p>Explanation: Instance name.</p> <p>Values: N/A</p>
datastore	object	<p>Explanation: Database information. For details, see Table 5-274.</p>

Table 5-274 datastore

Parameter	Type	Description
type	String	Explanation: Database type Values: N/A
version	String	Explanation: Database version Values: N/A

Example Requests

- **URI example**
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/054b61972980d4552f0bc00ac8d3f5cd/backups?instance_id=c0c9f155c7b7423a9d30f0175998b63bin01&offset=2&limit=2&begin_time=2018-07-06 10:41:14&end_time=2018-08-16 10:41:14&type=DatabaseTable
- **Incremental backups and their sizes**
GET https://{Endpoint}/v3/2900b7b8d03e4619b8db8d43bc6234ee/backups?offset=1&limit=5&backup_type=Incremental&instance_id=3149aee486d748f68db1ee81e95b9f56in06
- **Differential backups and their sizes**
GET https://{Endpoint}/v3/2900b7b8d03e4619b8db8d43bc6234ee/backups?offset=1&limit=5&backup_type=Differential&instance_id=3149aee486d748f68db1ee81e95b9f56in06
- **Example request body**
None

Example Responses

Status code: 200

Success

```
{
  "total_count" : 4,
  "backups" : [ {
    "id" : "43e4feaab48f11e89039fa163ebaa7e4br01",
    "name" : "backup-test",
    "instance_id" : "43e4feaab48f11e89039fa163ebaa7e4br01",
    "instance_name" : "cluster-test",
    "datastore" : {
      "type" : "cassandra",
      "version" : "3.4"
    },
    "type" : "Auto",
    "begin_time" : "2018-08-06 12:41:14",
    "end_time" : "2018-08-06 12:43:14",
    "status" : "COMPLETED",
    "size" : 2803,
    "description" : "backup description",
  }, {
    "id" : "43e4feaab48f11e89039fa163ebaa7e4br02",
    "name" : "backup-test-2",
    "instance_id" : "43e4feaab48f11e89039fa163ebaa7e4br02",
```

```

"instance_name" : "cluster-test",
"datastore" : {
  "type" : "cassandra",
  "version" : "3.4"
},
"type" : "Manual",
"begin_time" : "2018-08-06 12:41:14",
"end_time" : "2018-08-06 12:43:14",
"status" : "COMPLETED",
"size" : 2803,
"description" : "backup description",
}]
}

```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.2 Querying Backups (Recommended)

Function

This API is used to query backups based on specified conditions.

URI

GET https://{Endpoint}/v3.1/{project_id}/backups

Table 5-275 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Table 5-276 Query parameters

Parameter	Mandatory	Type	Description
instance_id	No	String	<p>Explanation: Instance ID</p> <p>Constraints: You can call the API used for querying instances to obtain the value. If there are no instances available, create one by calling the API used for creating an instance.</p> <p>Values: N/A</p> <p>Default value: N/A</p>
datastore_type	No	String	<p>Explanation: Database type</p> <p>Constraints: N/A</p> <p>Values:</p> <ul style="list-style-type: none"> • cassandra • redis • mongodb • influxdb <p>Default value: N/A</p>
backup_id	No	String	<p>Explanation: Backup ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Parameter	Mandatory	Type	Description
backup_type	No	String	<p>Explanation: Backup type</p> <p>Constraints: The value is case sensitive.</p> <p>Values:</p> <ul style="list-style-type: none"> • Auto, indicating that this backup is an automatic full backup. • Manual, indicating that this backup is a manual full backup. <p>Default value: If this parameter is not transferred, all automated and manual full backups (including database- and table-level backups) are queried by default.</p>
type	No	String	<p>Explanation: Backup policy type</p> <p>Constraints: N/A</p> <p>Values:</p> <ul style="list-style-type: none"> • Instance, indicating that an instance backup is queried. • DatabaseTable, indicating that a database or table backup is queried. <p>Default value: Instance</p>
limit	No	Integer	<p>Explanation: Maximum backup records to be queried.</p> <p>Constraints: N/A</p> <p>Values: 1 to 100</p> <p>Default value: If this parameter is not transferred, the first 100 instances are queried by default.</p>

Parameter	Mandatory	Type	Description
offset	No	Integer	<p>Explanation: Index position. The query starts from the next backup creation time indexed by this parameter under a specified project. If offset is set to <i>N</i>, the resource query starts from the <i>N</i>+1 piece of data.</p> <p>Constraints: N/A</p> <p>Values: The value must be greater than or equal to 0.</p> <p>Default value: If this parameter is not transferred, the index offset is 0 by default, indicating that the query starts from the latest created instance.</p>
begin_time	No	String	<p>Explanation: Time when a backup query starts</p> <p>Constraints: When end_time is not empty, begin_time is mandatory.</p> <p>Values: The value is in the format of <i>yyyy-mm-ddThh:mm:ssZ</i>. <i>T</i> indicates start time and <i>Z</i> a time zone offset.</p> <p>Default value: Null</p>

Parameter	Mandatory	Type	Description
end_time	No	String	<p>Explanation: Time when a backup query ends</p> <p>Constraints: When begin_time is not empty, end_time is mandatory.</p> <p>Values: The value is in the format of <i>yyyy-mm-ddThh:mm:ssZ</i>. <i>T</i> indicates start time and <i>Z</i> a time zone offset.</p> <p>Default value: Null</p>

Request Parameters

Table 5-277 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Response Parameters

Status code: 200

Table 5-278 Response body parameters

Parameter	Type	Description
total_count	Integer	Explanation: Total backups. Values: N/A
backups	Array of objects	Explanation: Backup details. For details, see Table 5-279 .

Table 5-279 Backup

Parameter	Type	Description
id	String	Explanation: Backup ID. Values: N/A
description	String	Explanation: Backup description. Values: N/A
instance_id	String	Explanation: ID of the instance that the backup is created for. Values: N/A
instance_name	String	Explanation: Name of the instance that the backup is created for. Values: N/A
datastore	object	Explanation: DB version information. For details, see Table 5-280 .
name	String	Explanation: Backup name. Values: N/A

Parameter	Type	Description
type	String	<p>Explanation: Backup type</p> <p>Values:</p> <ul style="list-style-type: none"> • Auto: indicates that this backup is an automatic full backup. • Manual: indicates that this backup is a manual full backup.
size	Double	<p>Explanation: Backup size, in KB.</p> <p>Values: N/A</p>
status	String	<p>Explanation: Backup status</p> <p>Values:</p> <ul style="list-style-type: none"> • BUILDING, indicating that the backup is in progress. • COMPLETED, indicating that the backup is completed. • FAILED, indicating that the backup failed.
begin_time	String	<p>Explanation: Backup start time</p> <p>Constraints: The value is in the format of <i>yyyy-mm-ddThh:mm:ssZ</i>. <i>T</i> indicates start time and <i>Z</i> a time zone offset.</p>
end_time	String	<p>Explanation: Backup end time</p> <p>Constraints: The value is in the format of <i>yyyy-mm-ddThh:mm:ssZ</i>. <i>T</i> indicates start time and <i>Z</i> a time zone offset.</p>
database_tables	Array of objects	<p>Explanation: Database and table information in the backup. For details, see Table 5-281.</p> <p>Constraints:</p> <ul style="list-style-type: none"> • Keep this parameter empty or ignore it when you query an instance backup. • Specify this parameter when you query a database or table backup (if any).

Table 5-280 Datastore

Parameter	Type	Description
type	String	Explanation: Database type Values: N/A
version	String	Explanation: Database version Values: N/A

Table 5-281 QueryDatabaseTableInfo

Parameter	Type	Description
database_name	String	Explanation: Database name. Values: N/A
table_names	Array of strings	Explanation: Table names. <ul style="list-style-type: none"> • If this parameter is empty, database names are queried. • If this parameter is not empty, table names are queried.

Example Requests

- URI example
GET `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3.1/054b61972980d4552f0bc00ac8d3f5cd/backups?instance_id=c0c9f155c7b7423a9d30f0175998b63bin01&offset=2&limit=2&begin_time=2019-05-27T03:38:51+0000&end_time=2019-05-28T03:38:51+0000&type=DatabaseTable`
- Example request body
None

Example Responses

Status code: 200

Success

```
{
  "total_count": 4,
  "backups": [ {
```

```
{
  "id": "43e4feaab48f11e89039fa163ebaa7e4br01",
  "name": "backup-test",
  "instance_id": "43e4feaab48f11e89039fa163ebaa7e4br01",
  "instance_name": "cluster-test",
  "datastore": {
    "type": "cassandra",
    "version": "3.4"
  },
  "type": "Auto",
  "begin_time": "2019-05-27T03:38:51+0000",
  "end_time": "2019-05-27T03:39:51+0000",
  "status": "COMPLETED",
  "size": 2803,
  "description": "backup description",
  "database_tables": [ {
    "database_name": "DATABASE_X",
    "table_names": [ "TABLE_A", "TABLE_B", "TABLE_C" ]
  }, {
    "database_name": "DATABASE_Y",
    "table_names": null
  } ]
}, {
  "id": "43e4feaab48f11e89039fa163ebaa7e4br02",
  "name": "backup-test-2",
  "instance_id": "43e4feaab48f11e89039fa163ebaa7e4br02",
  "instance_name": "cluster-test",
  "datastore": {
    "type": "cassandra",
    "version": "3.4"
  },
  "type": "Manual",
  "begin_time": "2019-05-27T03:38:51+0000",
  "end_time": "2019-05-27T03:39:51+0000",
  "status": "COMPLETED",
  "size": 2803,
  "description": "backup description",
  "database_tables": [ {
    "database_name": "DATABASE_X",
    "table_names": [ "TABLE_A", "TABLE_B", "TABLE_C" ]
  }, {
    "database_name": "DATABASE_Y",
    "table_names": null
  } ]
} ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.3 Querying an Automated Backup Policy

Function

This API is used to query an automated backup policy, including GeminiDB Cassandra databases and tables.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

URI

GET https://{Endpoint}/v3.1/{project_id}/instances/{instance_id}/backups/policy

Table 5-282 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Table 5-283 Query parameters

Parameter	Mandatory	Type	Description
type	No	String	Backup policy type. This parameter is available only to GeminiDB Cassandra. The value can be: <ul style="list-style-type: none"> • Instance: An instance backup is queried. • DatabaseTable: A database or table backup is queried. • The default value is Instance.

Request Parameters

Table 5-284 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

Status code: 200

Table 5-285 Response body parameters

Parameter	Type	Description
backup_policy	object	Backup policy objects, including the backup retention period (days) and start time. For details, see Table 5-286 .

Table 5-286 ShowBackupPolicyResult

Parameter	Type	Description
keep_days	Integer	Backup retention days.
start_time	String	Backup time window. Automated backup will be triggered during the backup time window.
period	String	Backup period. After a backup period is specified, data will be automatically backed up on the selected days every week.

Example Request

URI example

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3.1/054b61972980d4552f0bc00ac8d3f5cd/instances/764b8a2763d34414ad3f2d5495416cb5in06/backups/policy?type=DatabaseTable
```

Example Response

Status code: 200

Success

Enabling automated backup

```
{
  "backup_policy": {
    "keep_days": 7,
    "start_time": "19:00-20:00",
    "period": "1,2,4,5,6"
  }
}
```

Disabling automated backup

```
{
  "backup_policy": {
    "keep_days": 0
  }
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.4 Configuring an Automated Backup Policy

Function

This API is used to configuring an automated backup policy.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Influx
- GeminiDB Redis

URI

PUT `https://{Endpoint}/v3/{project_id}/instances/{instance_id}/backups/policy`

Table 5-287 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-288 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-289 Request body parameters

Parameter	Mandatory	Type	Description
backup_policy	Yes	object	Backup policy objects, including backup retention period (days) and start time. For details, see Table 5-290 .

Table 5-290 BackupPolicy

Parameter	Mandatory	Type	Description
keep_days	Yes	Integer	Backup retention days. The value ranges from 0 to 35 . The value 0 indicates that the automated backup policy is disabled.
start_time	No	String	Backup time window. Automated backup will be triggered during the backup time window. This parameter is mandatory if the automated backup policy is enabled. If the policy is disabled, you do not need to transfer this parameter. The value must be the UTC time in the hh:mm-HH:MM format. <ul style="list-style-type: none"> The HH value must be 1 greater than the hh value. The values of mm and MM must be the same and must be set to 00, 15, 30, or 45. Example value: 23:00-00:00

Parameter	Mandatory	Type	Description
period	No	String	<p>Backup period. After a backup period is specified, data will be automatically backed up on the selected days every week. This parameter is mandatory if the automated backup policy is enabled. If the policy is disabled, you do not need to transfer this parameter.</p> <p>The value is a list of digits separated by commas (.). Each digit indicates a day of the week. The restrictions on the backup period are as follows:</p> <ul style="list-style-type: none"> • If you set keep_days to 0, this parameter is not transferred. • If you set keep_days to 1 to 6, set this parameter to 1, 2, 3, 4, 5, 6, 7. • If you set keep_days to 7 to 35, select at least one day of the week for the backup cycle. Example value: 1,2,3,4

Response Parameters

Status code: 204

No response parameters

Example Requests

- URI example
PUT <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/9136fd2a9fcd405ea4674276ce36dae8in02/backups/policy>
- Example request body
Enabling or modifying the automated backup policy (Set **period** to **1, 2, 3, 4, 5**, and **6**, **start_time** to **01:00-02:00**, and **keep_days** to **7**.)

```
{
  "backup_policy": {
    "period": "1,2,3,4,5,6",
    "start_time": "01:00-02:00",
    "keep_days": 7
  }
}
```

Disabling automated backup

```
{
```

```
"backup_policy" : {
  "keep_days" : 0
}
```

Example Responses

Status code: 204

No Content

```
{}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.5 Querying Instances that Can Be Restored

Function

This API is used to query all instances that can be restored.

Constraints

This API supports only GeminiDB Mongo instances.

URI

GET https://{Endpoint}/v3/{project_id}/backups/{backup_id}/restorable-instances

Table 5-291 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a user in a region. To obtain this value, see Obtaining a Project ID .
backup_id	Yes	String	Backup file ID.

Table 5-292 Query parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Index offset. The value must be no less than 0 . If this parameter is not transferred, the index offset is 0 by default.
limit	No	Integer	Maximum instances that can be restored. The value ranges from 1 to 100 . If this parameter is not transferred, the first 100 records are queried by default.

Request Parameters

Table 5-293 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-294 Response body parameters

Parameter	Type	Description
total_count	Integer	Total number of instances that can be restored.
restorable_instances	Array of QueryRestoreList objects	Information about the instances that can be restored.

Table 5-295 QueryRestoreList

Parameter	Type	Description
instance_id	String	Instance ID.
instance_mode	String	Instance type.

Parameter	Type	Description
engine_name	String	API name.
engine_version	String	API version.
vpc_id	String	VPC ID.
subnet_ids	Array of strings	Subnet IDs.
security_group_ids	Array of strings	Security group IDs.

Example Requests

URI example

GET <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/backups/2b19a3348108411baf7c46cf8d668e90br10/restorable-instances?offset=0&limit=100>

Example Responses

Status code: 200

Success

```
{
  "total_count": 1,
  "restorable_instances": [ {
    "instance_id": "6aef2786115341b6a18a67e0ee6ef664in10",
    "instance_mode": "ReplicaSet",
    "engine_name": "geminimongodb",
    "engine_version": "4.0.3.11",
    "vpc_id": "674e9b42-cd8d-4d25-a2e6-5abcc565b961",
    "subnet_ids": [ "f1df08c5-71d1-406a-aff0-de435a51007" ],
    "security_group_ids": [ "7aa51dbf-5b63-40db-9724-dad3c4828b58" ]
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.6 Querying the Time Window When a Backup Can Be Restored

Function

This API is used to query the time window when a backup can be restored.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Influx

This API does not support single-node GeminiDB Influx instances.

Make sure that full backup, incremental backup, and automated backup have been enabled. To enable incremental backup, contact customer service. This function can be used only when the next automated backup is performed.

This API can be used to query the time point that a backup can be restored to, so values of **start_time** and **end_time** are the same.

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/backups/restorable-time-periods

Table 5-296 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a user in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID, which can be obtained by calling the API described in Querying Instances and Details . If there are no instances available, call the API described in Creating an Instance to create one.

Table 5-297 Query parameters

Parameter	Mandatory	Type	Description
start_time	No	String	Start time point that the backup can be restored to. The time point is in the yyyy-mm-ddThh:mm:ssZ format. T indicates the start time, and Z indicates the time zone offset.

Parameter	Mandatory	Type	Description
end_time	No	String	End time point that the backup can be restored to. The time point is in the yyyy-mm-ddThh:mm:ssZ format. T indicates the start time, and Z indicates the time zone offset.
offset	No	Integer	Offset. The records after this offset will be queried. The default value is 0 .
limit	No	Integer	Maximum number of records displayed on each page. The value ranges from 0 to 1000 . The default value is 1000 .

Request Parameters

Table 5-298 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-299 Response body parameters

Parameter	Type	Description
total_count	Integer	Total time windows when a backup can be restored.
restorable_time_periods	Array of restorableTime objects	Time windows when a backup can be restored.

Table 5-300 restorableTime

Parameter	Type	Description
start_time	Long	Start time of the restoration time range in the UNIX timestamp format. The start time is the UTC time in milliseconds.
end_time	Long	End time point of the restoration time range in the UNIX timestamp format. The end time is the UTC time in milliseconds.

Example Request

- URI example

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/054e292c9880d4992f02c0196d3ea468/instances/a4d8ea2584e047439a667703c0684119in06/backups/restorable-time-periods?start_time=2022-06-01T18:50:20+0800&end_time=2022-06-01T19:50:20+0800&offset=0&limit=1000
```

Example Response

Status code: 200

Success

```
{
  "total_count": 1,
  "restorable_time_periods": [ {
    "start_time": 1607731200000,
    "end_time": 1607731200000
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.7 Creating a Manual Backup

Function

This API is used to create a manual backup.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra

- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

URI

POST https://{Endpoint}/v3/{project_id}/instances/{instance_id}/backups

Table 5-301 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>
instance_id	Yes	String	<p>Explanation: Instance ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Request Parameters

Table 5-302 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Table 5-303 Request body parameters

Parameter	Mandatory	Type	Description
name	Yes	String	<p>Explanation: Manual backup name.</p> <p>Values: Can include 4 to 64 characters and must start with a letter. It is case-sensitive and can contain only letters, digits, hyphens (-), and underscores (_).</p>
description	Yes	String	<p>Explanation: Manual backup description.</p> <p>Values: Can include a maximum of 256 characters and cannot contain the following special characters: >!<"&'=</p>

Response Parameters

Status code: 202

Table 5-304 Response body parameters

Parameter	Type	Description
job_id	String	Explanation: Task ID. Values: N/A
backup_id	String	Explanation: Backup ID. Values: N/A

Example Requests

- URI example
POST `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/054e292c9880d4992f02c0196d3ea468/instances/a4d8ea2584e047439a667703c0684119in06/backups`
- Creating a manual backup

```
{
  "name": " Create a manual backup",
  "description": " Creating a manual backup",
}
```

Example Responses

Status code: 202

Success

```
{
  "job_id" : "8061ceaf-b319-4315-9338-7f3de8e26f05",
  "backup_id" : "646d88d0b03f4fd2ae944ae2a33bcb6ain06"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.8 Deleting a Manual Backup

Function

This API is used to delete a manual backup.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

URI

DELETE https://{Endpoint}/v3/{project_id}/backups/{backup_id}

Table 5-305 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>
backup_id	Yes	String	<p>Explanation: Backup file ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Request Parameters

Table 5-306 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Response Parameters

Status code: 202

Table 5-307 Response body parameters

Parameter	Type	Description
job_id	String	<p>Explanation: Task ID.</p> <p>Default value: N/A</p>

Example Request

- URI example
DELETE https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/054e292c9880d4992f02c0196d3ea468/backups/5b0ae36cb8a746b68685a8fb588d8a15br06

Example Response

Status code: 202

Accepted

```
{
  "job_id" : "f85104b5-4a9c-4e0f-9505-fc5409d8f7ae"
}
```


Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.9 Restoring Data to an Existing Instance

Function

This API is used to restore data to an existing instance.

Constraints

This API supports GeminiDB Mongo instances.

The destination instance cannot be an instance that is undergoing a restoration task.

The destination instance must be of a version no earlier than and the specifications no lower than the source and have storage space no smaller than the size of the used backup file.

URI

POST https://{Endpoint}/v3/{project_id}/instances/{instance_id}/recovery

Table 5-308 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a user in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID, which can be obtained by calling the API described in Querying Instances and Details . If there are no instances available, call the API described in Creating an Instance to create one.

Request Parameters

Table 5-309 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-310 Request body parameters

Parameter	Mandatory	Type	Description
backup_id	Yes	String	Backup file name. You can restore a backup file to an existing instance.
password	No	String	<p>Instance password.</p> <p>The password can include 8 to 32 characters and contain uppercase letters, lowercase letters, digits, and the following special characters: ~!@#%^*-_+=? The password of GeminiDB Redis instances can contain at least two types of the following characters: uppercase letters, lowercase letters, digits, and special characters (~!@#%\$%^&*()-_+=?).</p> <ul style="list-style-type: none"> • If this parameter is not specified, the password stored in the backup file will overwrite the password of the existing instance after restoration is complete. • If this parameter is specified, the configured password will overwrite the password of the existing instance after restoration is complete.

Response Parameters

Status code: 202

Table 5-311 Response body parameters

Parameter	Type	Description
job_id	String	Task ID.

Example Requests

- URI example
POST `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/054e292c9880d4992f02c0196d3ea468/instances/61f554065802400c9c33f87e8114f081in10/recovery`
- Restoring data to an existing instance

```
{
  "backup_id" : "a8114c9b30cf42b6ba7752bfa62dee0bbr10",
  "password" : "*****"
}
```

Example Responses

Status code: 202

Accepted

```
{
  "job_id" : "f85104b5-4a9c-4e0f-9505-fc5409d8f7ae"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.10 Querying the Recycling Policy

Function

This API is used to query the recycling policy.

URI

GET `https://{Endpoint}/v3/{project_id}/instances/recycle-policy`

Table 5-312 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a user in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-313 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-314 Response body parameters

Parameter	Type	Description
recycle_policy	object	Recycling policy. For details, see Table 5-315 .

Table 5-315 RecyclePolicy

Parameter	Type	Description
retention_period_in_days	Integer	Policy retention duration (1 to 7 days). The value is a positive integer. The default value is 7 .

Example Requests

- URI example
GET `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/054e292c9880d4992f02c0196d3ea468/instances/recycle-policy`
- Example request body
None

Example Responses

Status code: 200

Success

```
{
  "recycle_policy": {
    "retention_period_in_days": 7
  }
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.11 Modifying the Recycling Policy

Function

This API is used to change a retention period for deleted instances. The new retention period is available to only those instances deleted after the change, but not to the instances already moved to the recycle bin before the change.

Constraints

The retention period for deleted instances can be 1 to 7 days.

URI

PUT https://{Endpoint}/v3/{project_id}/instances/recycle-policy

Table 5-316 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-317 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-318 Request body parameters

Parameter	Mandatory	Type	Description
recycle_policy	Yes	object	Recycling policy. For details, see Table 5-319 .

Table 5-319 RecyclePolicy

Parameter	Mandatory	Type	Description
retention_period_in_days	No	Integer	Policy retention duration (1 to 7 days). The value is a positive integer. The default value is 7.

Response Parameters

Status code: 200

No response parameters

Example Requests

- URI example
PUT `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/recycle-policy`
- Setting the retention period of instances in the recycle bin to 3 days

```
{
  "recycle_policy": {
    "retention_period_in_days": 3
  }
}
```

Example Responses

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.12 Querying Instances in the Recycle Bin

Function

This API is used to query all instances in the recycle bin.

URI

GET `https://{Endpoint}/v3/{project_id}/recycle-instances`

Table 5-320 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Table 5-321 Query parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Index offset. <ul style="list-style-type: none"> If offset is set to <i>N</i>, the resource query starts from the <i>N</i>+1 piece of data. If action is set to filter, offset is 0 by default, indicating that the query starts from the first piece of data. The value must be a positive integer.
limit	No	Integer	Maximum records to be queried. <ul style="list-style-type: none"> The value ranges from 1 to 100. If this parameter is not transferred, the first 100 records are queried by default.

Request Parameters

Table 5-322 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-323 Response body parameters

Parameter	Type	Description
total_count	Integer	Total number of records.
instances	Array of objects	Instance information. For details, see Table 5-324 .

Table 5-324 RecycleInstance

Parameter	Type	Description
id	String	Instance ID.
name	String	Instance name.
mode	String	Instance type. The value can be: <ul style="list-style-type: none"> • Cluster: GeminiDB Cassandra, GeminiDB Influx, or GeminiDB Redis instance in a proxy cluster • RedisCluster: GeminiDB Redis instance in a Redis cluster • InfluxdbSingle: single-node GeminiDB Influx instance • ReplicaSet: GeminiDB Mongo instance in a replica set
datastore	object	Database information For details, see Table 5-325 .
charge_mode	String	Billing mode. Options: <ul style="list-style-type: none"> • prePaid: indicates that the billing mode is yearly/monthly. • postPaid: indicates that the billing mode is pay-per-use.
enterprise_project_id	String	Enterprise project ID. The value 0 indicates that the default enterprise project is used.
backup_id	String	Backup ID.
created_at	String	Instance creation time.
deleted_at	String	Instance deletion time.
retained_until	String	Retention end time.

Table 5-325 RecycleDatastore

Parameter	Type	Description
type	String	Database type <ul style="list-style-type: none"> ● cassandra: GeminiDB Cassandra instance ● mongodb: GeminiDB Mongo instance ● influxdb: GeminiDB Influx instance ● redis: GeminiDB Redis instance
version	String	Database version

Example Requests

- URI example
GET `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/recycle-instances?offset=0&limit=100`
- Example request body
None

Example Responses

Status code: 200

Success

```
{
  "total_count": 1,
  "instances": [ {
    "id": "07fc12a8e0e94df7a3fcf53d0b5e1605in06",
    "name": "test",
    "mode": "Cluster",
    "datastore": {
      "type": "cassandra",
      "version": "3.11"
    },
  },
  "charge_mode": "postPaid",
  "enterprise_project_id": "0",
  "backup_id": "bf9ee62a7f7044c583c6765c916c36edbr02",
  "created_at": "2022-01-01T10:00:00",
  "deleted_at": "2022-02-01T11:00:00",
  "retained_until": "2022-02-02T11:00:00"
} ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.13 Obtaining GeminiDB Cassandra Instance Database Information That Is Restored Using Tables

Function

This API is used to obtain GeminiDB Cassandra instance database information that is restored using tables.

Constraints

- This API supports only GeminiDB Cassandra API.

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/databases

Table 5-326 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Table 5-327 Query parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	<p>Index position.</p> <ul style="list-style-type: none"> • Index position. The query starts from the dedicated resource created after the dedicated resource indexed by this parameter in a specified project. If offset is set to <i>N</i>, the resource query starts from the <i>N</i>+1 piece of data. • The value must be no less than 0. • If this parameter is not transferred, offset is set to 0 by default, indicating that the query starts from the latest created dedicated resource.

Parameter	Mandatory	Type	Description
limit	No	Integer	Maximum of dedicated resources to be queried. <ul style="list-style-type: none"> The value ranges from 1 to 100. If this parameter is not transferred, the first 100 instance records are queried by default.

Request Parameters

Table 5-328 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-329 Response body parameters

Parameter	Type	Description
total_count	Integer	Total number of records.
database_names	Array of strings	Database names.

Example Requests

```
GET https://{endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/a6d3c8a9857b4c81b3c1fe4802dfa4d0in06/databases?offset=0&limit=10
```

Example Responses

Status code: 200

Success

```
{
  "total_count": 1,
  "database_names": [ "db01" ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.14 Obtaining GeminiDB Cassandra Instance Table Information That Is Restored Using Tables

Function

This API is used to obtain GeminiDB Cassandra instance table information that is restored using tables.

Constraints

- This API supports only GeminiDB Cassandra API.

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/tables

Table 5-330 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Table 5-331 Query parameters

Parameter	Mandatory	Type	Description
database_name	Yes	String	Database name.

Parameter	Mandatory	Type	Description
offset	No	Integer	<p>Index position.</p> <ul style="list-style-type: none"> Index position. The query starts from the dedicated resource created after the dedicated resource indexed by this parameter in a specified project. If offset is set to <i>N</i>, the resource query starts from the <i>N</i>+1 piece of data. The value must be no less than 0. If this parameter is not transferred, offset is set to 0 by default, indicating that the query starts from the latest created dedicated resource.
limit	No	Integer	<p>Maximum of dedicated resources to be queried.</p> <ul style="list-style-type: none"> The value ranges from 1 to 100. If this parameter is not transferred, the first 100 instances are queried by default.

Request Parameters

Table 5-332 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-333 Response body parameters

Parameter	Type	Description
total_count	Integer	Total number of records.

Parameter	Type	Description
table_names	Array of strings	Table names.

Example Requests

```
GET https://{endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/a6d3c8a9857b4c81b3c1fe4802dfa4d0in06/tables?offset=0&limit=10&database_name=db01
```

Example Responses

Status code: 200

Success

```
{
  "total_count" : 1,
  "table_names" : [ "tb01" ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.15 Restoring the Current Redis Instance to a Point in Time

Function

This API is used to restore the current Redis instance to a point in time.

Constraints

- This API supports only GeminiDB Redis instances in a proxy cluster.
- Some instances of earlier versions need to be upgraded to use this API.

URI

```
PUT https://{Endpoint}/v3/{project_id}/redis/instances/{instance_id}/pitr
```

Table 5-334 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID

Request Parameters

Table 5-335 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-336 Request body parameters

Parameter	Mandatory	Type	Description
restore_time	Yes	String	Time point to which data is restored in the yyyy-mm-ddThh:mm:ssZ format. T is the separator between calendar and hourly notation of time, and Z indicates the time zone offset. For details, see parameter values under restore_time in Table 5-350 .

Response Parameters

Status code: 202

Table 5-337 Response body parameters

Parameter	Type	Description
job_id	String	Task ID

Status code: 400

Table 5-338 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Example Request

```
PUT https://{endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/redis/instances/
e73893ef73754465a8bd2e0857bbf13ein12/pitr
{
  "restore_time": "2022-06-01T18:50:20+0800"
}
```

Example Response

Status code: 202

Success

```
{
  "job_id": "20f793fd-2703-4339-bfb5-f1c9065b15c8"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.16 Setting the Policy for Restoring Redis Data to a Specified Time Point

Function

This API is used to set the policy for restoring Redis data to a specified time point.

Constraints

- This API supports only GeminiDB Redis instances in a proxy cluster.
- Some instances of earlier versions need to be upgraded to use this API.

URI

PUT https://{Endpoint}/v3/{project_id}/redis/instances/{instance_id}/pitr/policy

Table 5-339 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID

Request Parameters

Table 5-340 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-341 Request body parameters

Parameter	Mandatory	Type	Description
enabled	Yes	Boolean	Whether to enable point-in-time restoration for a Redis instance. <ul style="list-style-type: none"> true: indicates that the point-in-time restoration is enabled for the instance. false: indicates that the point-in-time restoration is disabled for the instance.
interval	No	Integer	Interval for backing up data. The default value is 20 , in minutes, and the value ranges from 5 to 120. For example, if interval is set to 20 minutes, data will be recovered every 20 minutes. Shorter intervals may affect performance and require more storage space.
keep_days	No	Integer	Backup retention days. The default value is 1 , in days, and the value ranges from 1 to 7.

Response Parameters

Status code: 400

Table 5-342 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Example Request

```
PUT https://{endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/redis/instances/
e73893ef73754465a8bd2e0857bbf13ein12/pitr/policy
{
  "enabled" : true,
  "interval" : 20,
  "keep_days" : 7
}
```

Example Response

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.17 Querying the Policy for Restoring Redis Data to a Specified Time Point

Function

This API is used to query the policy for restoring Redis data to a specified time point.

Constraints

- This API supports only GeminiDB Redis instances in a proxy cluster.
- Some instances of earlier versions need to be upgraded to use this API.

URI

GET https://{Endpoint}/v3/{project_id}/redis/instances/{instance_id}/pitr/policy

Table 5-343 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-344 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

Status code: 200

Table 5-345 Response body parameters

Parameter	Type	Description
enabled	Boolean	Specifies whether to enable point-in-time restoration for a Redis instance. <ul style="list-style-type: none"> • true: indicates that the point-in-time restoration is enabled for the instance. • false: indicates that the point-in-time restoration is disabled for the instance.
interval	Integer	Interval for backing up data. This parameter is returned only when restoring Redis data to a specified time point is enabled.
keep_days	Integer	Days to retain the generated backups. This parameter is returned only when restoring Redis data to a specified time point is enabled.

Status code: 400

Table 5-346 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Example Request

```
GET https://{endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/redis/instances/e73893ef73754465a8bd2e0857bbf13ein12/pitr/policy
```

Example Response

Status code: 200

Success

```
{
  "enabled" : true,
  "interval" : 20,
  "keep_days" : 7
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.18 Querying the Restoration Time Range of a Redis Instance

Function

This API is used to query the restoration time range of a Redis instance.

Constraints

- This API supports only GeminiDB Redis instances in a proxy cluster.
- Some instances of earlier versions need to be upgraded to use this API.

URI

GET https://{Endpoint}/v3/{project_id}/redis/instances/{instance_id}/pitr/restorable-time-periods

Table 5-347 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Table 5-348 Query parameters

Parameter	Mandatory	Type	Description
start_time	Yes	String	Start time point that the backup can be restored to. The time point is in the yyyy-mm-ddThh:mm:ssZ format. T is the separator between calendar and hourly notation of time, and Z indicates the time zone offset.
end_time	Yes	String	End time point that the backup can be restored to. The time point is in the yyyy-mm-ddThh:mm:ssZ format. T is the separator between calendar and hourly notation of time, and Z indicates the time zone offset.
offset	No	Integer	Offset. The records after this offset will be queried. The default value is 0 .
limit	No	Integer	Maximum number of records displayed on each page. The value ranges from 0 to 100 . The default value is 300 .

Request Parameters

Table 5-349 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

Status code: 200

Table 5-350 Response body parameters

Parameter	Type	Description
restore_time	Array of strings	Time points that the Redis backup can be restored to. The time point is in the yyyy-mm-ddThh:mm:ssZ format. T is the separator between calendar and hourly notation of time, and Z indicates the time zone offset.
total_count	Integer	Total time points that a Redis backup can be restored to.

Status code: 400

Table 5-351 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Example Request

```
GET https://{endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/redis/instances/
e73893ef73754465a8bd2e0857bbf13ein12/pitr/restorable-time-periods?
start_time=2022-06-01T18:50:20+0800&end_time=2022-06-01T19:50:20+0800&offset=0&limit=100
```

Example Response

Status code: 200

Success

```
{
  "total_count": 2,
  "restore_time": [
    "2022-06-01T18:50:20+0800",
    "2022-06-01T18:55:20+0800"
  ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.19 Querying the Storage Space Used for Restoring a Redis Instance to a Specified Time Point

Function

This API is used to query the storage space used for restoring a Redis instance to a specified time point.

Constraints

- This API supports only GeminiDB Redis instances in a proxy cluster.
- Some instances of earlier versions need to be upgraded to use this API.

URI

GET https://{Endpoint}/v3/{project_id}/redis/instances/{instance_id}/pitr

Table 5-352 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-353 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

Status code: 200

Table 5-354 Response body parameters

Parameter	Type	Description
storage	String	Storage space required to restore a Redis instance to a specified time point. Unit: GB

Status code: 400

Table 5-355 Response body parameters

Parameter	Type	Description
error_code	String	Error code.
error_msg	String	Error message.

Example Request

```
GET https://{endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/redis/instances/
e73893ef73754465a8bd2e0857bbf13ein12/pitr
```

Example Response

Status code: 200

Success

```
{
  "storage": "10"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.20 Stopping a Backup

Function

This API is used to stop a backup in case of an emergency.

Constraints

- Currently, a backup can be stopped only in emergency scenarios of GeminiDB Cassandra instances.
- This API is available only for DB instances in the **Backing up** or **Uploading backup** state.

URI

```
PUT https://{Endpoint}/v3/{project_id}/backups/{backup_id}
```


Table 5-356 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
backup_id	Yes	String	Backup ID

Request Parameters

Table 5-357 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-358 Request body parameters

Parameter	Mandatory	Type	Description
action	Yes	String	Operation. stop indicates the operation for stopping a backup.

Response Parameters

Status code: 202

Table 5-359 Response body parameters

Parameter	Type	Description
job_id	String	ID of workflow for stopping a backup

Example Request

```
PUT https://{endpoint}/v3/054e292c9880d4992f02c0196d3ea468/backups/
ca8cb4638f3d42bfa6e942464339be89br14
{
  "action": "stop"
}
```

Example Response

Status code: 202

Success

```
{
  "job_id": "b52cc59d-e1e2-4e63-bb00-7694e126f66a"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.5.21 Deleting Manual Backups in Batches

Function

This API is used to delete manual backups of a DB instance in batches.

URI

DELETE https://{Endpoint}/v3/{project_id}/instances/backups

Table 5-360 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-361 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token
X-Language	No	String	Language

Table 5-362 Request body parameters

Parameter	Mandatory	Type	Description
backup_ids	Yes	Array of String	IDs of manual backups to be deleted in batches. No more than 10 manual backups can be deleted at a time.

Response Parameters

Status code: 200

No response parameters

Example Request

```
DELETE https://{endpoint}/v3/054e292c9880d4992f02c0196d3ea468/instances/backups
{
  "backup_ids": ["030acca1ad98444fb95a2e3e92d5b82fbr13", "be5b385ef95e48afa1222e9f2a954e58br13"]
}
```

Example Response

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.6 Parameter Templates

5.6.1 Obtaining Parameter Templates

Function

This API is used to obtain parameter templates, including all of the default and custom parameter templates.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

- GeminiDB Redis

URI

GET https://{Endpoint}/v3.1/{project_id}/configurations

Table 5-363 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Table 5-364 Query parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Index offset. <ul style="list-style-type: none"> • If offset is set to <i>N</i>, the resource query starts from the <i>N</i>+1 piece of data. If action is set to filter, offset is 0 by default, indicating that the query starts from the first piece of data. • The value must be a positive integer.
limit	No	Integer	Maximum number of instances that can be queried. <ul style="list-style-type: none"> • The value ranges from 1 to 100. • If this parameter is not transferred, the first 100 records are queried by default.

Request Parameters

Table 5-365 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-366 Response body parameters

Parameter	Type	Description
count	Integer	Total number of records.
quota	Integer	Maximum number of custom parameter templates that a user can create.
configurations	Array of ListConfigurations-Result objects	Parameter templates.

Table 5-367 ListConfigurationsResult

Parameter	Type	Description
id	String	Parameter template ID.
name	String	Parameter template name.
description	String	Parameter template description.
datastore_version_name	String	Database version name.
datastore_name	String	Database name.
created	String	Creation time in the <i>yyyy-MM-ddTHH:mm:ssZ</i> format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.
updated	String	Update time in the <i>yyyy-MM-ddTHH:mm:ssZ</i> format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.

Parameter	Type	Description
mode	String	Instance type. The value can be: Cluster : GeminiDB Cassandra instance in a cluster ReplicaSet : GeminiDB Mongo instance in a replica set Sharding : GeminiDB Mongo instance in a sharded cluster Cluster : GeminiDB Influx instance in a cluster InfluxdbSingle : single-node GeminiDB Influx instance Cluster : GeminiDB Redis instance in a proxy cluster Replication : primary/standby GeminiDB Redis instance
user_defined	Boolean	Whether the parameter template is a custom template. The value can be: <ul style="list-style-type: none"> false, indicating that the parameter template is a default parameter template. true, indicating that the parameter template is a custom template.

Example Requests

URI example

GET <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3.1/375d8d8fad1f43039e23d3b6c0f60a19/configurations?offset=0&limit=10>

Example Responses

Status code: 200

Success

```
{
  "count" : 2,
  "quota": 100,
  "configurations" : [ {
    "id" : "887ea0d1bb0843c49e8d8e5a09a95652pr06",
    "name" : "configuration_test",
    "description" : "configuration_test",
    "datastore_version_name" : "3.11",
    "datastore_name" : "cassandra",
    "created" : "2019-05-15T11:53:34+0000",
    "updated" : "2019-05-15T11:53:34+0000",
    "mode": "Cluster",
    "user_defined" : true
  }, {
    "id" : "3bc1e9cc0d34404b9225ed7a58fb284epr06",
    "name" : "Default-Cassandra-3.11",
    "description" : "Default parameter group for cassandra 3.11",
    "datastore_version_name" : "3.11",
    "datastore_name" : "cassandra",
    "created" : "2019-05-27T03:38:51+0000",
```

```
"updated" : "2019-05-27T03:38:51+0000",
"mode": "Cluster",
"user_defined" : false
}]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.6.2 Creating a Parameter Template

Function

This API is used to create a parameter template and configure the name, description, DB engine version, and parameter values in the parameter template.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

The new parameter template cannot have the same name as any existing parameter template.

For configuration item **values**, you can enter system-defined parameters that allow for modification.

URI

POST https://{Endpoint}/v3/{project_id}/configurations

Table 5-368 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-369 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-370 Request body parameters

Parameter	Mandatory	Type	Description
name	Yes	String	Parameter template name. It can include a maximum of 64 characters and can contain only uppercase letters, lowercase letters, digits, hyphens (-), underscores (_), and periods (.).
description	No	String	Parameter template description. It can contain a maximum of 256 characters except the following special characters: >!<"&'= The value is left blank by default.
values	No	Map<String,String>	Parameter values defined by users based on a default parameter template. Keep the parameter values unchanged by default.
datastore	Yes	ConfigurationDatastoreOption object	Database object.

Table 5-371 ConfigurationDatastoreOption

Parameter	Mandatory	Type	Description
type	Yes	String	Database type cassandra: GeminiDB Cassandra instance mongodb: GeminiDB Mongo instance influxdb: GeminiDB Influx instance

Parameter	Mandatory	Type	Description
version	Yes	String	Database version 3.11: GeminiDB Cassandra instance 3.11 4.0: GeminiDB Mongo instance 4.0 1.8: GeminiDB Influx instance 1.8
mode	No	String	Database deployment mode. This parameter is mandatory for GeminiDB Mongo. The options are as follows: <ul style="list-style-type: none"> • ReplicaSet, indicating that the instance is a GeminiDB Mongo replica set. • InfluxdbSingle, indicating that the instance is a single-node GeminiDB Influx instance.

Response Parameters

Status code: 200

Table 5-372 Response body parameters

Parameter	Type	Description
configuration	ConfigurationResult object	Parameter template information.

Table 5-373 ConfigurationResult

Parameter	Type	Description
id	String	Parameter template ID.
name	String	Parameter template name.
datastore_version_name	String	Database version name.
datastore_name	String	Database name.
description	String	Parameter template description

Parameter	Type	Description
created	String	Creation time in the <i>yyyy-MM-ddTHH:mm:ssZ</i> format.
updated	String	Update time in the <i>yyyy-MM-ddTHH:mm:ssZ</i> format.

Example Requests

- **URI example**
POST `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/configurations`
- **Creating a parameter template for GeminiDB Cassandra instances**

```
{
  "name": "configuration_test",
  "description": "configuration_test",
  "values": {
    "max_connections": "10",
    "autocommit": "OFF"
  },
  "datastore": {
    "type": "cassandra",
    "version": "3.11"
  }
}
```

Example Responses

Status code: 200

Success

```
{
  "configuration": {
    "id": "463b4b58d0e84e2b95605dea4552fdpr06",
    "name": "configuration_test",
    "datastore_version_name": "3.11",
    "datastore_name": "cassandra",
    "description": "configuration_test",
    "created": "2020-03-09T08:27:56+0800",
    "updated": "2020-03-09T08:27:56+0800"
  }
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.6.3 Modifying Parameters in a Parameter Template

Function

This API is used to modify parameters in a specified parameter template, including parameter names, descriptions, and values.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

The modified parameter template name must be different from the name of any existing or default parameter template.

Default parameter templates cannot be modified.

For configuration item **values**, you can enter system-defined parameters that allow for modification.

URI

PUT https://{Endpoint}/v3/{project_id}/configurations/{config_id}

Table 5-374 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
config_id	Yes	String	Parameter template ID.

Request Parameters

Table 5-375 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-376 Request body parameters

Parameter	Mandatory	Type	Description
name	No	String	Parameter template name. It contains a maximum of 64 characters and can contain only letters, digits, hyphens (-), underscores (_), and periods (.).
description	No	String	Parameter template description. The value is left blank by default. It contains a maximum of 256 characters except the following special characters: !<>=&"
values	No	Map<String,String>	Parameter values defined by users based on a default parameter template. If this parameter is not specified, its value will not be changed.

Response Parameters

Status code: 200

No response parameters

Example Request

- **URI example**
PUT <https://gaussdb-nosql.cn-north-4.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/configurations/e02e76567ae04662a2753492b77f965bpr06>
- **URI example**
PUT <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/configurations/e02e76567ae04662a2753492b77f965bpr06>
- **Modifying Parameters in a Parameter Template**

NOTE

At least one parameter in the request body must be specified. Otherwise, the request cannot be delivered.

```
{
  "name": "configuration_test",
  "description": "configuration_test",
  "values": {
    "concurrent_reads": "64"
  }
}
```

Example Response

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.6.4 Applying a Parameter Template

Function

This API is used to apply a parameter template to one or more instances.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

This API is an asynchronous API. A successful response does not indicate that the parameter template is successfully applied.

URI

PUT https://{Endpoint}/v3.1/{project_id}/configurations/{config_id}/apply

Table 5-377 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
config_id	Yes	String	Parameter template ID.

Request Parameters

Table 5-378 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-379 Request body parameters

Parameter	Mandatory	Type	Description
instance_ids	Yes	Array of strings	Instance IDs.

Response Parameters

Status code: 202

Table 5-380 Response body parameters

Parameter	Type	Description
job_id	String	ID of an asynchronous task that applies the parameter template
success	Boolean	Whether the task for applying the parameter template is successfully submitted. The value can be: <ul style="list-style-type: none"> true: The task is successfully submitted. false: The task failed to be submitted.

Example Request

- **URI example**
PUT `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3.1/375d8d8fad1f43039e23d3b6c0f60a19/configurations/e02e76567ae04662a2753492b77f965bpr06/apply`
- **Applying a Parameter Template**

```
{
  "instance_ids": [ "73ea2bf70c73497f89ee0ad4ee008aa2in06" ]
}
```

Example Response

Status code: 202

Success

```
{
  "job_id": "463b4b58-d0e8-4e2b-9560-5dea4552fde9",
  "success": true
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.6.5 Resetting a Custom Parameter Template

Function

This API is used to reset a custom parameter template.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

URI

POST https://{Endpoint}/v3/{project_id}/configurations/{config_id}/reset

Table 5-381 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
config_id	Yes	String	Parameter template ID.

Request Parameters

Table 5-382 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

None

Example Requests

URI example

POST https://{Endpoint}/v3/2900b7b8d03e4619b8db8d43bc6234ee/configurations/384b6bfbe562407d9e8c4dd6ed15cb42pr13/reset

Example Responses

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.6.6 Modifying Parameters of a Specified Instance

Function

This API is used to modify parameters of a specified instance.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

For configuration item **values**, you can enter system-defined parameters that allow for modification.

This API is an asynchronous API. A successful response does not indicate that the parameters are successfully modified.

URI

PUT https://{Endpoint}/v3.1/{project_id}/instances/{instance_id}/configurations

Table 5-383 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-384 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-385 Request body parameters

Parameter	Mandatory	Type	Description
values	Yes	Map<String,String>	Parameter values defined by users based on a default parameter template.

Response Parameters

Status code: 202

Table 5-386 Response body parameters

Parameter	Type	Description
job_id	String	ID of the asynchronous task for modifying instance parameters.
restart_required	Boolean	Whether the instance needs to be restarted. The value can be: <ul style="list-style-type: none"> true, indicating that the instance needs to be restarted. false, indicating that the instance does not need to be restarted.

Example Request

- **URI example**
PUT https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3.1/054e292c9880d4992f02c0196d3ea468/instances/392850e624504e1490901d50b585a60din06/configurations
- **Modifying Parameters of a Specified Instance**

```
{
  "values" : {
    "request_timeout_in_ms" : "10000"
  }
}
```

Example Response

Status code: 202

Success

```
{
  "job_id" : "463b4b58-d0e8-4e2b-9560-5dea4552fde9",
  "restart_required" : false
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.6.7 Querying Instance Parameter Settings

Function

This API is used to query instance parameter settings.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/configurations

Table 5-387 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Parameter	Mandatory	Type	Description
instance_id	Yes	String	Explanation: Instance ID. Constraints: N/A Values: N/A Default value: N/A

Request Parameters

Table 5-388 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication . Constraints: N/A Values: N/A Default value: N/A

Response Parameters

Status code: 200

Table 5-389 Response body parameters

Parameter	Type	Description
datastore_version_name	String	Explanation: Database version name. Values: N/A

Parameter	Type	Description
datastore_name	String	<p>Explanation: Database name.</p> <p>Values: N/A</p>
created	String	<p>Explanation: Creation time</p> <p>Constraints: Format: <i>yyyy-MM-ddTHH:mm:ssZ</i> T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.</p> <p>Values: N/A</p>
updated	String	<p>Explanation: Update time</p> <p>Constraints: Format: <i>yyyy-MM-ddTHH:mm:ssZ</i> T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.</p> <p>Values: N/A</p>
id	String	<p>Explanation: Parameter template ID.</p> <p>Values: N/A</p>

Parameter	Type	Description
mode	String	<p>Explanation: Instance type. The value can be:</p> <p>Values:</p> <ul style="list-style-type: none"> • Cluster: GeminiDB Cassandra instance in a cluster • ReplicaSet, indicating that the instance is of the GeminiDB Mongo replica set type. • Sharding, indicating that the instance is of the GeminiDB Mongo cluster type. • Cluster, indicating that the instance is of the GeminiDB Influx cluster type. • InfluxdbSingle: single-node GeminiDB Influx instance • Cluster: GeminiDB Redis instance in a proxy cluster • Replication: primary/standby GeminiDB Redis instance
configuration_parameters	Array of ConfigurationParameterResult objects	<p>Explanation: Parameters defined by users based on a default parameter template.</p>

Table 5-390 ConfigurationParameterResult

Parameter	Type	Description
name	String	<p>Explanation: Parameter name.</p> <p>Values: N/A</p>
value	String	<p>Explanation: Parameter value.</p> <p>Values: N/A</p>

Parameter	Type	Description
restart_required	Boolean	<p>Explanation: Whether the instance needs to be restarted. The value can be:</p> <p>Values:</p> <ul style="list-style-type: none"> • false, indicating that the instance does not need to be restarted. • true, indicating that the instance needs to be restarted.
readonly	Boolean	<p>Explanation: Whether the parameter is read-only. The value can be:</p> <p>Values:</p> <ul style="list-style-type: none"> • false, indicating that the parameter is not read-only. • true, indicating that the parameter is read-only.
value_range	String	<p>Explanation: Value range. For example, the value of the Integer type ranges from 0 to 1, and the value of the Boolean type is true or false.</p> <p>Values: N/A</p>
type	String	<p>Explanation: Parameter type</p> <p>Values: The value can be string, integer, boolean, list, or float.</p>
description	String	<p>Explanation: Parameter description.</p> <p>Values: N/A</p>

Example Requests

URI example

GET <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/9136fd2a9fcd405ea4674276ce36dae8in02/configurations>

Example Responses

Status code: 200

Success

```
{
  "datastore_version_name" : "3.11",
  "datastore_name" : "cassandra",
  "created" : "2020-03-21 11:40:44",
  "updated" : "2020-03-21 11:40:44",
  "id" : "9ad6bc82146e4043a50c963ab3bf09adpr06",
  "mode" : "Cluster",
  "configuration_parameters" : [ {
    "name" : "concurrent_reads",
    "value" : "64",
    "restart_required" : true,
    "readonly" : true,
    "value_range" : "4-512",
    "type" : "integer",
    "description" : "Number of concurrent read threads."
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.6.8 Obtaining Parameters of a Specified Parameter Template

Function

This API is used to obtain information about parameters of a specified parameter template.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

URI

GET https://{Endpoint}/v3/{project_id}/configurations/{config_id}

Table 5-391 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Parameter	Mandatory	Type	Description
config_id	Yes	String	Parameter template ID.

Request Parameters

Table 5-392 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-393 Response body parameters

Parameter	Type	Description
id	String	Parameter template ID.
name	String	Parameter template name.
description	String	Parameter template description.
datastore_version_name	String	Database version name.
datastore_name	String	Database name.
created	String	Creation time in the yyyy-MM-ddTHH:mm:ssZ format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.
updated	String	Update time in the yyyy-MM-ddTHH:mm:ssZ format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.
configuration_parameters	Array of ConfigurationParameterResult objects	Parameters defined by users based on a default parameter template.

Table 5-394 ConfigurationParameterResult

Parameter	Type	Description
name	String	Parameter name.
value	String	Parameter value.
restart_required	Boolean	Whether the instance needs to be restarted. The value can be: <ul style="list-style-type: none"> • false, indicating that the instance does not need to be restarted. • true, indicating that the instance needs to be restarted.
readonly	Boolean	Whether the parameter is read-only. The value can be: <ul style="list-style-type: none"> • false, indicating that the parameter is not read-only. • true, indicating that the parameter is read-only.
value_range	String	Value range. For example, the value of the Integer type ranges from 0 to 1 , and the value of the Boolean type is true or false .
type	String	Parameter type. The value can be string , integer , boolean , list , or float .
description	String	Parameter description.

Example Requests

URI example

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/configurations/e02e76567ae04662a2753492b77f965bpr06
```

Example Responses

Status code: 200

Success

```
{
  "id" : "07fc12a8e0e94df7a3fcf53d0b5e1605pr06",
  "name" : "default-cassandra-3.11",
  "datastore_version_name" : "3.11",
  "datastore_name" : "cassandra",
  "description" : "Default parameter group for cassandra 3.11",
  "created" : "2020-03-21T04:40:51+0800",
  "updated" : "2020-03-21T04:40:51+0800",
  "configuration_parameters" : [ {
    "name" : "concurrent_reads",
    "value" : "64",
    "restart_required" : true,
    "readonly" : true,
```

```
"value_range" : "4-512",
"type" : "integer",
"description" : "Number of concurrent read threads."
}]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.6.9 Deleting a Parameter Template

Function

This API is used to delete a specified parameter template.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

URI

DELETE https://{Endpoint}/v3/{project_id}/configurations/{config_id}

Table 5-395 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
config_id	Yes	String	Parameter template ID.

Request Parameters

Table 5-396 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

No response parameters

Example Requests

URI example

```
DELETE https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/configurations/e02e76567ae04662a2753492b77f965bpr06
```

Example Responses

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.6.10 Querying Instances that a Parameter Template Can Be Applied To

Function

This API is used to query instances that a parameter template can be applied to.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

URI

GET https://{Endpoint}/v3/{project_id}/configurations/{config_id}/applicable-instances

Table 5-397 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Parameter	Mandatory	Type	Description
config_id	Yes	String	Parameter template ID.

Table 5-398 Query parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Index offset. If offset is set to <i>N</i> , the resource query starts from the <i>N</i> +1 piece of data. If action is set to filter , offset is 0 by default, indicating that the query starts from the first piece of data. The offset value must be a number but cannot be a negative number.
limit	No	Integer	Maximum records to be queried. <ul style="list-style-type: none"> The value ranges from 1 to 100. If this parameter is not transferred, the first 100 records are queried by default.

Request Parameters

Table 5-399 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-400 Response body parameters

Parameter	Type	Description
instances	Array of ApplicableInstanceRsp objects	All instances.
count	Integer	Maximum number of instances that parameters can be applied to.

Table 5-401 ApplicableInstanceRsp

Parameter	Type	Description
id	String	Instance ID.
name	String	Instance name.

Example Requests

- URI example
GET <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/0549b4a43100d4f32f51c01c2fe4acdb/configurations/9e80bf6bbd7142f49761c07e9c32dd04pr06/applicable-instances?offset=0&limit=10>

Example Responses

Status code: 200

Success

```
{
  "instances" : [ {
    "id" : "f38e203908bd4fae82714e88f12600f6in06",
    "name" : "test"
  } ],
  "count" : 1000
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.6.11 Viewing Parameter Change History of an Instance

Function

This API is used to view change history of parameters of an instance.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

This API can be used to query only the past seven days of parameter changes.

GeminiDB Influx does not allow you to view parameter change history of single-node instances.

URI

GET `https://{Endpoint}/v3/{project_id}/instances/{instance_id}/configuration-histories`

Table 5-402 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>
instance_id	Yes	String	<p>Explanation: Instance ID</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Table 5-403 Query parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	<p>Explanation: Index offset.</p> <p>If offset is set to <i>N</i>, the resource query starts from the <i>N</i>+1 piece of data. If action is set to filter, offset is 0 by default, indicating that the query starts from the first piece of data.</p> <p>Constraints: The value must be a positive number or 0.</p> <p>Values: 0 or a positive integer</p> <p>Default value: 0</p>
limit	No	Integer	<p>Explanation: Maximum records to be queried</p> <p>Constraints: N/A</p> <p>Values: 1 to 100</p> <p>Default value: If this parameter is not transferred, the first 100 records are queried by default.</p>

Request Parameters

Table 5-404 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Response Parameters

Status code: 200

Table 5-405 Response body parameters

Parameter	Type	Description
histories	Array of ConfigurationHistoryRsp objects	<p>Explanation: Change history of parameters of an instance.</p>

Table 5-406 ConfigurationHistoryRsp

Parameter	Type	Description
parameter_name	String	<p>Explanation: Parameter name</p> <p>Values: N/A</p>
old_value	String	<p>Explanation: Old parameter value</p> <p>Values: N/A</p>

Parameter	Type	Description
new_value	String	Explanation: New parameter value Values: N/A
update_result	String	Explanation: Update result Values: <ul style="list-style-type: none"> • SUCCESS • FAILED
applied	Boolean	Explanation: Whether a change is applied Values: <ul style="list-style-type: none"> • true: A change is applied. • false: A change is not applied.
updated_at	String	Explanation: Update time Values: Format: <i>yyyy-MM-ddTHH:mm:ssZ</i> T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.
applied_at	String	Explanation: Time when a change is applied Values: Format: <i>yyyy-MM-ddTHH:mm:ssZ</i> T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.

Example Request

- URI example
GET <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/0549b4a43100d4f32f51c01c2fe4acdb/instances/c4e095105bc64797bc3be633ae7201eein10/configuration-histories?offset=0&limit=10>

Example Response

Status code: 200

Success

```
{
  "histories": [ {
    "parameter_name": "mongos.connPoolMaxShardedConnsPerHost",
    "old_value": "600",
    "new_value": "500",
    "update_result": "FAILED",
    "applied": true,
    "updated_at": "2022-09-20T11:17:04+0000",
    "applied_at": "2022-09-20T11:17:04+0000"
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.6.12 Viewing Application Records of a Parameter Template

Function

This API is used to view application records of a parameter template.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

After an instance is deleted, application records of the parameter template that the instance uses are also deleted.

URI

GET https://{Endpoint}/v3/{project_id}/configurations/{config_id}/applied-histories

Table 5-407 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
config_id	Yes	String	Parameter template ID.

Table 5-408 Query parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Index offset. If offset is set to <i>N</i> , the resource query starts from the <i>N</i> +1 piece of data. If action is set to filter , offset is 0 by default, indicating that the query starts from the first piece of data. The offset value must be a number but cannot be a negative number.
limit	No	Integer	Maximum records to be queried. <ul style="list-style-type: none"> The value ranges from 1 to 100. If this parameter is not transferred, the first 100 records are queried by default.

Request Parameters

Table 5-409 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-410 Response body parameters

Parameter	Type	Description
histories	Array of ApplyHistoryRsp objects	Application records of a parameter template.

Table 5-411 ApplyHistoryRsp

Parameter	Type	Description
instance_id	String	Instance ID.
instance_name	String	Instance name.
applied_at	String	Effective time in the yyyy-MM-ddTHH:mm:ssZ format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.
apply_result	String	<ul style="list-style-type: none"> • SUCCESS: indicates that the parameter template is applied to the corresponding instance. • Applying: indicates that the parameter template is being applied to the corresponding instance. • FAILED: indicates that the parameter template fails to be applied to the corresponding instance.
failure_reason	String	Failure cause.

Example Requests

- **URI example**
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/056f86e8d480d3cb2f43c00183f75e1f/configurations/e02e76567ae04662a2753492b77f965bpr06/applied-histories?offset=0&limit=10

Example Responses

Status code: 200

Success

```
{
  "histories": [
    {
      "instance_id": "a2d0cf32db3e4f2aa3a684240e10b457in06",
      "instance_name": "test",
      "applied_at": "2022-09-20T11:17:04+0000",
      "apply_result": "SUCCESS",
      "failure_reason": ""
    }
  ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.6.13 Comparing Parameter Templates

Function

This API is used to compare two parameter templates.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

This API only compares parameter templates with one of the same node type and DB engine to learn about configurations of the current template.

URI

POST https://{Endpoint}/v3/{project_id}/configurations/comparison

Table 5-412 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-413 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-414 Request body parameters

Parameter	Mandatory	Type	Description
source_configuration_id	Yes	String	ID of the source parameter template to be compared.

Parameter	Mandatory	Type	Description
target_configuration_id	Yes	String	ID of the destination parameter template for comparison.

Response Parameters

Status code: 202

Table 5-415 Response body parameters

Parameter	Type	Description
differences	Array of Table 5-416 objects	Differences between parameters.

Table 5-416 DiffDetails

Parameter	Type	Description
parameter_name	String	Parameter name.
source_value	String	Parameter value in the source parameter template.
target_value	String	Parameter value in the destination parameter template.

Example Requests

- URI example
POST `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/configurations/comparison`
- Comparing a source parameter template with the target parameter template

```
{
  "source_configuration_id": "0764fdcd949b411ba76c2b762b80c212pr06",
  "target_configuration_id": "fa42c57bb62844e490052f2ff9d5a264pr06"
}
```

Example Responses

Status code: 202

Accepted

```
{
  "differences": [ {
    "parameter_name": "batch_size_fail_threshold_in_kb",
    "source_value": "1000",
```

```
"target_value" : "5000"
}]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.6.14 Replicating a Parameter Template

Function

This API is used to replicate a parameter template.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

The parameter template generated after replication cannot have the same name as the default parameter template or an existing template.

Only custom parameter templates can be replicated.

URI

POST https://{Endpoint}/v3/{project_id}/configurations/{config_id}/copy

Table 5-417 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
config_id	Yes	String	Parameter template ID.

Request Parameters

Table 5-418 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-419 Request body parameters

Parameter	Mandatory	Type	Description
name	Yes	String	Name of the parameter template generated after replication. The name can include a maximum of 64 characters and can contain only uppercase letters, lowercase letters, digits, hyphens (-), underscores (_), and periods (.).
description	No	String	Parameter template description. The description can contain a maximum of 256 characters except the following special characters: >! <"&'= The value is left blank by default.

Response Parameters

Status code: 202

Table 5-420 Response body parameters

Parameter	Type	Description
config_id	String	ID of the replicated parameter template.

Example Requests

- URI example
POST `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/configurations/e02e76567ae04662a2753492b77f965bpr06/copy`
- Replicating a Parameter Template

```
{
  "name": "paramsGroup-2434",
  "description": "Replicating a parameter template"
}
```

Example Responses

Status code: 202

Accepted


```
{
  "config_id" : "7b4e07852bd54016906e89461b3182cdpr06"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.6.15 Querying API that Support Parameter Templates

Function

This API is used to query API that support parameter templates.

URI

GET https://{Endpoint}/v3/{project_id}/configurations/datastores

Table 5-421 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-422 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-423 Response body parameters

Parameter	Type	Description
datastores	Array of Table 5-424 objects	DB API information.

Table 5-424 DataStoreList

Parameter	Type	Description
datastore_name	String	DB API.
version	String	DB API version.
mode	String	Instance type. The value can be: Cluster: GeminiDB Cassandra, GeminiDB Influx, or GeminiDB Redis instance in a proxy cluster InfluxdbSingle: single-node GeminiDB Influx instance ReplicaSet: GeminiDB Mongo instance in a replica set

Example Requests

- URI example
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/configurations/datastores
- Example request body
None

Example Responses

Status code: 200

Success

```
{
  "datastores": [ {
    "datastore_name": "mongodb",
    "mode": "ReplicaSet",
    "version": "4.0"
  }, {
    "datastore_name": "influxdb",
    "mode": "InfluxdbCluster",
    "version": "1.8"
  }, {
    "datastore_name": "cassandra",
    "mode": "ReplicaCircle",
    "version": "3.11"
  }
]
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.7 Managing Databases and Accounts

5.7.1 Creating a Database Account

Function

This API is used to create a database account for a GeminiDB Redis instance.

Constraints

- Only the GeminiDB Redis API is supported.
- This operation cannot be performed when the instance is in any of the following states: creating, changing specifications, changing database port, frozen, or abnormal.

URI

POST https://{Endpoint}/v3/{project_id}/redis/instances/{instance_id}/db-users

Table 5-425 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-426 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-427 Request body parameters

Parameter	Mandatory	Type	Description
users	No	Array of objects	All accounts to be created. For details, see Table 5-428 .

Table 5-428 UserForCreation

Parameter	Mandatory	Type	Description
name	Yes	String	Account name. This parameter cannot be empty. It must start with a letter and contain up to 36 characters, including only numbers, letters, hyphens (-), and underscores (_).
password	Yes	String	<ul style="list-style-type: none"> Account password. The password can contain 8 to 32 characters. The password must contain at least two of the following types: uppercase letters, lowercase letters, digits, and special characters. The following special characters are allowed: ~!@#%^*_-=+?\$()& Minimum length: 8 characters Maximum length: 32 characters
databases	Yes	Array of strings	All databases that the account has operation permissions for. Specify at least one database or set this parameter to all , indicating that all databases are selected.
privilege	Yes	String	Account permissions. The value can be: <ul style="list-style-type: none"> ReadOnly, indicating that the account has read-only permissions. ReadWrite, indicating that the account has read and write permissions.

Response Parameters

Status code: 200

Table 5-429 Response body parameters

Parameter	Type	Description
job_id	String	Task ID.

Example Requests

- **URI example**
POST `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/0549b4a43100d4f32f51c01c2fe4acdb/redis/054e292c9880d4992f02c0196d3ein12/db-users`
- **Creating 2 database accounts (Set `name` of the first account to `redis1`, `password` to `****`, `privilege` to `ReadOnly`, and `databases` to [`"1"`, `"2"`].**

```
{
  "users": [ {
    "name": "redis1",
    "password": "****",
    "privilege": "ReadOnly",
    "databases": [ "1", "2" ]
  }, {
    "name": "redis2",
    "password": "****",
    "privilege": "ReadOnly",
    "databases": [ "1", "2" ]
  } ]
}
```

Example Responses

Status code: 200

Success

```
{
  "job_id": "f85104b5-4a9c-4e0f-9505-fc5409d8f7ae"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.7.2 Changing Permissions for a Database Account

Function

This API is used to change permissions for a GeminiDB Redis database account.

Constraints

- Only the GeminiDB Redis API is supported.

- This operation cannot be performed when the instance is in any of the following states: creating, changing specifications, changing database port, frozen, or abnormal.

URI

PUT https://{Endpoint}/v3/{project_id}/redis/instances/{instance_id}/db-users/privilege

Table 5-430 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-431 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-432 Request body parameters

Parameter	Mandatory	Type	Description
users	No	Array of objects	Request body for modifying permissions of a database account. For details, see Table 5-433 .

Table 5-433 ModifyDbUserPrivilegeRequestBody

Parameter	Mandatory	Type	Description
name	Yes	String	Account name.

Parameter	Mandatory	Type	Description
privilege	Yes	String	Account permissions. The value can be: <ul style="list-style-type: none"> • ReadOnly, indicating that the account has read-only permissions. • ReadWrite, indicating that the account has read and write permissions.
databases	No	Array of strings	All databases that the account has access permissions for. If this parameter is not transferred, the databases remain unchanged.

Response Parameters

Status code: 202

Table 5-434 Response body parameters

Parameter	Type	Description
job_id	String	Task ID.

Example Requests

- URI example
PUT <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/054e292c9880d4992f02c0196d3ea468/redis/instances/3d39c18788b54a919bab633874c159dfin12/db-users/privilege>
- Modifying permissions of the two database accounts (Set **privilege** of user **test1** to **ReadOnly**, **databases** to ["1", "2"], **privilege** of user **test2** to **ReadWrite**, and **databases** to ["3", "4"].)

```
{
  "users": [ {
    "name": "test1",
    "privilege": "ReadOnly",
    "databases": [ "1", "2" ]
  }, {
    "name": "test2",
    "privilege": "ReadWrite",
    "databases": [ "3", "4" ]
  } ]
}
```

Example Responses

Status code: 202

Accepted

```
{
  "job_id" : "f85104b5-4a9c-4e0f-9505-fc5409d8f7ae"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.7.3 Resetting the Password of a Database Account

Function

This API is used to reset the password for a GeminiDB Redis database account.

Constraints

- Only the GeminiDB Redis API is supported.
- This operation cannot be performed when the instance is in any of the following states: creating, changing specifications, changing database port, frozen, or abnormal.

URI

PUT https://{Endpoint}/v3/{project_id}/redis/instances/{instance_id}/db-users/
password

Table 5-435 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-436 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-437 Request body parameters

Parameter	Mandatory	Type	Description
name	Yes	String	Account name.
password	Yes	String	Password to be reset. The password: <ul style="list-style-type: none"> • Can contain 8 to 32 characters. • Must contain at least two of the following types: uppercase letters, lowercase letters, digits, and special characters. The following special characters are allowed: ~!@#%^*-_+=? \$()&

Response Parameters

Status code: 204

No response parameters

Example Requests

- URI example
 PUT https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/054e292c9880d4992f02c0196d3ea468/redis/instances/3d39c18788b54a919bab633874c159dfin12/db-users/password
- Resetting the password of user **db_user1**

```
{
  "name": "db_user1",
  "password": "*****"
}
```

Example Responses

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.7.4 Deleting a Database Account

Function

This API is used to delete a database account of a GeminiDB Redis instance.

Constraints

- Only the GeminiDB Redis API is supported.
- This operation cannot be performed when the instance is in any of the following states: creating, changing specifications, changing database port, frozen, or abnormal.

URI

DELETE https://{Endpoint}/v3/{project_id}/redis/instances/{instance_id}/db-users

Table 5-438 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-439 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-440 Request body parameters

Parameter	Mandatory	Type	Description
names	Yes	Array of strings	All database accounts to be deleted.

Response Parameters

Status code: 200

Table 5-441 Response body parameters

Parameter	Type	Description
job_id	String	Task ID.

Example Requests

- URI example
DELETE https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/redis/instances/9136fd2a9fcd405ea4674276ce36dae8in12/db-users
- Example request body

```
{
  "names": [ "test1", "test2" ]
}
```

Example Responses

Status code: 200

Success

```
{
  "job_id": "f85104b5-4a9c-4e0f-9505-fc5409d8f7ae"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.7.5 Obtaining the Database Account List

Function

This API is used to obtain the GeminiDB Redis database account list.

Constraints

- Only the GeminiDB Redis API is supported.
- This operation cannot be performed when the instance is in any of the following states: creating, changing specifications, changing database port, frozen, or abnormal.

URI

GET https://{Endpoint}/v3/{project_id}/redis/instances/{instance_id}/db-users

Table 5-442 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Table 5-443 Query parameters

Parameter	Mandatory	Type	Description
name	No	String	Database account name. If this parameter is transferred, information about a specified account is queried. Otherwise, information about all database accounts is returned.
offset	No	Integer	Index offset. <ul style="list-style-type: none"> The query starts from the next piece of data indexed by this parameter. The value is 0 by default. The value must be a positive integer.
limit	No	Integer	Maximum records to be queried. <ul style="list-style-type: none"> The value ranges from 1 to 100. If this parameter is not transferred, the first 100 records are queried by default.

Request Parameters

Table 5-444 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-445 Response body parameters

Parameter	Type	Description
users	Array of objects	All database accounts. For details, see Table 5-446 .
total_count	Integer	Total number of records.

Table 5-446 DbUserInfo

Parameter	Type	Description
name	String	Account name.
type	String	Account type. The value can be: <ul style="list-style-type: none"> • rwuser, indicating that the account is an administrator account • acluser, indicating that the account is a common account
privilege	String	Account permissions. The value can be: <ul style="list-style-type: none"> • ReadOnly, indicating that the account has read-only permissions • ReadWrite, indicating that the account has read and write permissions
databases	Array of strings	All databases that the account has access permissions for.

Example Requests

URI example

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/redis/instances/9136fd2a9fcd405ea4674276ce36dae8in12/db-users
```

Example Responses

Status code: 200

Success

```
{
  "users": [ {
    "name": "db_user1",
    "type": "rwuser",
    "privilege": "ReadWrite",
    "databases": [ "1", "2" ]
  } ]
}
```

```

}, {
  "name" : "db_user2",
  "type" : "acluser",
  "privilege" : "ReadWrite",
  "databases" : [ "1", "2" ]
}],
"total_count" : 2
}

```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.7.6 Obtaining All Databases in an Instance

Function

This API is used to obtain the database list for a GeminiDB Redis instance.

Constraints

- Only the GeminiDB Redis API is supported.
- This operation cannot be performed when the instance is in any of the following states: creating, changing specifications, changing database port, frozen, or abnormal.
- Only databases of instances that contain data are returned.

URI

GET https://{Endpoint}/v3/{project_id}/redis/instances/{instance_id}/databases

Table 5-447 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Table 5-448 Query parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Index offset <ul style="list-style-type: none"> The query starts from the next piece of data indexed by this parameter. The value is 0 by default. The value must be a positive integer.
limit	No	Integer	Maximum records to be queried <ul style="list-style-type: none"> The value ranges from 1 to 100. If this parameter is not transferred, the first 100 records are queried by default.

Request Parameters

Table 5-449 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-450 Response body parameters

Parameter	Type	Description
databases	Array of strings	All databases in the GeminiDB Redis instance.
total_count	Integer	Total number of records.

Example Requests

URI example

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/
/v3/375d8d8fad1f43039e23d3b6c0f60a19/redis/instances/9136fd2a9fcd405ea4674276ce36dae8in12/
databases
```

Example Responses

Status code: 200

Success

```
{
  "databases": [ "1", "2", "3" ],
  "total_count": 3
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.7.7 Performing Operations on a GeminiDB Instance Database

Function

This function allows you to performing operations on a GeminiDB instance database.

Constraints

This API supports GeminiDB Redis instances.

URI

PUT https://{Endpoint}/v3/{project_id}/instances/{instance_id}/databases

Table 5-451 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID

Request Parameters

Table 5-452 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-453 Request body parameters

Parameter	Mandatory	Type	Description
action	Yes	String	Operation on the instance: Flush : Delete data.
db_id	No	Integer	ID of a database to be cleared. This parameter is valid only when action is set to Flush .

Response Parameters

Status code: 200

Table 5-454 Response body parameters

Parameter	Type	Description
job_id	String	Task ID

Status code: 400

Table 5-455 Response body parameters

Parameter	Type	Description
error_code	String	Error code
error_msg	String	Error message

Status code: 500

Table 5-456 Response body parameters

Parameter	Type	Description
error_code	String	Error code
error_msg	String	Error message

Example Requests

- Clearing data in a specified database

```
PUT https://{endpoint}/v3/0549b4a43100d4f32f51c01c2fe4acdb/instances/e73893ef73754465a8bd2e0857bbf13ein12/databases
```

```
{
  "action": "flush",
  "db_id": 1
}
```

- **Clearing all data**

PUT <https://{{endpoint}}/v3/0549b4a43100d4f32f51c01c2fe4acdb/instances/e73893ef73754465a8bd2e0857bbf13ein12/databases>

```
{
  "action": "flush"
}
```

Example Responses

Status code: 200

Accepted.

```
{
  "job_id": "04efe8e2-9255-44ae-a98b-d87cae411890"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.8 Tags

5.8.1 Querying an Instance by Tag

Function

This API is used to query a specified instance by tag.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

A maximum of 20 tags can be added to a DB instance. The tag key must be unique.

URI

POST https://{{Endpoint}}/v3/{{project_id}}/instances/resource-instances/action

Table 5-457 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-458 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-459 Request body parameters

Parameter	Mandatory	Type	Description
offset	No	String	Index offset. The query starts from the next piece of data indexed by this parameter. <ul style="list-style-type: none"> If action is set to count, this parameter does not need to be transferred. If action is set to filter, the parameter value must be a positive integer. The default value is 0, indicating that the query starts from the first piece of data. '
limit	No	String	Number of records to be queried. <ul style="list-style-type: none"> If action is set to count, this parameter does not need to be transferred. If action is set to filter, the value ranges from 1 to 100. If this parameter is not transferred, the first 100 instances are queried by default.

Parameter	Mandatory	Type	Description
action	Yes	String	Operation identifier. <ul style="list-style-type: none"> If action is set to filter, instances are queried based on tag filters. If action is set to count, only the total number of records is returned.
matches	No	Array of MatchOption objects	Search parameter. <ul style="list-style-type: none"> If this parameter is not specified, the query is not based on the instance name or ID. This parameter cannot be left blank.
tags	No	Array of TagOption objects	Included tags. Each tag contains a maximum of 20 keys.

Table 5-460 MatchOption

Parameter	Mandatory	Type	Description
key	Yes	String	Query criteria. The value can be instance_name or instance_id , indicating that the query is based on the instance name or instance ID.
value	Yes	String	Name or ID of the instance to be queried

Table 5-461 TagOption

Parameter	Mandatory	Type	Description
key	Yes	String	Tag key. It can contain a maximum of 36 Unicode characters. The key value cannot be null, an empty string, or spaces. Before using key , delete spaces before and after the value. NOTE The character set of this parameter is not verified during search.
values	Yes	Array of strings	Tag values. Each tag value can contain a maximum of 43 Unicode characters and cannot contain spaces. Before using values , delete spaces before and after the value. If the values is not specified, any parameter value can be queried. All values are in the OR relationship.

Response Parameters

Status code: 200

Table 5-462 Response body parameters

Parameter	Type	Description
instances	Array of InstanceResult objects	All instances.
total_count	Integer	Total number of records.

Table 5-463 InstanceResult

Parameter	Type	Description
instance_id	String	Instance ID.
instance_name	String	Instance name.

Parameter	Type	Description
tags	Array of InstanceTagResult objects	All tags. If there are no tags, tags is taken as an empty array by default.

Table 5-464 InstanceTagResult

Parameter	Type	Description
key	String	Tag key. The tag key must be specified and can include a maximum of 36 Unicode characters. It is case-sensitive and can contain digits, letters, underscores (_), and hyphens (-).
value	String	Tag value. The tag value can contain a maximum of 43 Unicode characters and can be an empty string. It is case-sensitive and can contain digits, letters, underscores (_), periods (.), and hyphens (-).

Example Requests

- URI example
POST <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/resource-instances/action>

- Example request body
Querying an instance by name (Set **offset** to **100** and **limit** to **100**.)

```
{
  "offset": 100,
  "limit": 100,
  "action": "filter",
  "matches": [{
    "key": "instance_name",
    "value": "test-single"
  }],
  "tags": [{
    "key": "key1",
    "values": [ "value1", "value2" ]
  }]
}
```

Querying total records

```
{
  "action": "count",
  "tags": [ {
    "key": "key1",
    "values": [ "value1", "value2" ]
  }, {
    "key": "key2",
    "values": [ "value1", "value2" ]
  } ],
  "matches": [ {
    "key": "instance_name",
```

```
"value" : "test-single"
}, {
  "key" : "instance_id",
  "value" : "958693039f284d6ebfb177375711072ein06"
}]
}
```

Example Responses

Status code: 200

Success

```
{
  "total_count": 1,
  "instances": [{
    "instance_id": "2acbf2223caf3bac3c33c6153423c3ccin06",
    "instance_name": "test-single",
    "tags": [ {
      "key": "key1",
      "value": "value1"
    }, {
      "key": "key2",
      "value": "value1"
    } ]
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.8.2 Adding or Deleting Resource Tags in Batches

Function

This API is used to add tags to or delete tags from a specified DB instance in batches.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

A maximum of 20 tags can be added to an instance. The tag key must be unique.

If the request body contains duplicated keys, an error message will be reported when the API is called.

If the key in the request body is the same as an existing key in a specified instance, the value of the **value** parameter that corresponds to the existing key is overwritten.

If the tag to be deleted does not exist, the system deems the deletion operation successful by default but does not check whether the tag key and value meets character set rules.

URI

POST https://{Endpoint}/v3/{project_id}/instances/{instance_id}/tags/action

Table 5-465 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-466 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-467 Request body parameters

Parameter	Mandatory	Type	Description
action	Yes	String	Operation identifier. The value can be: <ul style="list-style-type: none"> create, indicating that tags are added. delete, indicating that tags are deleted.
tags	Yes	Array of BatchTagActionTagOption objects	All tags.

Table 5-468 BatchTagActionTagOption

Parameter	Mandatory	Type	Description
key	Yes	String	<p>Tag key. It can contain a maximum of 128 Unicode characters. The key value cannot be null, an empty string, or spaces. Before using key, delete spaces before and after the value.</p> <p>It is case-sensitive and can contain digits, letters, underscores (_), and hyphens (-).</p>
value	No	String	<p>Tag value. The tag value can contain a maximum of 255 Unicode characters and can be an empty string.</p> <p>It is case-sensitive and can contain digits, letters, underscores (_), periods (.), and hyphens (-).</p> <ul style="list-style-type: none"> • If action is set to create, this parameter is mandatory. • If action is set to delete, this parameter is optional. <p>NOTE If value is specified, tags are deleted by key and value. If value is not specified, tags are deleted by key.</p>

Response Parameters

Status code: 200

No response parameters

Example Requests

- **URI example**
POST <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/9136fd2a9fcd405ea4674276ce36dae8in02/tags/action>
- **Adding two tags**

```
{
  "action": "create",
  "tags": [ {
    "key": "key1",
    "value": "value1"
  }, {
```

```
"key" : "key2",  
"value" : "value2"  
}]  
}
```

- Deleting two tags

```
{  
  "action" : "delete",  
  "tags" : [{  
    "key" : "key1"  
  }, {  
    "key" : "key2",  
    "value" : "value3"  
  }]  
}
```

Example Responses

Status code: 200

Success

```
{}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.8.3 Querying Tags of an Instance

Function

This API is used to query tags of a specified instance.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

A maximum of 20 tags can be added to a DB instance. The tag key must be unique.

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/tags

Table 5-469 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>
instance_id	Yes	String	<p>Explanation: Instance ID</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Request Parameters

Table 5-470 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token</p> <p>You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Response Parameters

Status code: 200

Table 5-471 Response body parameters

Parameter	Type	Description
tags	Array of ListInstanceTagsResult objects	Explanation: Tags of the instance.

Table 5-472 ListInstanceTagsResult

Parameter	Type	Description
key	String	Explanation: Tag key. The tag key can contain a maximum of 36 Unicode characters and must be specified. Values: It is case-sensitive and can contain digits, letters, underscores (_), and hyphens (-).
value	String	Explanation: Tag value. The tag value can contain a maximum of 43 Unicode characters and can be an empty string. Values: It is case-sensitive and can contain digits, letters, underscores (_), periods (.), and hyphens (-).

Example Request

URI example

GET <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/9136fd2a9fcd405ea4674276ce36dae8in02/tags>

Example Response

Status code: 200

Success

```
{
  "tags": [ {
    "key": "key1",
    "value": "value1"
  }, {
    "key": "key2",
```

```
"value" : "value2"
} ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.8.4 Querying Tags of a Specified Project

Function

This API is used to query tags of a specified project.

URI

GET https://{Endpoint}/v3/{project_id}/tags

Table 5-473 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Table 5-474 Query parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Index offset. <ul style="list-style-type: none"> The query starts from the next piece of data indexed by this parameter. The value is 0 by default. The value must be a positive integer.
limit	No	Integer	Maximum records to be queried. <ul style="list-style-type: none"> The value ranges from 1 to 100. If this parameter is not transferred, the first 100 records are queried by default.

Request Parameters

Table 5-475 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-476 Response body parameters

Parameter	Type	Description
tags	Array of Tag objects	All tags.
total_count	Integer	Total number of records.

Table 5-477 Tag

Parameter	Type	Description
type	String	Tag type. The value can be: <ul style="list-style-type: none"> • user • system
key	String	Tag key. The tag key must be specified and can include a maximum of 36 Unicode characters. The key is case-sensitive and can contain digits, uppercase letters, lowercase letters, underscores (_), and hyphens (-).
values	Array of strings	Tag values. The value can include a maximum of 43 Unicode characters and can also be an empty string. The value is case-sensitive and can contain digits, uppercase letters, lowercase letters, underscores (_), periods (.), and hyphens (-).

Example Requests

URI example

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/0549b4a43100d4f32f51c01c2fe4acdb/tags?offset=1&limit=10
```

Example Responses

Status code: 200

Success

```
{
  "tags": [ {
    "key": "key1",
    "values": [ "value1", "value2" ],
    "type": "user"
  }, {
    "key": "key2",
    "values": [ "value1", "value2" ],
    "type": "system"
  } ],
  "total_count": 2
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.9 Logs

5.9.1 Querying Database Slow Logs

Function

This API is used to query the latest 2,000 slow query logs of an instance. Searching by keyword is not supported.

Constraints

- This API supports GeminiDB Cassandra instances.
- This API can be used to query only the latest 2000 slow query logs in a specified time range.

URI

```
GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/slowlog
```

Table 5-478 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>
instance_id	Yes	String	Instance ID

Table 5-479 Query parameters

Parameter	Mandatory	Type	Description
start_date	Yes	String	<p>Explanation: Start time in the <i>yyyy-mm-ddThh:mm:ssZ</i> format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset. The start time is at most 30 days earlier than the current time.</p>
end_date	Yes	String	<p>Explanation: End time in the <i>yyyy-mm-ddThh:mm:ssZ</i> format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset. You can only query slow query logs generated in the last one month. The end time cannot be later than the current time.</p>

Parameter	Mandatory	Type	Description
offset	No	Integer	<p>Explanation: Index offset. Its value ranges from 0 to 1999.</p> <p>If offset is set to N, the resource query starts from the $N+1$ piece of data. The value is 0 by default, indicating that the query starts from the first piece of data. The value cannot be a negative number.</p>
limit	No	Integer	<p>Explanation: Number of records to be queried. The value ranges from 1 to 100.</p> <p>The sum of values of limit and offset must be 2000 or lower.</p>
node_id	No	String	<p>Explanation: Node ID. If this parameter is not specified, all nodes of the instance are queried.</p> <p>For details about the value, see id in Table 5-47.</p>
type	No	String	<p>Explanation: SQL statement type. If this parameter is not specified, all types of SQL statements are queried. You can also specify the following log type:</p> <ul style="list-style-type: none"> • SELECT

Request Parameters

Table 5-480 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Response Parameters

Status code: 200

Table 5-481 Response body parameters

Parameter	Type	Description
slow_log_list	Array of SlowlogResult objects	<p>Explanation: Information about slow query logs.</p>
total_record	Integer	<p>Explanation: Total number of records.</p>

Table 5-482 SlowlogResult

Parameter	Type	Description
time	String	<p>Explanation: Execution time.</p>
database	String	<p>Explanation: Database which slow query logs belong to.</p>

Parameter	Type	Description
query_sample	String	Explanation: Execution syntax.
type	String	Explanation: SQL statement type.
start_time	String	Explanation: UTC time when logs are generated.

Example Request

- URI example

Querying slow query logs:

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/0483b6b16e954cb88930a360d2c4e663/instances/6ade8143870047b8999aba8f1891b48ein06/slowlog?start_date=2018-08-06T10:41:14+0800&end_date=2018-08-07T10:41:14+0800
```

- URI example

Querying slow query logs based on specified conditions:

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/0549b4a43100d4f32f51c01c2fe4acdb/instances/6ade8143870047b8999aba8f1891b48ein06/slowlog?type=SELECT&offset=1&limit=20&node_id=a7c84462483642798cf159237343135fno06&start_date=2018-08-06T10:41:14+0800&end_date=2018-08-07T10:41:14+0800
```

Example Response

Status code: 200

Success

```
{
  "total_record" : 1,
  "slow_log_list" : [ {
    "time" : "513 ms",
    "database" : "cassandra",
    "query_sample" : "SELECT * FROM cassandra.sz_user LIMIT 100;",
    "type" : "SELECT",
    "start_time" : "2020-11-15T22:49:38.643000Z"
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.9.2 Querying Slow Query Logs of a GeminiDB Redis Instance

Function

This API is used to query slow query logs of GeminiDB Redis instances. Searching for slow query logs by keyword is supported.

Constraints

- Operators &, ||, AND, OR, NOT, *, ?, :, >, <, =, ≥, and ≤ cannot be used to search for logs.
- The query must be within the period specified by **start_time** and **end_time**.
- The value of **line_num** must be obtained from the log information returned for the last query. This parameter is optional for the first query.
- A maximum of 100 records can be queried on each page, and slow query logs can be stored for a maximum of 30 days.

URI

POST https://{Endpoint}/v3/{project_id}/redis/instances/{instance_id}/slow-logs

Table 5-483 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>
instance_id	Yes	String	<p>Explanation: Instance ID</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Request Parameters

Table 5-484 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Table 5-485 Request body parameters

Parameter	Mandatory	Type	Description
start_time	Yes	String	<p>Explanation: Start time</p> <p>Constraints: The start time cannot be 30 days earlier than the current time.</p> <p>Values: Format: <i>yyyy-mm-ddThh:mm:ssZ</i> T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.</p> <p>Default value: N/A</p>

Parameter	Mandatory	Type	Description
end_time	Yes	String	<p>Explanation: End time</p> <p>Constraints: The end time cannot be later than the current time.</p> <p>Values: Format: <i>yyyy-mm-ddThh:mm:ssZ</i> T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.</p> <p>Default value: N/A</p>
limit	Yes	Integer	<p>Explanation: Number of logs queried each time</p> <p>Constraints: N/A</p> <p>Values: 1 to 100</p> <p>Default value: N/A</p>
line_num	No	String	<p>Explanation: Sequence number of a single log line. The current query starts from the next log of line_num, excluding the log of line_num.</p> <p>Constraints: This parameter is not required for the first query. The value can be obtained from the response of the last query.</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Parameter	Mandatory	Type	Description
operate_type	No	String	<p>Explanation: Statement type</p> <p>Constraints: N/A</p> <p>Values: If this parameter is left blank, all statement types are queried.</p> <p>The following statement types can be queried: SET, GET, DEL, INCR, INCRBY, INCRBYFLOAT, DECR, DECRBY, GETSET, APPEND, MGET, KEYS, SETNX, SETEX, PSETEX, DELVX, MSET, MSETNX, GETRANGE, SUBSTR, SETRANGE, STRLEN, EXISTS, EXPIRE, PEXPIRE, EXPIREAT, PEXPIREAT, TTL, PTTL, PERSIST, TYPE, SCANX, PKSETEXAT, SORT, HDEL, HSET, HGET, HGETALL, HEXISTS, HINCRBY, HINCRBYFLOAT, HKEYS, HLEN, HMGET, HMSET, HSETNX, HSTRLEN, HVALS, HSCAN, HSCANX, PKHSCANRANGE, PKHRSCANRANGE, LINDEX, LINSERT, LLEN, LPOP, LPUSH, LPUSHX, LRANGE, LREM, LSET, LTRIM, RPOP, RPOPLPUSH, RPUSH, RPUSHX, ZADD, ZCARD, ZSCAN, ZINCRBY, ZRANGE, ZREVRANGE, ZRANGEBYSCORE, ZREVRANGEBYSCORE, ZCOUNT, ZREM, ZUNIONSTORE, ZINTERSTORE, ZRANK, ZREVRANK, ZSCORE, ZRANGEBYLEX, ZREVRANGEBYLEX, ZLEXCOUNT, ZREMRANGEBYRANK, ZREMRANGEBYSCORE, ZREMRANGEBYLEX, ZPOPMAX, ZPOPMIN, SADD, SPOP, SCARD, SMEMBERS, SSCAN, SREM, SUNION, SUNIONSTORE, SINTER, SINTERSTORE, SISMEMBER,</p>

Parameter	Mandatory	Type	Description
			SDIFF, SDIFFSTORE, SMOVE, SRANDMEMBER, BITSET, BITGET, BITCOUNT, BITPOS, BITOP, BITFIELD, PFADD, PFCOUNT, PFMERGE, GEOADD, GEORADIUSBYMEMBER, GEORADIUS, GEOHASH, GEODIST, GEOPOS, XADD, XACK, XGROUP, XDEL, XTRIM, XLEN, XRANGE, XREVRANGE, XCLAIM, XPENDING, XINFO, XREAD, and XREADGROUP Default value: N/A
node_id	No	String	Explanation: Node ID Constraints: N/A Values: If this parameter is left blank, all nodes in the instance can be queried. For details about the value, see field id in Table 5-47 . Default value: N/A
keywords	No	Array of strings	Explanation: Full-text log search based on multiple keywords, indicating that all keywords are matched. Constraints: <ul style="list-style-type: none"> • A maximum of 10 keywords are supported. • Each keyword can contain a maximum of 512 characters.

Parameter	Mandatory	Type	Description
max_cost_time	No	Double	<p>Explanation: Logs can be searched based on the maximum execution duration. The unit is millisecond.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>
min_cost_time	No	Double	<p>Explanation: Logs can be searched based on the minimum execution duration. The unit is millisecond.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Response Parameters

Status code: 200

Table 5-486 Response body parameters

Parameter	Type	Description
slow_logs	Array of objects	<p>Explanation: Slow log information. For details, see Table 5-487.</p>

Table 5-487 RedisSlowLogDetail

Parameter	Type	Description
node_name	String	Explanation: Node name Values: N/A
node_id	String	Explanation: Node ID Values: N/A
whole_message	String	Explanation: Statement Values: N/A
operate_type	String	Explanation: Statement type Values: N/A
cost_time	Double	Explanation: Execution time (ms) Values: N/A
log_time	String	Explanation: UTC time when a log is generated Values: The format is <i>yyyy-mm-ddThh:mm:ssZ</i> . T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.
line_num	String	Explanation: Sequence number of a log event. Values: N/A

Example Request

```
POST
https://{Endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/redis/instances/
a6d3c8a9857b4c81b3c1fe4802dfa4d0in12/slow-logs
{
```

```
"start_time" : "2022-09-06T10:41:14+0800",
"end_time" : "2022-09-16T10:41:14+0800",
"limit" : 100,
"line_num" : "1595659490239433658",
"operate_type" : "set",
"node_id" : "2997329fe3cb4b3faedcade16df6966eno12",
"keywords" : [ "log", "test" ],
"max_cost_time" : 100.12,
"min_cost_time" : 50.12
}
```

Example Response

Status code: 200

Success

```
{
  "slow_logs": [{
    "node_name": "test_worker_node_1",
    "node_id": "dbd0b65ed0c34125a0b04d4e5ba67e66no02",
    "whole_message": "testDb",
    "operate_type": "SET",
    "cost_time": 60.12,
    "log_time": "2022-09-15T22:49:38.643000Z",
    "line_num": "1595659490239433659"
  }]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.9.3 Querying Slow Query Logs of a GeminiDB Influx Instance

Function

This API is used to query slow query logs of GeminiDB Influx instances. Searching for slow query logs by keyword is supported.

Constraints

- Operators &, ||, AND, OR, NOT, *, ?, :, >, <, =, ≥, and ≤ cannot be used to search for logs.
- The query must be within the period specified by **start_time** and **end_time**.
- The value of **line_num** must be obtained from the log information returned for the last query. This parameter is optional for the first query.
- A maximum of 100 records can be queried on each page, and slow query logs can be stored for a maximum of 30 days.

URI

POST https://{Endpoint}/v3/{project_id}/influxdb/instances/{instance_id}/slow-logs

Table 5-488 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-489 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-490 Request body parameters

Parameter	Mandatory	Type	Description
start_time	Yes	String	Start time in the yyyy-mm-ddThh:mm:ssZ format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset. The start time cannot be 30 days earlier than the current time.
end_time	Yes	String	End time in the yyyy-mm-ddThh:mm:ssZ format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset. The end time cannot be later than the current time.
limit	Yes	Integer	Number of logs to be queried each time. The maximum value is 100 .

Parameter	Mandatory	Type	Description
line_num	No	String	Sequence number of a log event. This parameter is not required for the first query, but is required for the next query. The value can be obtained from the response of the last query. The current query starts from the next log of line_num , excluding the log of line_num .
operate_type	No	String	Statement type. If this parameter is left empty, all statement types are queried. Value options: <ul style="list-style-type: none"> • SELECT • DELETE • SHOW • DROP • CREATE • ALTER
node_id	No	String	Node ID. If this parameter is left blank, all nodes in the instance can be queried. For details about the value, see field id in Table 5-47 .
keywords	No	String	Full-text log search based on multiple keywords, indicating that all keywords are matched. A maximum of 10 keywords are supported. Each keyword can contain up to 512 characters.
max_cost_time	No	Integer	Logs can be searched based on the maximum execution duration. Unit: ms
min_cost_time	No	Integer	Logs can be searched based on the minimum execution duration. Unit: ms
database	No	String	Database name.
retention_policy	No	String	Retention policy.

Response Parameters

Status code: 200

Table 5-491 Response body parameters

Parameter	Type	Description
slow_logs	Array of objects	GeminiDB Influx slow query logs. For details, see Table 5-492 .

Table 5-492 InfluxdbSlowLogDetail

Parameter	Type	Description
node_id	String	Node ID.
node_name	String	Node name.
whole_message	String	Statement.
operate_type	String	SQL statement type.
cost_time	String	Execution time. Unit: ms
log_time	String	UTC time when a log is generated. The format is yyyy-mm-ddThh:mm:ssZ. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.
database	String	Database name.
retention_policy	String	Retention policy.
line_num	String	Sequence number of a log event.

Example Requests

Querying 100 slow query logs that consist of SELECT statements from 2023-07-06 10:41:14 to 2023-07-10 10:41:14.

```
POST https://{Endpoint}/v3/2900b7b8d03e4619b8db8d43bc6234ee/influxdb/instances/4466bd981011413cb98514e1cc20fdb0in13/slow-logs
```

```
{
  "start_time": "2023-07-06T10:41:14+0800",
  "end_time": "2023-07-10T10:41:14+0800",
  "limit": 100,
  "line_num": "1595659490239433658",
  "operate_type": "SELECT",
  "node_id": "8b76c35e91eb4c9d82fe25417d750cb5no13",
  "max_cost_time": 100,
  "min_cost_time": 50,
  "database": "dbname",
```

```
"retention_policy" : "monitor"
}
```

Example Responses

Status code: 200

Success

```
{
  "slow_logs" : [ {
    "node_name" : "test_priam_node_1",
    "node_id" : "8b76c35e91eb4c9d82fe25417d750cb5no13",
    "whole_message" : "SELECT tag,field FROM dbname..mst WHERE time >= '1' AND time < '2'",
    "operate_type" : "SELECT",
    "cost_time" : 60,
    "log_time" : "2023-07-04T05:42:32.000+0000",
    "database" : "dbname",
    "line_num" : "1595659490239433659",
    "retention_policy" : "monitor"
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.9.4 Querying Slow Query Logs of a GeminiDB Cassandra Instance

Function

This API is used to query slow query logs of GeminiDB Cassandra instances. Searching for slow query logs by keyword is supported.

Constraints

- The following operators are supported: &&, ||, AND, OR, NOT, *, ?, :, >, <, =, ≥, and ≤
- The query must be within the period specified by **start_time** and **end_time**.
- The value of **line_num** must be obtained from the log information returned for the last query. This parameter is optional for the first query.
- A maximum of 100 records can be queried on each page, and slow query logs can be stored for a maximum of 30 days.

URI

POST https://{Endpoint}/v3/{project_id}/cassandra/instances/{instance_id}/slow-logs

Table 5-493 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-494 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-495 Request body parameters

Parameter	Mandatory	Type	Description
start_time	Yes	String	Start time in the yyyy-mm-ddThh:mm:ssZ format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset. The start time cannot be 30 days earlier than the current time.
end_time	Yes	String	End time in the yyyy-mm-ddThh:mm:ssZ format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset. The end time cannot be later than the current time.
limit	Yes	Integer	Number of logs to be queried each time. The maximum value is 100 .

Parameter	Mandatory	Type	Description
line_num	No	String	Sequence number of a log event. This parameter is not required for the first query, but is required for the next query. The value can be obtained from the response of the last query. The current query starts from the next log of line_num , excluding the log of line_num .
operate_type	No	String	Statement type. If this parameter is left empty, all statement types are queried. The value can be select .
node_id	No	String	Node ID. If this parameter is left blank, all nodes in the instance can be queried. For details about the value, see field id in Table 5-47 .
keywords	No	Array of strings	Full-text log search based on multiple keywords, indicating that all keywords are matched. <ul style="list-style-type: none"> • A maximum of 10 keywords are supported. • Each keyword can contain a maximum of 512 characters.
keyspace_key words	No	Array of strings	Fuzzy search for logs based on multiple keyspace keywords, indicating that at least one keyword is matched. <ul style="list-style-type: none"> • Only fuzzy search by keyword prefix is supported. A maximum of 10 keywords are supported. • Each keyword can contain a maximum of 48 characters.

Parameter	Mandatory	Type	Description
table_keywords	No	Array of strings	Fuzzy search for logs based on multiple database table name keywords, indicating that at least one keyword is matched. <ul style="list-style-type: none"> Only fuzzy search by keyword prefix is supported. A maximum of 10 keywords are supported. Each keyword can contain a maximum of 48 characters.
max_cost_time	No	Integer	Logs can be searched based on the maximum execution duration. Unit: ms
min_cost_time	No	Integer	Logs can be searched based on the minimum execution duration. Unit: ms

Response Parameters

Status code: 200

Table 5-496 Response body parameters

Parameter	Type	Description
slow_logs	Array of objects	Slow log information. For details, see Table 5-497 .

Table 5-497 CassandraSlowLogDetail

Parameter	Type	Description
node_name	String	Node name.
node_id	String	Node ID.
whole_message	String	Statement.
operate_type	String	Statement type.
cost_time	Integer	Execution time. Unit: ms
keyspace	String	Database keyspace.
table	String	Table name.

Parameter	Type	Description
log_time	String	UTC time when a log is generated. The format is yyyy-mm-ddThh:mm:ssZ. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.
line_num	String	Sequence number of a log event.

Example Requests

```
POST https://{Endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/cassandra/instances/853f97101a9a4f618202f281cda82e92in06/slow-logs
```

```
{
  "start_time" : "2023-01-06T10:41:14+0800",
  "end_time" : "2023-01-10T10:41:14+0800",
  "limit" : 100,
  "line_num" : "1595659490239433658",
  "operate_type" : "select",
  "node_id" : "8b76c35e91eb4c9d82fe25417d750cb5no06",
  "keywords" : [ "log", "test" ],
  "keyspace_keywords" : [ "system", "system_auth" ],
  "table_keywords" : [ "test" ],
  "max_cost_time" : 100,
  "min_cost_time" : 50
}
```

Example Responses

Status code: 200

Success

```
{
  "slow_logs" : [ {
    "node_name" : "test_priam_node_1",
    "node_id" : "8b76c35e91eb4c9d82fe25417d750cb5no06",
    "keyspace" : "system",
    "table" : "test",
    "whole_message" : "GqKc6_S",
    "operate_type" : "select",
    "cost_time" : 60,
    "log_time" : "2023-01-09T05:42:32.000+0000",
    "line_num" : "1595659490239433659"
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.9.5 Querying Slow Query Logs of a GeminiDB Mongo Instance

Function

This API is used to query slow query logs of GeminiDB Mongo instances. Searching for slow query logs by keyword is supported.

Constraints

- The following operators are not supported: &&, ||, AND, OR, NOT, *, ?, :, >, <, =, ≥, and ≤
- The query must be within the period specified by **start_time** and **end_time**.
- The value of **line_num** must be obtained from the log information returned for the last query. This parameter is optional for the first query.
- A maximum of 100 records can be queried on each page, and slow query logs can be stored for a maximum of 30 days.

URI

POST https://{Endpoint}/v3/{project_id}/mongodb/instances/{instance_id}/slow-logs

Table 5-498 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-499 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-500 Request body parameters

Parameter	Mandatory	Type	Description
start_time	Yes	String	Start time in the yyyy-mm-ddThh:mm:ssZ format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset. The start time cannot be 30 days earlier than the current time.
end_time	Yes	String	End time in the yyyy-mm-ddThh:mm:ssZ format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset. The end time cannot be later than the current time.
limit	Yes	Integer	Number of logs to be queried each time. The maximum value is 100 .
line_num	No	String	Sequence number of a log event. This parameter is not required for the first query, but is required for the next query. The value can be obtained from the response of the last query. The current query starts from the next log of line_num , excluding the log of line_num .
operate_type	No	String	Statement type. If this parameter is left empty, all statement types are queried. Value options: <ul style="list-style-type: none"> • insert • query • update • remove • getmore • command • killcursors

Parameter	Mandatory	Type	Description
node_id	No	String	Node ID. If this parameter is left blank, all nodes in the instance can be queried. For details about the value, see field id in Table 5-47 .
keywords	No	Array of strings	Full-text log search based on multiple keywords, indicating that all keywords are matched. <ul style="list-style-type: none"> • A maximum of 10 keywords are supported. • Each keyword can contain a maximum of 512 characters.
database_key words	No	Array of strings	Fuzzy search for logs based on multiple database name keywords, indicating that at least one keyword is matched. <ul style="list-style-type: none"> • Only fuzzy search by keyword prefix is supported. A maximum of 10 keywords are supported. • Each keyword can contain a maximum of 64 characters.
collection_key words	No	Array of strings	Fuzzy search for logs based on multiple database collection name keywords, indicating that at least one keyword is matched. <ul style="list-style-type: none"> • Only fuzzy search by keyword prefix is supported. A maximum of 10 keywords are supported. • Each keyword can contain a maximum of 128 characters.
max_cost_time	No	Integer	Logs can be searched based on the maximum execution duration. Unit: ms
min_cost_time	No	Integer	Logs can be searched based on the minimum execution duration. Unit: ms

Response Parameters

Status code: 200

Table 5-501 Response body parameters

Parameter	Type	Description
slow_logs	Array of objects	Slow log information. For details, see Table 5-502 .

Table 5-502 MongodbSlowLogDetail

Parameter	Type	Description
node_name	String	Node name.
node_id	String	Node ID.
whole_message	String	Statement.
operate_type	String	Statement type.
cost_time	Integer	Execution time. Unit: ms
lock_time	Integer	Lock wait time. Unit: μ s
docs_returned	Integer	Number of documents returned by a slow query.
docs_scanned	Integer	Number of documents scanned during a slow query.
database	String	Database name.
collection	String	Database collection name.
log_time	String	UTC time when a log is generated. The format is yyyy-mm-ddThh:mm:ssZ. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.
line_num	String	Sequence number of a log event.

Example Requests

```
POST https://{Endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/mongodb/instances/916de02c35fe4abaa3e707ebd916d3f8in10/slow-logs
```

```
{
  "start_time" : "2021-01-06T10:41:14+0800",
  "end_time" : "2023-01-10T10:41:14+0800",
  "limit" : 100,
  "line_num" : "1595659490239433658",
  "operate_type" : "query",
```

```
"node_id" : "dfe41edcb2a1483c96ddf61cc0cee237no10",
"keywords" : [ "response", "oplog" ],
"database_keywords" : [ "testDB", "system_auth" ],
"collection_keywords" : [ "testCollection" ],
"max_cost_time" : 100,
"min_cost_time" : 50
}
```

Example Responses

Status code: 200

Success

```
{
  "slow_logs" : [ {
    "node_name" : "test_replica_node_1",
    "node_id" : "dfe41edcb2a1483c96ddf61cc0cee237no10",
    "database" : "testDB",
    "collection" : "testCollection",
    "whole_message" : "{\"responseLength\": 230, \"ts\": {\"$date\": 1605480486800}, \"ninserted\": 1, \"locks\": {\"oplog\": {\"acquireCount\": {\"w\": 1}}, \"Global\": {\"acquireCount\": {\"r\": 3, \"w\": 2}}, \"Collection\": {\"acquireCount\": {\"w\": 2}}, \"Database\": {\"acquireCount\": {\"w\": 3}}}, \"numYield\": 0, \"ns\": \"geographySpace.tiles\"\",
    \"operate_type\" : \"query\",
    \"cost_time\" : 60,
    \"lock_time\" : 10,
    \"docs_returned\" : 0,
    \"docs_scanned\" : 0,
    \"log_time\" : \"2023-01-09T06:42:32.000+0000\",
    \"line_num\" : \"1595659490239433659\"
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.9.6 Querying Database Error Logs

Function

This API is used to query error logs of an instance. Searching for error logs by keyword is not supported.

Constraints

This API supports GeminiDB Mongo instances.

The latest 2,000 error logs can be queried.

The past one month of error logs can be queried.

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/error-log

Table 5-503 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID, which can be obtained by calling the API described in Querying Instances and Details . If there are no instances available, call the API described in Creating an Instance to create one.

Table 5-504 Query parameters

Parameter	Mandatory	Type	Description
start_time	Yes	String	Start time in the yyyy-mm-ddThh:mm:ssZ format. T is the separator between the calendar and the hourly notation of time. Z indicates the time zone offset. The start time is at most 30 days earlier than the current time.
end_time	Yes	String	End time in the yyyy-mm-ddThh:mm:ssZ format. T is the separator between the calendar and the hourly notation of time. Z indicates the time zone offset. Only error logs generated in the past month can be queried. The end time cannot be later than the current time.
node_id	No	String	Node ID. If this parameter is left blank, all nodes in the instance can be queried.

Parameter	Mandatory	Type	Description
type	No	String	Statement type. If this parameter is left empty, all statement types are queried. The following log types are supported: <ul style="list-style-type: none"> • Warning • Error
offset	No	Integer	Index offset. <ul style="list-style-type: none"> • Its value ranges from 0 to 1999. If offset is set to N, the resource query starts from the N+1 piece of data. If action is set to filter, offset is 0 by default, indicating that the query starts from the first piece of data. • The value must be a positive integer.
limit	No	Integer	Number of records to be queried. <ul style="list-style-type: none"> • The value ranges from 1 to 100. The default value is 10, indicating that 10 records are returned by default. • The sum of limit and offset values must be less than or equal to 2000.

Request Parameters

Table 5-505 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 202

Table 5-506 Response body parameters

Parameter	Type	Description
total_count	Integer	Total number of records.
error_log_list	Array of Table 5-507 objects	Error log information.

Table 5-507 ErrorLogList

Parameter	Type	Description
node_name	String	Node name.
level	String	Log level.
time	String	UTC time when logs are generated.
content	String	Log content.

Example Requests

- URI example
GET <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/9136fd2a9fcd405ea4674276ce36dae8in10/error-log>
- Example request body
None

Example Responses

Status code: 202

Success

```
{
  "error_log_list": [ {
    "node_name": "Test_replica_node_2",
    "level": "WARNING",
    "time": "2020-12-15T08:53:01.868+0000",
    "content": "W NETWORK [LogicalSessionCacheReap] Unable to reach primary for set replica"
  } ],
  "total_count": 1
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.9.7 Querying Error Logs of a GeminiDB Mongo Instance

Function

This API is used to query error logs of GeminiDB Mongo instances. Searching for error logs by keyword is supported.

Constraints

- The following operators are not supported: &&, ||, AND, OR, NOT, *, ?, :, >, <, =, ≥, and ≤
- The query must be within the period specified by **start_time** and **end_time**.
- The value of **line_num** must be obtained from the log information returned for the last query. This parameter is optional for the first query.
- A maximum of 100 records can be queried on each page, and error logs can be stored for a maximum of 30 days.

URI

POST https://{Endpoint}/v3/{project_id}/mongodb/instances/{instance_id}/error-logs

Table 5-508 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-509 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-510 Request body parameters

Parameter	Mandatory	Type	Description
start_time	Yes	String	Start time in the yyyy-mm-ddThh:mm:ssZ format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset. The start time cannot be 30 days earlier than the current time.
end_time	Yes	String	End time in the yyyy-mm-ddThh:mm:ssZ format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset. The end time cannot be later than the current time.
limit	Yes	Integer	Number of logs to be queried each time. The maximum value is 100 .
line_num	No	String	Sequence number of a log event. This parameter is not required for the first query, but is required for the next query. The value can be obtained from the response of the last query. The current query starts from the next log of line_num , excluding the log of line_num .
severity	No	String	Log level. If this parameter is left blank, logs of all levels can be queried. Value options: <ul style="list-style-type: none"> ● Warning ● Error
node_id	No	String	Node ID. If this parameter is left blank, all nodes in the instance can be queried. For details about the value, see field id in Table 5-47 .

Parameter	Mandatory	Type	Description
keywords	No	Array of strings	Full-text log search based on multiple keywords, indicating that all keywords are matched. <ul style="list-style-type: none"> • A maximum of 10 keywords are supported. • Each keyword can contain a maximum of 512 characters.

Response Parameters

Status code: 200

Table 5-511 Response body parameters

Parameter	Type	Description
error_logs	Array of objects	Error log details. For details, see Table 5-512 .

Table 5-512 MongoDBErrorLogDetail

Parameter	Type	Description
node_name	String	Node name.
node_id	String	Node ID.
raw_message	String	Error description.
severity	String	Log level.
log_time	String	UTC time when a log is generated. The format is yyyy-mm-ddThh:mm:ssZ. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.
line_num	String	Sequence number of a log event.

Example Requests

```
POST https://{Endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/mongodb/instances/916de02c35fe4abaa3e707ebd916d3f8in10/error-logs
{
  "start_time" : "2023-01-06T10:41:14+0800",
  "end_time" : "2023-01-10T10:41:14+0800",
  "limit" : 100,
  "line_num" : "1595659490239433658",
```

```
"severity" : "Warning",
"node_id" : "dfe41edcb2a1483c96ddf61cc0cee237no10",
"keywords" : [ "error" ]
}
```

Example Responses

Status code: 200

Success

```
{
  "error_logs" : [ {
    "node_name" : "test_replica_node_1",
    "node_id" : "dfe41edcb2a1483c96ddf61cc0cee237no10",
    "raw_message" : "W TEST [test001] test error log",
    "severity" : "Warning",
    "log_time" : "2023-01-09T12:55:39.398+0000",
    "line_num" : "1595659490239433659"
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.9.8 Setting the Desensitization Status of Slow Query Logs

Function

This API is used to set the desensitization status of slow query logs.

Constraints

This API supports the following types of instances:

- GeminiDB Mongo

After desensitization is disabled, it cannot be enabled.

URI

PUT https://{Endpoint}/v3/{project_id}/instances/{instance_id}/slowlog-desensitization

Table 5-513 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Parameter	Mandatory	Type	Description
instance_id	Yes	String	Instance ID, which can be obtained by calling the API described in Querying Instances and Details . If there are no instances available, call the API described in Creating an Instance to create one.

Request Parameters

Table 5-514 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-515 Request body parameters

Parameter	Mandatory	Type	Description
desensitization_status	Yes	String	Whether slow query log desensitization is enabled. The value can be: <ul style="list-style-type: none"> off, indicating that desensitization is disabled.

Response Parameters

Status code: 204

No response parameters

Example Requests

- **URI example**
PUT `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/9136fd2a9fcd405ea4674276ce36dae8in10/slowlog-desensitization`
- **Disabling data masking**

```
{
  "desensitization_status": "off"
}
```

Example Responses

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.9.9 Querying the Desensitization Status of Slow Query Logs

Function

This API is used to query the desensitization status of slow query logs.

Constraints

This API supports GeminiDB Mongo instances.

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/slowlog-desensitization

Table 5-516 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-517 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-518 Response body parameters

Parameter	Type	Description
desensitization_status	String	Whether slow query log desensitization is enabled. The value can be: <ul style="list-style-type: none"> on, indicating that desensitization is enabled. off, indicating that desensitization is disabled.

Example Requests

- URI example
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/0549b4a43100d4f32f51c01c2fe4acdb/instances/094424666ef04f79a2dfbe9f5b8b31a5in06/slowlog-desensitization
- Example request body
None

Example Responses

Status code: 200

Success

```
{
  "desensitization_status": "on"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.9.10 Associating Instances with an LTS Log Stream

Function

- After you have associated instances with a Log Tank Service (LTS) log stream, logs of these instances are automatically uploaded to the associated LTS log stream.
- You will be billed for log reporting. See LTS pricing details.
- After a specific log stream is selected, the system creates structuring configurations of the required log type for it. If there are already structuring configurations of another log type in the log stream, they will be overwritten.

Constraints

- Only slow query logs of GeminiDB Redis and GeminiDB MongoDB instances and audit logs of GeminiDB Redis instances can be configured.
- This function is only available to instances of the latest kernel version.
- Logs of GeminiDB Redis instances in a Redis cluster cannot be audited.
- This operation cannot be performed on instances in the creating, deleted, or frozen state.
- A maximum of 100 instances can be associated with an LTS log stream in a batch.

URI

POST https://{Endpoint}/v3/{project_id}/instances/logs/lts-configs

Table 5-519 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-520 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.
X-Language	No	String	Language.

Table 5-521 Request body parameters

Parameter	Mandatory	Type	Description
instance_ids	Yes	Array of strings	IDs of the instances to be associated with an LTS log stream.
log_type	Yes	String	Log type. The value can be: <ul style="list-style-type: none"> • slow_log, indicating that the log is a slow query log. • audit_log, indicating that the log is an audit log.

Parameter	Mandatory	Type	Description
lts_group_id	Yes	String	LTS log group ID. You can obtain the value using the LTS API for querying all log groups under an account.
lts_stream_id	Yes	String	LTS log stream ID. You can obtain the value using the LTS API for querying all log streams in a specified log group.

Response Parameters

Status code: 200

No response parameters

Example Requests

- URI example
POST `https://{Endpoint}/v3/0549b4a43100d4f32f51c01c2fe4acdb/instances/logs/lts-configs`
- Associating instances with an LTS log stream that supports slow query logs

```
{
  "instance_ids": [ "fe45815ce8a6413ab7fa290da1db1614in12",
                    "362b42f769ff45d884e1866f8ee00bb4in12" ],
  "log_type": "slow_log",
  "lts_group_id": "65442dd4-b8de-4e9f-8a80-b60aa0698354",
  "lts_stream_id": "5992dd29-fd11-4228-a852-5f6e578789c5"
}
```
- Associating instances with an LTS log stream that supports audit logs

```
{
  "instance_ids": [ "fe45815ce8a6413ab7fa290da1db1614in12",
                    "362b42f769ff45d884e1866f8ee00bb4in12" ],
  "log_type": "audit_log",
  "lts_group_id": "65442dd4-b8de-4e9f-8a80-b60aa0698354",
  "lts_stream_id": "5992dd29-fd11-4228-a852-5f6e578789c5"
}
```

Example Responses

Status code: 200

Success

```
{}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.9.11 Disassociating Instances from an LTS Log Stream

Function

After you have disassociated instances from an LTS log stream, logs of these instances are not automatically uploaded to the associated LTS log stream.

Constraints

- This function supports only slow query logs and audit logs of GeminiDB Redis instances.

URI

DELETE https://{Endpoint}/v3/{project_id}/instances/logs/lts-configs

Table 5-522 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-523 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.
X-Language	No	String	Language.

Table 5-524 Request body parameters

Parameter	Mandatory	Type	Description
instance_ids	Yes	Array of strings	IDs of the instances to be disassociated from an LTS log stream.
log_type	Yes	String	Log type. The value can be: <ul style="list-style-type: none"> slow_log, indicating that the log is a slow query log. audit_log, indicating that the log is an audit log.

Response Parameters

Status code: 200

No response parameters

Example Requests

- **URI example**
DELETE https://{Endpoint}/v3/0549b4a43100d4f32f51c01c2fe4acdb/instances/logs/lts-configs
- **Disassociating instances from an LTS log stream that supports slow query logs**

```
{
  "instance_ids": [ "fe45815ce8a6413ab7fa290da1db1614in12",
    "362b42f769ff45d884e1866f8ee00bb4in12" ],
  "log_type": "slow_log"
}
```
- **Disassociating instances from an LTS log stream that supports audit logs**

```
{
  "instance_ids": [ "fe45815ce8a6413ab7fa290da1db1614in12",
    "362b42f769ff45d884e1866f8ee00bb4in12" ],
  "log_type": "audit_log"
}
```

Example Responses

Status code: 200

Success

```
{}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.9.12 Querying LTS Log Configurations

Function

This API is used to query the LTS configuration of instances.

Constraints

- This function supports only slow query logs and audit logs of GeminiDB Redis instances.

URI

GET https://{Endpoint}/v3/{project_id}/instances/logs/lts-configs

Table 5-525 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID. To obtain this value, see Obtaining a Project ID .

Table 5-526 Query parameters

Parameter	Mandatory	Type	Description
offset	No	Integer	Index offset. The query starts from the next piece of data indexed by this parameter. The offset value must be a positive integer. The default value is 0 , indicating that the query starts from the first piece of data.
limit	No	Integer	Maximum returned records. The value must be an integer ranging from 1 to 100 . If this parameter is not transferred, information of 100 instances is queried by default.
instance_id	No	String	Instance ID used for exact search.
instance_name	No	String	Instance name used for fuzzy search.
enterprise_project_id	No	String	Enterprise project ID for exact search.

Request Parameters

Table 5-527 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.
X-Language	No	String	Language.

Response Parameters

Status code: 200

Table 5-528 Response body parameters

Parameter	Type	Description
total_count	Integer	Total instance records.
instance_lts_configs	Array of objects	LTS log configurations of the instance. For details, see Table 5-529 .

Table 5-529 InstanceLogConfig

Parameter	Type	Description
instance	object	Instance information. For details, see Table 5-530 .
lts_configs	Array of objects	LTS log configuration details. If the LTS log stream is not configured, this parameter is not returned. For details, see Table 5-532 .

Table 5-530 LogInstanceInfo

Parameter	Type	Description
id	String	Instance ID.
name	String	Instance name.
status	String	Instance status. The value can be: <ul style="list-style-type: none"> • normal, indicating that the instance is running normally. • abnormal, indicating that the instance is abnormal. • creating, indicating that the instance is being created. • frozen, indicating that the instance is frozen. • data_disk_full, indicating that the instance disk is full. • createfail, indicating that the instance failed to be created. • enlargefail, indicating that nodes failed to be added to the instance.

Parameter	Type	Description
mode	String	Instance type. The value can be: <ul style="list-style-type: none"> • RedisCluster: GeminiDB Redis instance in proxy and Redis clusters • RedisReplication: primary/standby GeminiDB Redis instance • ReplicaSet: GeminiDB MongoDB instance in a replica set
datastore	object	Database information. For details, see Table 5-531 .
actions	Array of strings	Operation that is executed on the instance.
enterprise_project_id	String	Enterprise project ID.
supported_log_types	Array of strings	Supported log type. The value can be: <ul style="list-style-type: none"> • slow_log, indicating that the log is a slow query log. • audit_log, indicating that the log is an audit log.

Table 5-531 InstancesDatastoreResult

Parameter	Type	Description
type	String	DB API.
version	String	DB version number.

Table 5-532 InstanceLogConfigDetail

Parameter	Type	Description
log_type	String	Log type. The value can be: <ul style="list-style-type: none"> • slow_log, indicating that the log is a slow query log. • audit_log, indicating that the log is an audit log.
lts_group_id	String	ID of the LTS log group. If enabled is set to false , this parameter indicates the ID of the latest associated LTS log group.

Parameter	Type	Description
lts_stream_id	String	ID of the LTS log stream. If enabled is set to false , this parameter indicates the ID of the latest associated LTS log stream.
enabled	Boolean	Whether LTS log stream is enabled. <ul style="list-style-type: none"> ● true: enabled ● false: disabled

Example Requests

```
GET https://{Endpoint}/v3/0549b4a43100d4f32f51c01c2fe4acdb/instances/logs/lts-configs?
limit=10&offset=0&instance_id=362b42f769ff45d884e1866f8ee00bb4in12&instance_name=c316&enterprise_
project_id=0
```

Example Responses

Status code: 200

Success

```
{
  "total_count" : 1,
  "instance_lts_configs" : [ {
    "instance" : {
      "id" : "362b42f769ff45d884e1866f8ee00bb4in12",
      "name" : "nosql-c316",
      "mode" : "RedisCluster",
      "datastore" : {
        "version" : "5.0",
        "type" : "redis"
      },
      "status" : "normal",
      "actions" : [ "GROWING" ],
      "enterprise_project_id" : "0",
      "supported_log_types" : [ "slow_log", "audit_log" ]
    },
    "lts_configs" : [ {
      "log_type" : "slow_log",
      "lts_group_id" : "65442dd4-b8de-4e9f-8a80-b60aa0698354",
      "lts_stream_id" : "5992dd29-fd11-4228-a852-5f6e578789c5",
      "enabled" : true
    } ]
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.10 Quotas

5.10.1 Querying Quota

Function

This API is used to query GeminiDB resource quotas of a tenant.

URI

GET https://{Endpoint}/v3/{project_id}/quotas

Table 5-533 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	<p>Explanation: Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Request Parameters

Table 5-534 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	<p>Explanation: User token You can obtain the token by calling the IAM API by following Obtaining a User Token Through Password Authentication.</p> <p>Constraints: N/A</p> <p>Values: N/A</p> <p>Default value: N/A</p>

Response Parameters

Status code: 200

Table 5-535 Response body parameters

Parameter	Type	Description
quotas	ShowResourcesListResponseBody object	Explanation: Quota information.

Table 5-536 ShowResourcesListResponseBody

Parameter	Type	Description
resources	Array of ShowResourcesDetailResponseBody objects	Explanation: All resources.

Table 5-537 ShowResourcesDetailResponseBody

Parameter	Type	Description
type	String	Explanation: Quota resource type Constraints: This parameter is configured only by instance type. Values: N/A
quota	Integer	Explanation: Current quota. If this parameter is set to 0 , no quantity limit is set for resources. Values: N/A
used	Integer	Explanation: Number of used resources. Values: N/A

Example Requests

URI example

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/0549b4a43100d4f32f51c01c2fe4acdb/quotas
```

Example Responses

Status code: 200

Success

```
{
  "quotas" : {
    "resources" : [ {
      "type" : "instance",
      "quota" : 200,
      "used" : 58
    } ]
  }
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.11 Disaster Recovery

5.11.1 Querying Regions Where a Dual-Active Relationship Can Be Created Between Two Instances

Function

This API is used to query the regions where a dual-active relationship can be created between two instances.

Constraints

This API can be used on GeminiDB Cassandra instances.

URI

```
GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/disaster-recovery/regions
```

Table 5-538 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-539 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

Status code: 200

Table 5-540 Response body parameters

Parameter	Type	Description
region_codes	Array of strings	Region where a dual-active relationship can be created between two instances. The creation of a dual-active instance is contingent on network configurations and necessary conditions. A dual-active relationship is not necessarily created between two instances even though there are available regions as such.

Example Request

URI example

```
GET https://{Endpoint}/v3/2900b7b8d03e4619b8db8d43bc6234ee/instances/3149aee486d748f68db1ee81e95b9f56in06/disaster-recovery/regions
```

Example Response

Status code: 200

Success.

```
{  
  "region_codes" : [ "cn-north-4" ]  
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.11.2 Setting the Percentage of Faulty Nodes to Be Taken Over

Function

This API is used to set the percentage of faulty nodes to be taken over in case of a failover.

Constraints

This API can be used only on GeminiDB Cassandra instances.

URI

PUT https://{Endpoint}/v3/{project_id}/instances/disaster-recovery/settings

Table 5-541 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-542 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-543 Request body parameters

Parameter	Mandatory	Type	Description
disaster_recovery_settings	No	Array of SwitchoverRatio objects	Faulty nodes to be taken over

Table 5-544 SwitchoverRatio

Parameter	Mandatory	Type	Description
instance_id	Yes	String	Instance ID
switchover_ratio	No	Integer	Percentage of faulty nodes to be taken over. The value ranges from 50% to 100% . 10% nodes can be added each time. The default value is 100% .

Response Parameters

Status code: 200

Table 5-545 Response body parameters

Parameter	Type	Description
successed_instance_ids	Array of strings	Instances for which the percentage of faulty nodes to be taken over is successfully set

Status code: 400

Table 5-546 Response body parameters

Parameter	Type	Description
error_code	String	Error code
error_msg	String	Error message

Status code: 500

Table 5-547 Response body parameters

Parameter	Type	Description
error_code	String	Error code
error_msg	String	Error message

Example Request

Setting the percentage of faulty nodes to be taken over in case of a failover

```
PUT https://{endpoint}/v3/0549b4a43100d4f32f51c01c2fe4acdb/instances/disaster-recovery/settings
{
  "disaster_recovery_settings": [ {
    "instance_id": "d74162b942604190b991d1c39772a66ain06",
    "switchover_ratio": 50
  }, {
    "instance_id": "e74162b942604190b991d1c39772a66ain06",
    "switchover_ratio": 100
  } ]
}
```

Example Response

Status code: 200

Success.

```
{
  "succeeded_instance_ids": [ "d74162b942604190b991d1c39772a66ain06",
    "e74162b942604190b991d1c39772a66ain06" ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.11.3 Querying the Percentage of Faulty Nodes to Be Taken Over

Function

This API is used to query the percentage of faulty nodes to be taken over in case of a failover.

Constraints

This API can be used only on GeminiDB Cassandra instances.

URI

GET https://{Endpoint}/v3/{project_id}/instances/disaster-recovery/settings

Table 5-548 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Table 5-549 Query parameters

Parameter	Mandatory	Type	Description
instance_id	No	String	Instance ID
offset	No	Integer	Index position. If offset is set to N, the resource query starts from the N+1 piece of data. The value must be greater than or equal to 0. If this parameter is not transferred, the value is 0 by default.
limit	No	Integer	Maximum number of instances that can be queried. The value ranges from 1 to 50. If this parameter is not transferred, 50 instances are queried by default.

Request Parameters

Table 5-550 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

Status code: 200

Table 5-551 Response body parameters

Parameter	Type	Description
disaster_recovery_settings	Array of SwitchoverRatio objects	Faulty nodes to be taken over
total_count	Integer	Total number of records

Table 5-552 SwitchoverRatio

Parameter	Type	Description
instance_id	String	Instance ID
switchover_ratio	Integer	Percentage of faulty nodes to be taken over. The value ranges from 50 to 100 . 10 nodes can be added each time. The default value is 100 .

Status code: 400

Table 5-553 Response body parameters

Parameter	Type	Description
error_code	String	Error Codes
error_msg	String	Error message

Status code: 500

Table 5-554 Response body parameters

Parameter	Type	Description
error_code	String	Error code
error_msg	String	Error message

Example Request

- Querying the percentage of faulty nodes to be taken over on 50 instances in case of a failover
GET [https://\[endpoint\]/v3/0549b4a43100d4f32f51c01c2fe4acdb/instances/disaster-recovery/settings](https://[endpoint]/v3/0549b4a43100d4f32f51c01c2fe4acdb/instances/disaster-recovery/settings)
- Querying the percentage of faulty nodes to be taken over in case of a failover

```
GET https://{endpoint}/v3/0549b4a43100d4f32f51c01c2fe4acdb/instances/disaster-recovery/settings?offset=0&limit=50&instance_id=d74162b942604190b991d1c39772a66ain06&instance_id=e74162b942604190b991d1c39772a66ain06
```

Example Response

Status code: 200

Success.

```
{
  "total_count": 2,
  "disaster_recovery_settings": [ {
    "instance_id": "d74162b942604190b991d1c39772a66ain06",
    "switchover_ratio": 50
  }, {
    "instance_id": "e74162b942604190b991d1c39772a66ain06",
    "switchover_ratio": 100
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.12 Task Management

5.12.1 Querying Tasks and Details

Function

This API is used to query tasks (by default) and details.

URI

GET https://{Endpoint}/v3/{project_id}/jobs

Table 5-555 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Table 5-556 Query parameters

Parameter	Mandatory	Type	Description
id	No	String	Task ID.
start_time	No	String	Query start time in the "yyyy-mm-ddThh:mm:ssZ" format. The default value is 30 days before the current date. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.
end_time	No	String	End time (the current time by default) in the "yyyy-mm-ddThh:mm:ssZ" format. It must be later than the start time and the time span cannot exceed 30 days. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.
status	No	String	Task status. The value can be: Running , indicating that the task is being executed. Completed , indicating that the task is completed. Failed , indicating that the task fails.

Parameter	Mandatory	Type	Description
name	No	String	<p>Task name.</p> <ul style="list-style-type: none"> • CreateInstance: Create an instance. • RestoreNewInstance: Restore data to a new instance. • EnlargeInstance: Add nodes. • ReduceInstance: Delete nodes. • RestartInstance: Restart an instance. • RestartNode: Restart a node. • EnlargeInstanceVolume: Scale up storage space of an instance. • ReduceInstanceVolume: Scale in storage space of an instance. • ResizeInstance: Change the specifications of an instance. • UpgradeDbVersion: Upgrade the engine version. • BindPublicIP: Bind an EIP to an instance. • UnbindPublicIP: Unbind an EIP from an instance. • DeleteInstance: Delete an instance. • EnlargeInstanceColdVolume: Scale up cold storage of an instance. • AddInstanceColdVolume: Enable cold storage for an instance. • ModifySecurityGroup: Modify a security group. • ModifyCcmCert: Modify a CCM certificate. • ModifyPort: Change a port.

Parameter	Mandatory	Type	Description
			<ul style="list-style-type: none"> • ConstructDisasterRecovery: Establish a DR relationship. • DeConstructDisasterRecovery: Remove a DR relationship. • SwitchOverDisasterRecovery: Switch a DR relationship. • BuildBiActiveInstance: Create an instance with a dual-active DR relationship. • ReleaseBiActiveInstance: Remove a dual-active relationship from an instance. • BackupInstance: Back up an instance.
offset	No	Integer	Index offset. If offset is set to N , the resource query starts from the $N+1$ piece of data. The value is 0 by default, indicating that the query starts from the first piece of data. The value cannot be a negative number.
limit	No	Integer	Number of records to be queried. The value can be 10 , 20 , or 50 . The default value is 50 .

Request Parameters

Table 5-557 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-558 Response body parameters

Parameter	Type	Description
jobs	Array of objects (Table 5-559)	Task list.
total_count	Integer	Total number of tasks in the task list.

Table 5-559 JobDetail

Parameter	Type	Description
id	String	Task ID.
name	String	Task name.
status	String	Task execution status. The value can be: Running , indicating that the task is being executed. Completed indicating that the task has been successfully executed. Failed indicating that the task fails to be executed.
start_time	String	Creation time in the "yyyy-mm-ddThh:mm:ssZ" format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset. For example, in the Beijing time zone, the offset is +0800 .
end_time	String	End time in the yyyy-mm-ddThh:mm:ssZ format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset. For example, in the Beijing time zone, the offset is +0800 .
progress	String	Task execution progress. NOTE The execution progress (such as " 60% ", indicating the task execution progress is 60%) is displayed only when the task is being executed. Otherwise, "" is returned.
instance	Objects in Table 5-560	Details of the instance associated with the task.
fail_reason	String	Task failure information.

Table 5-560 JobInstanceInfo

Parameter	Type	Description
id	String	Instance ID.
name	String	Instance name.

Example Requests

- URI example
GET https://{endpoint}/v3/0549b4a43100d4f32f51c01c2fe4acdb/jobs?id=89a0cde6-9c46-4b89-a92c-573e1083ff23

Example Responses

Status code: 200

Success

```
{
  "total_count": 1,
  "jobs": [ {
    "id": "6f85e061-04dd-42e7-86d6-d3b1e40aac2e",
    "name": "CreateCassandra",
    "status": "Running",
    "start_time": "2023-09-12T06:44:01+0000",
    "end_time": "2023-09-12T06:44:03+0000",
    "progress": "14%",
    "instance": {
      "id": "27a045b6bf9e46f691f81366d398cb04in06",
      "name": "nosql-12f5"
    },
    "fail_reason": ""
  } ]
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.12.2 Querying the Maintenance Period of an Instance

Function

This function is used to query the maintenance period of an instance.

Constraints

This API supports the following types of instances:

- GeminiDB Influx

- GeminiDB Redis
- GeminiDB Mongo
- GeminiDB Cassandra

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/ops-window

Table 5-561 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a user in a region
instance_id	Yes	String	Instance ID

Request Parameters

Table 5-562 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

Status code: 200

Table 5-563 Response body parameters

Parameter	Type	Description
maintenance_window	String	Maintenance time window of an instance

Example Request

Query the maintenance period of an instance.

```
GET
https://{Endpoint}/v3/619d3e78f61b4be68bc5aa0b59edcf7b/instances/
e73893ef73754465a8bd2e0857bbf13ein02/ops-window
```

Example Response

Status code: 200

Success

```
{
  "maintenance_window" : "02:00-06:00"
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.12.3 Querying Scheduled Tasks

Function

This function is used to query scheduled tasks based on specified conditions.

Constraints

This API supports the following types of instances:

- GeminiDB Influx
- GeminiDB Redis
- GeminiDB Mongo
- GeminiDB Cassandra

URI

GET https://{Endpoint}/v3/{project_id}/scheduled-jobs

Table 5-564 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a user in a region

Table 5-565 Query parameters

Parameter	Mandatory	Type	Description
job_name	No	String	Task name. This parameter is left blank by default. The value can be: <ul style="list-style-type: none"> • REBOOT: Restart the instance. • RESIZE_FLAVOR: Change the CPU or memory of an instance. • UPGRADE_DATABASE: patch upgrade
job_status	No	String	Task execution status. This parameter is left blank by default. The value can be: <ul style="list-style-type: none"> • Pending: The task is not executed. • Running: The task is being executed. • Completed: The task has been successfully executed. • Failed: The task fails to be executed. • Canceled: The task is canceled.
instance_id	No	String	Instance ID. If this parameter is not specified, all instances that meet the search criteria are queried by default.
start_time	No	String	Task creation time in the <i>yyyy-mm-ddThh:mm:ssZ</i> format. T indicates the start time, and Z indicates the time zone offset. For example, if the time zone offset is one hour, +0100 is displayed. If this parameter is not transferred, the time zone offset is seven days before the current time by default.

Parameter	Mandatory	Type	Description
end_time	No	String	End time for creating a task, in the <i>yyyy-mm-ddThh:mm:ssZ</i> format. T indicates the start time, and Z indicates the time zone offset. For example, if the time zone offset is one hour, +0100 is displayed. If this parameter is not transferred, the current time is used by default.
offset	No	Integer	Index offset. The query starts from the next piece of data indexed by this parameter. The value is 0 by default.
limit	No	Integer	Number of records to be queried. If this parameter is not transferred, the default value 10 is used. The value ranges from 1 to 100 .

Request Parameters

Table 5-566 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

Status code: 200

Table 5-567 Response body parameters

Parameter	Type	Description
total_count	Integer	Total number of records.
schedules	Array of objects	Task details. For details, see Table 5-568 .

Table 5-568 ScheduleDetailInfo

Parameter	Type	Description
job_id	String	Task ID
job_name	String	Task name. The value can be: <ul style="list-style-type: none"> ● REBOOT: Restart the instance. ● RESIZE_FLAVOR: Change the CPU or memory of an instance. ● UPGRADE_DATABASE: patch upgrade
job_status	String	Task execution status. The value can be: The value can be: <ul style="list-style-type: none"> ● Pending: The task is not executed. ● Running: The task is being executed. ● Completed: The task has been successfully executed. ● Failed: The task fails to be executed. ● Canceled: The task is canceled.
instance_id	String	Instance ID
instance_name	String	Instance name
instance_status	String	Instance status. The value can be: <ul style="list-style-type: none"> ● Createfail, indicating that the instance failed to be created. ● Creating, indicating that the instance is being created. ● Normal: The instance is normal. ● Abnormal: The instance is abnormal. ● Deleted: The instance has been deleted.
datastore_type	String	Database type
create_time	String	Task creation time in the <i>yyyy-mm-ddThh:mm:ssZ</i> format. T is the separator between the calendar and the hourly notation of time. Z indicates the time zone offset. For example, if the time zone offset is one hour, the value of Z is +0100 .
start_time	String	Task start time in the <i>yyyy-mm-ddThh:mm:ssZ</i> format. T is the separator between the calendar and the hourly notation of time. Z indicates the time zone offset. For example, if the time zone offset is one hour, the value of Z is +0100 .

Parameter	Type	Description
end_time	String	Task end time in the <i>yyyy-mm-ddThh:mm:ssZ</i> format. T is the separator between the calendar and the hourly notation of time. Z indicates the time zone offset. For example, if the time zone offset is one hour, the value of Z is +0100 .

Example Request

This function is used to query scheduled tasks based on specified conditions.

```
GET https://{Endpoint}/v3/0549b4a43100d4f32f51c01c2fe4acdb/scheduled-jobs?offset=2&limit=2&start_time=2019-05-27T03:38:51+0000&job_status=Pending
```

Example Response

Status code: 200

Success

```
{
  "schedules" : [ {
    "job_id" : "56d3c1138dcf4f1da73b0170700c78d0",
    "job_name" : "REBOOT",
    "job_status" : "Pending",
    "instance_id" : "8a6117ece8a0457a90060441fa372e8ein06",
    "instance_name" : "gauss-nosql",
    "instance_status" : "normal",
    "datastore_type" : "GeminiDB Cassandra",
    "create_time" : "2019-05-26T10:41:14+0000",
    "start_time" : "2019-05-26T12:41:14+0000",
    "end_time" : "2019-05-26T12:51:14+0000"
  } ],
  "total_count" : 1
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.12.4 Canceling a Scheduled Task

Function

This function is used to cancel a scheduled task based on the task ID.

Constraints

This API supports the following types of instances:

- GeminiDB Influx
- GeminiDB Redis
- GeminiDB Mongo
- GeminiDB Cassandra

URI

DELETE https://{Endpoint}/v3/{project_id}/scheduled-jobs/{job_id}

Table 5-569 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a user in a region
job_id	Yes	String	Task ID. The value is the same as the job_id field returned in the scheduled task list.

Request Parameters

Table 5-570 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

Status code: 204

None

Example Request

Cancel a scheduled task.

```
DELETE https://{Endpoint}/v3/0549b4a43100d4f32f51c01c2fe4acdb/scheduled-jobs/56d3c1138dcf4f1da73b0170700c78d0
```

Example Response

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.13 Enterprise Projects

5.13.1 Querying Enterprise Project Quotas

Function

This API is used to query enterprise project quotas.

URI

GET https://{Endpoint}/v3/{project_id}/enterprise-projects/quotas

Table 5-571 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Table 5-572 Query parameters

Parameter	Mandatory	Type	Description
enterprise_project_name	No	String	Enterprise project name. Fuzzy search is supported. If this parameter is not specified, all enterprise project quotas are returned.
offset	No	Integer	Index offset. <ul style="list-style-type: none"> The query starts from the next piece of data indexed by this parameter. The value is 0 by default. The value must be a positive integer.

Parameter	Mandatory	Type	Description
limit	No	Integer	<p>Maximum records to be queried.</p> <ul style="list-style-type: none"> The value ranges from 1 to 100. If this parameter is not transferred, the first 100 records are queried by default.

Request Parameters

Table 5-573 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

Status code: 200

Table 5-574 Response body parameters

Parameter	Type	Description
total_count	Integer	Total number of records
quotas	objects	Enterprise project quotas. For details, see Table 5-575 .

Table 5-575 NoSqlQueryEpsQuotalInfo

Parameter	Type	Description
enterprise_project_id	String	Enterprise project ID
enterprise_project_name	String	Enterprise project name
quota	object	Enterprise project quotas. For details, see Table 5-576 .
used	object	Enterprise project quota used. For details, see Table 5-577 .

Table 5-576 NoSqlEpsQuotaTotal

Parameter	Type	Description
instance	Integer	Instance quota
vcpus	Integer	vCPU quota
ram	Integer	RAM quota

Table 5-577 NoSqlEpsQuotaUsed

Parameter	Type	Description
instance	Integer	Used instance quota
vcpus	Integer	Used vCPU quota
ram	Integer	Used RAM quota

Example Requests

- URI example
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/0549b4a43100d4f32f51c01c2fe4acdb/enterprise-projects/quotas?enterprise_project_name=test&offset=1&limit=10
- Example request body
None

Example Responses

Status code: 200

Success.

```
{
  "quotas" : [ {
    "enterprise_project_id" : "c0348bb1-d09d-4ee2-8edd-53e496fe6b52",
    "enterprise_project_name" : "test1",
    "quota" : {
      "instance" : 500,
      "vcpus" : 1000,
      "ram" : 2000
    },
    "used" : {
      "instance" : 15,
      "vcpus" : 88,
      "ram" : 256
    }
  }, {
    "enterprise_project_id" : "780a6b1f-58b8-4df6-a85e-326d052de704",
    "enterprise_project_name" : "test2",
    "quota" : {
      "instance" : 500,
      "vcpus" : 1000,
      "ram" : 2000
    },
    "used" : {
```

```

"instance" : 36,
"vcpus" : 64,
"ram" : 192
}
}],
"total_count" : 2
}
    
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.13.2 Modifying Enterprise Project Quotas

Function

This API is used to modify enterprise project quotas.

URI

PUT https://{Endpoint}/v3/{project_id}/enterprise-projects/quotas

Table 5-578 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

Table 5-579 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 5-580 Request body parameters

Parameter	Mandatory	Type	Description
quotas	Yes	objects	Enterprise quotas to be modified. For details, see Table 5-581 .

Table 5-581 NoSqlRequestEpsQuota

Parameter	Mandatory	Type	Description
enterprise_project_id	Yes	String	Enterprise project ID
quota	Yes	object	Enterprise quotas to be modified. For details, see Table 5-582 . NOTE At least one of parameters instance , vcpus , and ram must be transferred.

Table 5-582 NoSqlEpsQuotaRequestInfo

Parameter	Mandatory	Type	Description
instance	No	Integer	Instance quota
vcpus	No	Integer	vCPU quota
ram	No	Integer	RAM quota

Response Parameters

Status code: 204

No response parameters

Example Requests

- **URI example**
PUT <https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/054e292c9880d4992f02c0196d3ea468/enterprise-projects/quotas>

- **Modifying quotas of an enterprise project (Set **instance** to **1000**, **vcpus** to **500**, and **ram** to **1024**.)**

```
{
  "quotas" : [ {
    "enterprise_project_id" : "4d05638e-d4c6-477c-9b51-9620fa257a11",
    "quota" : {
      "instance" : 1000,
      "vcpus" : 500,
      "ram" : 1024
    }
  }, {
    "enterprise_project_id" : "92450d0e-8c4b-48e1-9909-4d9d2f086ce4",
    "quota" : {
      "ram" : 512
    }
  }
]
```

Example Responses

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.14 Instance Load Balancing Management

5.14.1 Configure IP Address Blacklist/Whitelist for Instance Load Balancing

Function

This API is used to configure IP Address blacklist or whitelist for instance load balancing. Each time this API is called, the previously configured value is overwritten. After this function is disabled, the source IP address of the connection is not restricted.

Constraints

This API supports the following types of instances:

- GeminiDB Redis

URI

PUT https://{Endpoint}/v3/{project_id}/instances/{instance_id}/lb/access-control

Table 5-583 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-584 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Table 5-585 Request body parameters

Parameter	Mandatory	Type	Description
enabled	Yes	Boolean	The value true indicates that this function is enabled, and false indicates that this function is disabled.
type	Yes	String	Type. The value can be: <ul style="list-style-type: none"> • whiteList, indicating that only specified IP addresses or CIDR blocks are allowed. • blackList, indicating that specified IP addresses or CIDR blocks are not allowed.
ip_groups	Yes	Array of objects	IP addresses or CIDR blocks in the IP address group. For details, see Table 5-586 .

Table 5-586 IpGroupItem

Parameter	Mandatory	Type	Description
ip	Yes	String	IP address or CIDR block. IPv4 and IPv6 addresses are supported.
description	No	String	Supplementary information. The value consists of 0 to 255 characters.

Response Parameters

Status code: 200

None

Example Requests

- Adding the CIDR block 123.123.123.0/24 to the blacklist
PUT <https://{Endpoint}/v3/054e292c9880d4992f02c0196d3ea468/instances/054e292c9880d4992f02c0196d3ea468/lb/access-control>

```
{
  "type" : "blackList",
  "enabled" : true,
  "ip_groups" : [ {
    "ip" : "123.123.123.0/24",
    "description" : "test"
  } ]
}
```

- Disabling the load balancing blacklist and whitelist
PUT <https://{Endpoint}/v3/054e292c9880d4992f02c0196d3ea468/instances/054e292c9880d4992f02c0196d3ea468/lb/access-control>

```
{
  "type" : "blackList",
  "enabled" : false,
  "ip_groups" : [ {
    "ip" : "123.123.123.123",
    "description" : "test"
  } ]
}
```

Example Responses

None

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

5.14.2 Querying IP Address Blacklist/Whitelist for Instance Load Balancing

Function

This API is used to query IP address blacklist or whitelist for instance load balancing

Constraints

This API supports the following type of instances:

- GeminiDB Redis

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/lb/access-control

Table 5-587 Path parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID.

Request Parameters

Table 5-588 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token.

Response Parameters

Status code: 200

Table 5-589 Response body parameters

Parameter	Type	Description
enabled	Boolean	The value true indicates that this function is enabled, and false indicates that this function is disabled.
type	String	Type. The value can be: <ul style="list-style-type: none"> • whiteList, indicating that specified IP addresses or CIDR blocks are allowed. • blackList, indicating that specified IP addresses or CIDR blocks are not allowed.
ip_groups	Array of objects	IP addresses or CIDR blocks in the IP address group. For details, see Table 5-590 .

Table 5-590 IpGroupItem

Parameter	Type	Description
ip	String	IP address or CIDR block. IPv4 and IPv6 addresses are supported.

Parameter	Type	Description
description	String	Supplementary information. The value consists of 0 to 255 characters.

Example Requests

```
GET https://{Endpoint}/v3/054e292c9880d4992f02c0196d3ea468/instances/  
054e292c9880d4992f02c0196d3ein12/lb/access-control
```

Example Responses

Status code: 200

Success

```
{  
  "type" : "blackList",  
  "ip_groups" : [ {  
    "ip" : "123.123.123.0/24",  
    "description" : "test"  
  } ],  
  "enabled" : true  
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

6 API v3 (Unavailable Soon)

6.1 Instance Specifications

Function

This API is used to query all instance specifications under a specified condition.

NOTICE

This API will be unavailable on March 31, 2025. You are advised to switch workloads to the new API ([Querying Instance Specifications](#)) before then.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

URI

- URI format
GET `https://{Endpoint}/v3/{project_id}/flavors?region={region}&engine_name={engine_name}`
- URI example
`https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/flavors?region=ap-southeast-1&engine_name=cassandra`
- Required parameters

Table 6-1 Parameter description

Parameter	Mandatory	Description
project_id	Yes	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
region	No	Region where the instance is deployed. The value can be: Must be specified. For details, see Regions and Endpoints .
engine_name	No	Database type. The value can be: <ul style="list-style-type: none"> • cassandra, indicating that the instances are of the GeminiDB Cassandra type. • mongodb, indicating that the instances are of the GeminiDB Mongo type. • influxdb, indicating that the instances are of the GeminiDB Influx type. • redis, indicating that the instances are of the GeminiDB Redis type. • If this parameter is not transferred, the default value is cassandra.

Request Parameters

None

Response Parameters

- Normal response

Table 6-2 Parameter description

Parameter	Type	Description
total_count	Integer	Total number of records.
flavors	Array of objects	Instance specifications. For more information, see Table 6-3 .

Table 6-3 Data structure description of parameter **flavors**

Parameter	Type	Description
engine_name	String	API name.
engine_version	String	API version.
vcpus	String	Number of vCPUs.
ram	String	Memory size in megabytes (MB).
spec_code	String	Resource specification code. Example: geminidb.cassandra.8xlarge.4 NOTE <ul style="list-style-type: none"> • geminidb.cassandra indicates that the instance of the GeminiDB Cassandra type. • 8xlarge.4 indicates node specifications.
availability_zone	Array of strings	ID of the AZ that supports the specifications. NOTE <ul style="list-style-type: none"> • This parameter has been discarded. Do not use it.
az_status	Object	Status of specifications in an AZ. The value can be: <ul style="list-style-type: none"> • normal, indicating that the specifications are on sale. • unsupported, indicating that the specifications are not supported. • sellout, indicating that the specifications are sold out.

 **NOTE**

The parameter values under **az_status** are example values and only for reference.

- Example normal response

```
{
  "total_count": 4,
  "flavors": [
    {
      "engine_name": "cassandra",
      "engine_version": "3.11",
      "vcpus": "4",
      "ram": "16",
      "spec_code": "geminidb.cassandra.xlarge.4",
      "availability_zone": [
        "az1",
        "az2"
      ],
    }
  ],
}
```

```

    "az_status": {
      "az1": "normal",
      "az2": "unsupported"
    }
  },
  {
    "engine_name": "cassandra",
    "engine_version": "3.11",
    "vcpus": "8",
    "ram": "32",
    "spec_code": "geminidb.cassandra.2xlarge.4",
    "availability_zone": [
      "az1",
      "az2"
    ],
    "az_status": {
      "az1": "unsupported",
      "az2": "normal"
    }
  },
  {
    "engine_name": "cassandra",
    "engine_version": "3.11",
    "vcpus": "16",
    "ram": "64",
    "spec_code": "geminidb.cassandra.4xlarge.4",
    "availability_zone": [
      "az1",
      "az2"
    ],
    "az_status": {
      "az1": "normal",
      "az2": "sellout"
    }
  },
  {
    "engine_name": "cassandra",
    "engine_version": "3.11",
    "vcpus": "32",
    "ram": "128",
    "spec_code": "geminidb.cassandra.8xlarge.4",
    "availability_zone": [
      "az1",
      "az2"
    ],
    "az_status": {
      "az1": "normal",
      "az2": "normal"
    }
  }
]
}

```

- Abnormal response
For details, see [Abnormal Request Results](#).

Status Codes

For more information, see [Status Codes](#).

Error Codes

For more information, see [Error Codes](#).

6.2 Backup and Restoration

6.2.1 Querying an Automated Backup Policy

Function

This API is used to query an automated backup policy, including GeminiDB Cassandra databases and tables.

NOTICE

This API will be unavailable. You are advised to switch workloads to the new API described in [Querying an Automated Backup Policy](#).

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx
- GeminiDB Redis

URI

GET https://{Endpoint}/v3/{project_id}/instances/{instance_id}/backups/policy

Table 6-4 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID

Table 6-5 Query parameters

Parameter	Mandatory	Type	Description
type	No	String	Backup policy type. This parameter is available only to GeminiDB Cassandra API. The value can be: <ul style="list-style-type: none"> • Instance: An instance backup is queried. • DatabaseTable: A database or table backup is queried. • The default value is Instance.

Request Parameters

Table 6-6 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Response Parameters

Status code: 202

Table 6-7 Response body parameters

Parameter	Type	Description
backup_policy	object	Backup policy objects, including the backup retention period (days) and start time. For details, see Table 6-8 .

Table 6-8 ShowBackupPolicyResult

Parameter	Type	Description
keep_days	Integer	Backup retention days
start_time	String	Backup time window. Automated backup will be triggered during the backup time window.

Parameter	Type	Description
period	String	Backup period. After a backup period is specified, data will be automatically backed up on the selected days every week.

Example Request

URI example

```
GET https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/054b61972980d4552f0bc00ac8d3f5cd/instances/764b8a2763d34414ad3f2d5495416cb5in06/backups/policy?type=DatabaseTable
```

Example Response

Status code: 202

Success

Enabling automated backup

```
{
  "backup_policy": {
    "keep_days": 7,
    "start_time": "19:00-20:00",
    "period": "1,2,4,5,6"
  }
}
```

Disabling automated backup

```
{
  "backup_policy": {
    "keep_days": 0
  }
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

6.3 Parameter Templates

6.3.1 Obtaining Parameter Templates

Function

This API is used to obtain parameter templates, including default and custom parameter templates of all instances.

NOTICE

This API will be unavailable on March 31, 2025. You are advised to switch workloads to the new API ([Obtaining Parameter Templates](#)) before then.

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

URI

- URI format
GET https://{Endpoint}/v3/{project_id}/configurations
- URI example
https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/configurations
- Required parameters

Table 6-9 Parameter description

Parameter	Mandatory	Description
project_id	Yes	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

None

Response Parameters

- Normal response

Table 6-10 Parameter description

Parameter	Type	Description
count	Integer	Total number of records.
configurations	Array of objects	Parameter templates For details, see Table 6-11 .

Table 6-11 Data structure description of parameter configurations

Parameter	Type	Description
id	String	Parameter template ID.
name	String	Parameter template name.
description	String	Parameter template description.
datastore_version_name	String	Database version name.
datastore_name	String	Database name
created	String	Creation time in the yyyy-MM-ddTHH:mm:ssZ format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.
updated	String	Update time in the yyyy-MM-ddTHH:mm:ssZ format. T is the separator between calendar and hourly notation of time. Z indicates the time zone offset.
user_defined	Boolean	Whether the parameter template is a custom template. The value can be: <ul style="list-style-type: none"> • false: indicates that the parameter template is a default parameter template. • true: indicates that the parameter template is a custom template.

- Example normal response

```
{
  "count": 2,
  "configurations": [{
    "id": "887ea0d1bb0843c49e8d8e5a09a95652pr06",
    "name": "configuration_test",
    "description": "configuration_test",
    "datastore_version_name": "3.11",
    "datastore_name": "cassandra",
    "created": "2019-05-15T11:53:34+0000",
    "updated": "2019-05-15T11:53:34+0000",
    "user_defined": true
  },
  {
    "id": "3bc1e9cc0d34404b9225ed7a58fb284epr06",
    "name": "Default-Cassandra-3.11",
    "description": "Default parameter group for cassandra 3.11",
```

```
"datastore_version_name": "3.11",  
"datastore_name": "cassandra",  
"created": "2020-03-21T03:38:51+0000",  
"updated": "2019-03-21T03:38:51+0000",  
"user_defined": false  
}  
]  
}
```

- Abnormal response
For details, see [Abnormal Request Results](#).

Status Codes

For more information, see [Status Codes](#).

Error Codes

For more information, see [Error Codes](#).

6.3.2 Modifying Parameters of a Specified Instance

Function

This API is used to modify parameters of a specified instance.

NOTICE

This API will be unavailable. You are advised to switch workloads to the new API described in [Modifying Parameters of a Specified Instance](#).

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

For configuration item **values**, you can enter system-defined parameters that allow for modification.

This API is an asynchronous API. A successful response does not indicate that the parameters are successfully modified.

URI

PUT https://{Endpoint}/v3/{project_id}/instances/{instance_id}/configurations

Table 6-12 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
instance_id	Yes	String	Instance ID

Request Parameters

Table 6-13 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 6-14 Request body parameters

Parameter	Mandatory	Type	Description
values	Yes	Map<String,String>	Parameter values defined by users based on a default parameter template.

Response Parameters

Status code: 200

Table 6-15 Response body parameters

Parameter	Type	Description
job_id	String	ID of an asynchronous task for modifying instance parameters
restart_required	Boolean	Whether an instance needs to be restarted. The value can be: <ul style="list-style-type: none"> true: An instance needs to be restarted. false: An instance does not need to be restarted.

Example Request

- URI example
PUT `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/054e292c9880d4992f02c0196d3ea468/instances/392850e624504e1490901d50b585a60din06/configurations`
- Modifying parameters of a specified instance

```
{  
  "values" : {  
    "request_timeout_in_ms" : "10000"  
  }  
}
```

Example Response

Status code: 200

Success

```
{  
  "job_id" : "463b4b58-d0e8-4e2b-9560-5dea4552fde9",  
  "restart_required" : false  
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

6.3.3 Applying a Parameter Template

Function

This API is used to apply a parameter template to one or more instances.

NOTICE

This API will be unavailable. You are advised to switch workloads to the new API described in [Applying a Parameter Template](#).

Constraints

This API supports the following types of instances:

- GeminiDB Cassandra
- GeminiDB Mongo
- GeminiDB Influx

This API is an asynchronous API. A successful response does not indicate that the parameter template is successfully applied.

URI

PUT https://{Endpoint}/v3/{project_id}/configurations/{config_id}/apply

Table 6-16 URI parameters

Parameter	Mandatory	Type	Description
project_id	Yes	String	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .
config_id	Yes	String	Parameter template ID

Request Parameters

Table 6-17 Request header parameters

Parameter	Mandatory	Type	Description
X-Auth-Token	Yes	String	User token

Table 6-18 Request body parameters

Parameter	Mandatory	Type	Description
instance_ids	Yes	Array of strings	Instance IDs

Response Parameters

Status code: 200

Table 6-19 Response body parameters

Parameter	Type	Description
job_id	String	ID of an asynchronous task that applies the parameter template
success	Boolean	Whether the task for applying the parameter template is successfully submitted. The value can be: <ul style="list-style-type: none"> true: The task is successfully submitted. false: The task failed to be submitted.

Example Request

- URI example
PUT `https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/configurations/e02e76567ae04662a2753492b77f965bpr06/apply`
- Applying a parameter template

```
{  
  "instance_ids": [ "73ea2bf70c73497f89ee0ad4ee008aa2in06" ]  
}
```

Example Response

Status code: 200

Success

```
{  
  "job_id": "463b4b58-d0e8-4e2b-9560-5dea4552fde9",  
  "success": true  
}
```

Status Codes

For details, see [Status Codes](#).

Error Codes

For details, see [Error Codes](#).

6.4 Tags

6.4.1 Querying an Instance by Tag

Function

This API is used to query a specified instance by tag.

NOTICE

This API will be unavailable on March 31, 2025. You are advised to switch workloads to the new API ([Querying an Instance by Tag](#)) before then.

Constraints

- This API supports the following types of instances:
 - GeminiDB Cassandra
 - GeminiDB Mongo
 - GeminiDB Influx
 - GeminiDB Redis

- A maximum of 20 tags can be added to a DB instance. The tag key must be unique.

URI

- URI format
POST https://{Endpoint}/v3/{project_id}/instances/resource_instances/action
- URI example
https://gaussdb-nosql.ap-southeast-1.myhuaweicloud.com/v3/375d8d8fad1f43039e23d3b6c0f60a19/instances/resource_instances/action
- Required parameters

Table 6-20 Parameter description

Parameter	Mandatory	Description
project_id	Yes	Project ID of a tenant in a region. To obtain this value, see Obtaining a Project ID .

Request Parameters

- Required parameters

Table 6-21 Parameter description

Parameter	Mandatory	Type	Description
offset	No	String	<p>Index position. The query starts from the next piece of data indexed by this parameter.</p> <ul style="list-style-type: none"> • If action is set to count, this parameter does not need to be transferred. • If action is set to filter, the parameter value must be a positive integer. The default value is 0, indicating that the query starts from the first piece of data.

Parameter	Mandatory	Type	Description
limit	No	String	<p>Number of records to be queried.</p> <ul style="list-style-type: none"> • If action is set to count, this parameter does not need to be transferred. • If action is set to filter, the value ranges from 1 to 100. If this parameter is not transferred, the first 100 instances are queried by default.
action	Yes	String	<p>Operation identifier.</p> <ul style="list-style-type: none"> • If action is set to filter, instances are queried based on tag filters. • If action is set to count, only the total number of records is returned.
matches	No	Array of objects	<p>Search parameter.</p> <ul style="list-style-type: none"> • If this parameter is not specified, the query is not based on the instance name or ID. • If the parameter is specified, see parameter values in Table 6-23.
tags	No	Array of objects	<p>Included tags. Each tag contains up to 20 keys. For more information, see Table 6-22.</p>

Table 6-22 Data structure description of parameter **tags**

Parameter	Mandatory	Type	Description
key	Yes	String	Tag key. It can contain a maximum of 36 Unicode characters. The key value cannot be null, an empty string, or spaces. Before using key , delete spaces before and after the value. NOTE The character set of this parameter is not verified during search.
values	Yes	Array of strings	Tag values. Each tag value can contain a maximum of 43 Unicode characters and cannot contain spaces. Before using values , delete spaces before and after the value. If the values is not specified, any parameter value can be queried. All values are in the OR relationship.

Table 6-23 Data structure description of parameter **matches**

Parameter	Mandatory	Type	Description
key	Yes	String	Query criteria. The value can be instance_name or instance_id , indicating that the query is based on the instance name or instance ID.
value	Yes	String	Name or ID of the instance to be queried

- Example request body
Querying an instance by tag:

```
{
  "offset": "100",
  "limit": "100",
  "action": "filter",
  "matches": [
    {
      "key": "instance_name",
      "value": "test-af07"
    }
  ],
}
```

```
"tags": [
  {
    "key": "key1",
    "values": [
      "value1",
      "value2"
    ]
  }
]
```

Querying the total number of records:

```
{
  "action": "count",
  "tags": [
    {
      "key": "key1",
      "values": [
        "value1",
        "value2"
      ]
    },
    {
      "key": "key2",
      "values": [
        "value1",
        "value2"
      ]
    }
  ],
  "matches": [
    {
      "key": "instance_name",
      "value": "test-af07"
    },
    {
      "key": "instance_id",
      "value": "958693039f284d6ebfb177375711072ein06"
    }
  ]
}
```

Response Parameters

- Normal response

Table 6-24 Parameter description

Parameter	Type	Description
instances	Array of objects	All instances
total_count	Integer	Total number of records

Table 6-25 Data structure description of parameter **instance**

Parameter	Type	Description
instance_id	String	Instance ID
instance_name	String	Instance name
tags	Array of objects	All tags. If there are no tags, tags is taken as an empty array by default. For more information, see Table 6-26 .

Table 6-26 Data structure description of parameter **tags**

Parameter	Type	Description
key	String	Tag key. The tag key can contain a maximum of 36 Unicode characters and must be specified. It is case-sensitive and can contain digits, letters, underscores (_), and hyphens (-).
value	String	Tag value. The tag value can contain a maximum of 43 Unicode characters and can be an empty string. It is case-sensitive and can contain digits, letters, underscores (_), periods (.), and hyphens (-).

- Example normal response

Returning a specified instance by tag:

```
{
  "instances": [
    {
      "instance_id": "2acbf2223caf3bac3c33c6153423c3ccin06",
      "instance_name": "test-single",
      "tags": [
        {
          "key": "key1",
          "value": "value1"
        },
        {
          "key": "key2",
          "value": "value1"
        }
      ]
    }
  ]
}
```

Returning total records:

```
{  
  "total_count": 4  
}
```

- Abnormal response
For details, see [Abnormal Request Results](#).

Status Codes

For more information, see [Status Codes](#).

Error Codes

For more information, see [Error Codes](#).

7 Permission Policies and Supported Actions

7.1 Introduction

You can use Identity and Access Management (IAM) for fine-grained management of the permissions for your GeminiDB databases. If your account does not need individual IAM users, then you may skip over this section.

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

You can grant users permissions by using [roles](#) and [policies](#). Roles are a type of coarse-grained authorization mechanism that defines permissions related to user responsibilities. Policies define API-based permissions for operations on specific resources under certain conditions, allowing for more fine-grained, secure access control of cloud resources.

NOTE

Policy-based authorization is useful if you want to allow or deny the access to an API.

An account has all the permissions required to call all APIs, but IAM users must be assigned the required permissions. The permissions required for calling an API are determined by the actions supported by the API. Only users who have been granted permissions allowing the actions can call the API successfully. For example, if an IAM user queries GeminiDB Cassandra instances using an API, the user must have been granted the permissions that allow the `nosql:instance:list` action.

Supported Actions

GeminiDB provides system-defined policies that can be directly used in IAM. You can also create custom policies and use them to supplement system-defined policies, implementing more refined access control. Operations supported by policies are specific to APIs. The following are common concepts related to policies:

- Permissions: Statements in a policy that allow or deny certain operations.
- APIs: REST APIs that can be called in a custom policy
- Actions: Added to a custom policy to control permissions for specific operations.
- IAM or enterprise projects: Type of projects for which an action will take effect. Policies that contain actions for both IAM and enterprise projects can be used and take effect for both IAM and Enterprise Management. Policies that only contain actions for IAM projects can be used and only take effect for IAM. For the differences between IAM and enterprise projects, see [Differences Between IAM and Enterprise Management](#).

For details about the custom actions supported by GeminiDB, see [GeminiDB Actions](#).

7.2 GeminiDB Actions

Table 7-1 Instance management actions

Permission	API	Action	IAM Project	Enterprise Project
Creating a DB instance	POST /v3/{project_id}/instances	nosql:instance:create	√	√
Deleting a DB instance	DELETE /v3/{project_id}/instances/{instance_id}	nosql:instance:delete	√	√
Querying DB instances	GET /v3/{project_id}/instances?id={id}&name={name}&mode={mode}&datastore_type={datastore_type}&vpc_id={vpc_id}&subnet_id={subnet_id}&offset={offset}&limit={limit}	nosql:instance:list	√	√
Scaling up storage space of an instance	POST /v3/{project_id}/instances/{instance_id}/extend-volume	nosql:instance:modifyStorageSize	√	√

Permission	API	Action	IAM Project	Enterprise Project
Adding nodes for a cluster instance	POST /v3/{project_id}/instances/{instance_id}/enlarge-node	nosql:instance:extendNode	√	√
Deleting nodes from a cluster instance	POST /v3/{project_id}/instances/{instance_id}/reduce-node	nosql:instance:reduceNode	√	√
Changing specifications of an instance	PUT /v3/{project_id}/instances/{instance_id}/resize	nosql:instance:modifySpecification	√	√
Changing the administrator password	PUT /v3/{project_id}/instances/{instance_id}/password	nosql:instance:modifyPasswd	√	√
Editing the name of an instance	PUT /v3/{project_id}/instances/{instance_id}/name	nosql:instance:rename	√	√
Changing the security group of an instance	PUT /v3/{project_id}/instances/{instance_id}/security-group	nosql:instance:modifySecurityGroup	√	√
Upgrading minor version	POST /v3/{project_id}/instances/{instance_id}/db-upgrade	nosql:instance:upgradeDatabaseVersion	√	√
Patching databases in batches	/v3/{projectId}/instances/db-upgrade	nosql:instance:batchUpgradeDatabaseVersion	√	√

Permission	API	Action	IAM Project	Enterprise Project
Creating cold storage	POST /v3/{project_id}/instances/{instance_id}/cold-volume	nosql:instance:modifyStorageSize	√	√
Scaling up cold storage	PUT /v3/{project_id}/instances/{instance_id}/cold-volume	nosql:instance:modifyStorageSize	√	√
Binding or unbinding an EIP	POST /v3/{project_id}/instances/{instance_id}/nodes/{node_id}/public-ip	nosql:instance:bindPublicIp	√	√
Enabling or disabling SSL	POST /v3/{project_id}/instances/{instance_id}/ssl-option	nosql:instance:switchSSL	√	√
Restarting a DB instance	POST /v3/{project_id}/instances/{instance_id}/restart	nosql:instance:restart	√	√
Configuring an autoscaling policy for storage space	PUT /v3/{project_id}/instances/disk-auto-expansion	nosql:instance:modifyStorageSize	√	√
Modifying high-risk commands	PUT /v3/{projectId}/instances/{instanceId}/high-risk-commands	nosql:instances:modifyHighRiskCommands	√	√
Setting the maintenance period of an instance	PUT /v3/{project_id}/instances/{instance_id}/maintenance-window	nosql:instance:modifyMaintenanceWindow	√	√

Permission	API	Action	IAM Project	Enterprise Project
Querying the password-free configuration of a GeminiDB Redis instance	Get /v3/{project_id}/instances/{instance_id}/passwordless-config	nosql:instance:getPasswordlessConfig	√	√
Modifying the password-free configuration of a GeminiDB Redis instance	PUT /v3/{project_id}/instances/{instance_id}/passwordless-config	nosql:instance:setPasswordlessConfig	√	√
Updating an agency permission policy	x	nosql:instance:updateAgency	√	√
Querying agency permission details	x	nosql:instance:listAgency	√	√
Querying memory acceleration rules and details	GET /v3/{project_id}/dbcache/mappings	nosql:instance:listDBCACHEMappings	√	√
Creating a memory acceleration rule	POST /v3/{project_id}/dbcache/rule	nosql:instance:listDBCACHERules	√	√
Deleting a memory mapping	DELETE /v3/{project_id}/dbcache/mapping	nosql:instance:deleteDBCACHEMapping	√	√
Creating a memory mapping	POST /v3/{project_id}/dbcache/mapping	nosql:instance:createDBCACHEMapping	√	√
Modifying a memory acceleration rule	PUT /v3/{project_id}/dbcache/rule	nosql:instance:updateDBCACHERule	√	√

Permission	API	Action	IAM Project	Enterprise Project
Querying memory acceleration rules and details	GET /v3/{project_id}/dbcache/rules	nosql:instance:listDBCACHERules	√	√
Deleting a memory acceleration rule	DELETE /v3/{project_id}/dbcache/rule	nosql:instance:deleteDBCACHERule	√	√
Enabling or disabling instance data export	PUT /v3/{project_id}/instances/{instance_id}/data-dump	nosql:instance:operateDataDump	√	√
Enabling or disabling second-level monitoring	PUT /v3/{project_id}/instances/{instance_id}/monitoring-by-seconds/switch	nosql:instance:secondLevelMonitoring	√	√
Querying second-level monitoring configurations	GET /v3/{project_id}/instances/{instance_id}/monitoring-by-seconds/switch	nosql:instance:secondLevelMonitoring	√	√
Querying hot keys of a GeminiDB Redis instance	GET /v3/{project_id}/instances/{instance_id}/hot-keys	nosql:instance:getHotKeys	√	√
Performing a primary/standby switchover for GeminiDB Redis instances	PUT /v3/{project_id}/instance/{instance_id}/switchover	nosql:instance:switchover	√	√

Permission	API	Action	IAM Project	Enterprise Project
Querying disabled commands for a GeminiDB Redis instance	GET /v3/{project_id}/redis/instances/{instance_id}/disabled-commands	nosql:instance:queryRedisDisabledCommands	√	√
Setting disabled commands for a GeminiDB Redis instance	POST /v3/{project_id}/redis/instances/{instance_id}/disabled-commands	nosql:instance:setRedisDisabledCommands	√	√
Deleting disabled commands for a GeminiDB Redis instance	DELETE /v3/{project_id}/redis/instances/{instance_id}/disabled-commands	nosql:instance:setRedisDisabledCommands	√	√

Table 7-2 Actions for backups and restorations

Permission	API	Action	IAM Project	Enterprise Project
Querying an automated backup policy	GET /v3.1/{project_id}/instances/{instance_id}/backups/policy	nosql:backup:list	√	√
Configuring an automated backup policy	PUT /v3/{project_id}/instances/{instance_id}/backups/policy	nosql:instance:modifyBackupPolicy	√	√
Querying instances that can be restored	GET /v3/{project_id}/backups/{backup_id}/restorable-instances	nosql:instance:list	√	√

Permission	API	Action	IAM Project	Enterprise Project
Querying the time window when a backup can be restored	GET /v3/{project_id}/instances/{instance_id}/backups/restorable-time-periods	nosql:backup:list	√	√
Creating a manual backup	POST /v3/{project_id}/instances/{instance_id}/backups	nosql:backup:create	√	√
Deleting a manual backup	DELETE /v3/{project_id}/backups/{backup_id}	nosql:backup:delete	√	√
Restoring data to an existing instance	POST /v3/{project_id}/instances/{instance_id}/recovery	nosql:backup:refreshInstanceFromBackup	√	√
Restoring data of the current GeminiDB Redis instance to a point in time	PUT /v3/{project_id}/redis/instances/{instance_id}/pitr	nosql:instance:redisPitrRestore	√	√
Setting a policy for restoring data of GeminiDB Redis instance to a point in time	PUT /v3/{project_id}/redis/instances/{instance_id}/pitr/policy	nosql:instance:setRedisPitrPolicy	√	√
Querying a policy for restoring data of GeminiDB Redis instance to a point in time	GET /v3/{project_id}/redis/instances/{instance_id}/pitr/policy	nosql:instance:showRedisPitrPolicy	√	√

Permission	API	Action	IAM Project	Enterprise Project
Querying the time range for restoring data of GeminiDB Redis instance to a point in time	GET /v3/{project_id}/redis/instances/{instance_id}/pitr/restorable-time-periods	nosql:instance:queryRedisPitrRestoreTime	√	√
Querying the storage space used for restoring data of GeminiDB Redis instance to a point in time	GET /v3/{project_id}/redis/instances/{instance_id}/pitr	nosql:instance:showRedisPitrInfo	√	√
Stopping a backup	PUT /v3/{project_id}/backups/{backup_id}	nosql:backup:stop	√	√
Deleting manual backups in batches	DELETE /v3/{project_id}/instances/backups	nosql:backup:delete	√	√

Table 7-3 Parameter template management actions

Permission	API	Action	IAM Project	Enterprise Project
Obtaining parameter templates	GET /v3/{project_id}/configurations	nosql:param:list	√	√
Creating a parameter template	POST /v3/{project_id}/configurations	nosql:param:create	√	√
Modifying parameters in a parameter template	PUT /v3/{project_id}/configurations/{config_id}	nosql:param:modify	√	√

Permission	API	Action	IAM Project	Enterprise Project
Applying a parameter template	PUT /v3.1/{project_id}/configurations/{config_id}/apply	nosql:instance:modifyParameter	√	√
Modifying parameters of a specified instance	PUT /v3.1/{project_id}/instances/{instance_id}/configurations	nosql:instance:modifyParameter	√	√
Obtaining parameters of a specified instance	GET /v3/{project_id}/instances/{instance_id}/configurations	nosql:param:list	√	√
Obtaining parameters of a specified parameter template	GET /v3/{project_id}/configurations/{config_id}	nosql:param:list	√	√
Deleting a parameter template	DELETE /v3/{project_id}/configurations/{config_id}	nosql:param:delete	√	√
Querying instances that a parameter template can be applied to	GET /v3/{project_id}/configurations/{config_id}/applicable-instances	nosql:instance:list	√	√
Viewing parameter change history of an instance	GET /v3/{project_id}/instances/{instance_id}/configuration-histories	nosql:param:list	√	√
Viewing application records of a parameter template	GET /v3/{project_id}/configurations/{config_id}/applied-histories	nosql:param:list	√	√

Table 7-4 Tag management actions

Permission	API	Action	IAM Project	Enterprise Project
Querying an instance by tag	POST /v3/{project_id}/instances/resource_instances/action	<ul style="list-style-type: none"> • nosql:instance:list • nosql:tag:list 	√	√
Adding or deleting resource tags in batches	POST /v3/{project_id}/instances/{instance_id}/tags/action	nosql:instance:tag	√	√
Querying tags of an instance	GET /v3/{project_id}/instances/{instance_id}/tags	<ul style="list-style-type: none"> • nosql:instance:list • nosql:tag:list 	√	√

Table 7-5 Log management actions

Permission	API	Action	IAM Project	Enterprise Project
Querying slow query logs of an instance	GET /v3/{project_id}/instances/{instance_id}/slowlog?start_date={start_date}&end_date={end_date}	nosql:instance:list	√	√

Table 7-6 Quota management actions

Permission	API	Action	IAM Project	Enterprise Project
Querying resource quotas	GET /v3/{project_id}/quotas	nosql:instance:list	√	√

Table 7-7 Actions for disaster recovery management

Permission	API	Action	IAM Project	Enterprise Project
Querying regions where a dual-active relationship can be created between two instances	GET /v3/{project_id}/instances/{instance_id}/disaster-recovery/regions	nosql:instance:list	√	√
Setting the percentage of faulty nodes to be taken over in case of a failover	PUT /v3/{project_id}/instances/disaster-recovery/settings	nosql:dr:setDisasterRecoverySettings	√	√
Querying the percentage of faulty nodes to be taken over in case of a failover	GET /v3/{project_id}/instances/disaster-recovery/settings	nosql:dr:listDisasterRecoverySettings	√	√

Table 7-8 Task management

Permission	API	Action	IAM Project	Enterprise Project
Querying the maintenance period of an instance	GET /v3/{project_id}/instances/{instance_id}/ops-window	nosql:instance:maintenanceWindow	√	√
Canceling a scheduled task	DELETE /v3/{project_id}/scheduled-jobs/{job_id}	nosql:instance:cancelScheduledJob	√	√
Querying a scheduled task	GET /v3/{projectId}/scheduled-jobs	nosql:task:list	√	√

 **NOTE**

The check mark (√) indicates that the action takes effect. The cross mark (x) indicates that the action does not take effect.

8 Appendixes

8.1 Abnormal Request Results

- Abnormal Response

Table 8-1 Parameter description

Parameter	Mandatory	Type	Description
error_code	Yes	String	Error code returned when a task submission exception occurs.
error_msg	Yes	String	Error description returned when a task submission exception occurs.

- Example abnormal response

```
{
  "error_code": "DBS.200001",
  "error_msg": "Parameter error"
}
```

8.2 Status Codes

- Normal

Status Code	Encoding	State Description
200	OK	Request succeeded.
204	No Content	Request succeeded, but no response is returned.
202	Accepted	Asynchronous request submitted successfully.

- Abnormal

Status Code	Encoding	State Description
400	Bad Request	Invalid request. Do not retry the request before modification.
401	Unauthorized	The authorization information provided by the client is incorrect or invalid. Check the username and password.
403	Forbidden	The request is rejected. The server understood the request, but is refusing to fulfill it. The client should not repeat the request without modifications.
404	Not Found	The requested resource could not be found. Do not retry the request before modification.
405	Method Not Allowed	The method specified in the request is not supported for the requested resource. Do not retry the request before modification.
409	Conflict	The request could not be processed due to a conflict. The resource that the client attempts to create already exists, or the update request fails to be processed because of a conflict.
413	Request Entity Too Large	The request is larger than that a server is able to process. The server may close the connection to prevent the client from continuing the request. If the server temporarily cannot process the request, the response will contain a Retry-After header field.
415	Unsupported Media Type	The server is unable to process the media format in the request.
422	Unprocessable Entity	The request is well-formed but is unable to be processed due to semantic errors.

Status Code	Encoding	State Description
500	Internal Server Error	The server is able to receive the request but unable to understand the request.
501	Not Implemented	The server does not support the requested function.
503	Service Unavailable	The requested service is unavailable. Do not retry the request before modification.

8.3 Error Codes

If an error occurs during API calling, no results will be returned. You can locate the error cause based on error codes of each API. If an error occurs, an HTTP status code is returned. The returned message body contains a specific error code and error message.

Error Code Description

If an error code starting with **APIGW** is returned after you call an API, rectify the fault by referring to the instructions provided in [Error Codes](#).

Table 8-2 Error code description

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.200001	Parameter error.	Parameter error.	Check whether transferred parameters or URLs are correct.
404	DBS.200002	The DB instance does not exist.	Instance not found.	Check whether the instance and its ID are correct and whether the instance exists.
400	DBS.200010	Authentication failed.	Authentication failed.	Check whether the tenant and instance match.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
403	DBS.200011	This instance {0} status makes it not be allowed to do this {1} operation now.	Operation cannot be performed in current state of the instance.	Check whether the instance status or the ongoing operation on the instance conflicts with the request.
404	DBS.200013	This node does not exist.	Node not found.	Check whether the node ID or group ID is correct.
403	DBS.200018	This instance's status or its node's status makes it not be allowed to do this operation now.	Unavailable instance.	Check whether the instance status or the ongoing operation on the instance conflicts with the request.
403	DBS.200019	Another operation is being performed on the DB instance or the DB instance is faulty.	Operation cannot be performed in current state of the instance.	Check whether the instance status or the ongoing operation on the instance conflicts with the request.
400	DBS.200022	The DB instance name {0} already exists.	Instance name already exists.	Check whether the instance name exists.
400	DBS.200024	The region is unavailable.	Unavailable region.	Check whether the region name is correct and whether the region is available.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.200025	Invalid AZ.	Invalid AZ.	Check whether the AZ name is correct and whether the AZ is available.
403	DBS.200028	Volume size reach limit.	Maximum storage space has been reached.	Check whether the storage space exceeds the upper limit.
400	DBS.200029	Invalid username and password.	Invalid username and password.	Check whether the username and password match and whether the password meets password strength requirements.
400	DBS.200041	Invalid datastore version.	Invalid database version.	Check whether the database version is supported.
404	DBS.200042	Invalid database engine.	Invalid DB API.	Check whether the DB version is supported.
400	DBS.200047	Operation cannot be executed in current state of the DB instance or node.	Operation cannot be executed in current state of the instance or node.	Check whether the instance status or the ongoing operation on the instance conflicts with the request.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.200048	Invalid VPC.	Invalid VPC.	Check whether the VPC ID and name are correct and meet the requirements.
400	DBS.200049	Invalid subnet.	Invalid subnet.	Check whether the subnet ID and name are correct and meet the requirements.
400	DBS.200050	Invalid security group.	Invalid security group.	Check whether the security group ID and name are correct and meet the requirements.
400	DBS.200052	Invalid password.	Invalid password.	Check whether the username and password match and whether the password meets password strength requirements.
400	DBS.200053	The DB instance specifications do not exist.	Instance specifications not found.	Check whether the specifications are correct and supported in the current AZ.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.200054	Invalid DB instance specifications.	Invalid instance specifications.	Check whether the specifications are correct and supported in the current AZ.
400	DBS.200057	Invalid parameter group ID.	Invalid parameter template.	Check whether the parameter template is supported.
404	DBS.200058	Parameter template does not exist.	Parameter template not found.	Check whether the parameter template exists.
400	DBS.200059	Invalid database port.	Invalid database port.	Check whether the database port is missing or valid.
400	DBS.200060	The database port number is out of the specified range.	Database port is not in the specified range.	Check whether the database port is valid.
400	DBS.200063	Invalid DB instance type.	Invalid instance type.	Check whether the instance type is valid.
400	DBS.200068	Weak password.	Weak password.	Password is too easy to guess. Change it to a strong password.
400	DBS.200072	Invalid storage space.	Invalid storage space.	Check whether the storage space exceeds the upper limit.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.200075	Invalid node role.	Invalid node role.	Check whether the role of the node meets the requirements and whether the instance is normal.
403	DBS.200076	Operation cannot be executed in current state of the DB instance.	Operation cannot be performed in current state of the instance.	Check whether the instance status or the ongoing operation on the instance conflicts with the request.
400	DBS.200077	Failed to change the password.	Updating password failed.	Check whether the username is correct.
400	DBS.200091	Invalid IP address.	Invalid IP address.	Check whether the required IP address is missing or whether the input IP address is valid.
400	DBS.200092	The IP address already exists.	IP address already exists.	Check whether the IP address exists.
400	DBS.200095	Invalid parameter.	Parameter error.	Check whether parameters in the request and URLs are correct.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
500	DBS.200165	Insufficient enterprise project instance quota range.	Enterprise project quota of the instance is insufficient.	Increase the enterprise project quota on the GeminiDB console.
500	DBS.200166	Insufficient enterprise project cpu quota range.	CPU quota of the enterprise project is insufficient.	Increase the CPU quota on the GeminiDB console.
500	DBS.200167	Insufficient enterprise project mem quota range.	Memory quota of the enterprise project is insufficient.	Increase the memory quota on the GeminiDB console.
400	DBS.200302	The storage space increase must be a positive integer.	Storage space increase is not a positive integer.	Check whether the storage space increase is a positive integer.
400	DBS.200303	The maximum number of times that the storage space can be scaled up has been reached.	Storage space has reached its upper limit and cannot be scaled up anymore.	Contact technical support to scale up the storage space.
400	DBS.200304	The storage space can be scaled up for a maximum of four times.	Storage space can be scaled up at most four times.	Check how many times the instance has been scaled up.
400	DBS.200306	Invalid storage space.	Invalid storage space.	Check whether the storage space is correct and meets the requirements.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.200311	Scaling up storage space is not allowed in current state of the node.	Scaling up storage space not allowed on the current node.	Check whether the node type, instance type, and node ID are correct.
400	DBS.200434	Failed to restart the DB instance.	Restarting instance failed.	Check whether the instance is available or ongoing other operations.
400	DBS.200451	The node does not exist.	Node not found.	Check whether the node ID is correct.
400	DBS.200462	The database port is the same as the current port.	Database port is the same as current port.	Check whether the new port number is the same as the original port number.
400	DBS.200470	Invalid AZ.	Invalid AZ.	Check whether the AZ is correct.
400	DBS.200501	The subnet does not exist.	Subnet deleted.	Check whether the subnet ID and name exist and whether the subnet matches the VPC.
400	DBS.200502	The security group does not exist.	Security group is not found or does not belong to the VPC.	Check whether the security group ID and name exist and whether the security group matches the VPC.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.200503	The VPC does not exist.	VPC deleted.	Check whether the VPC is available to the tenant.
400	DBS.200506	The encryption key does not exist.	Storage encryption key ID not found.	Check whether the disk encryption key ID exists.
400	DBS.200507	The encryption key is not available.	Storage encryption key unavailable.	Check whether the disk encryption key is available.
400	DBS.200604	The instance is not owned by the current user.	Instance does not belong to the current user.	Check whether the project ID is subordinate to the instance ID.
400	DBS.200700	The EIP status does not allow EIP binding.	EIP is being bound and cannot be bound again.	Check whether there is an EIP being bound to the instance.
400	DBS.200701	The EIP status does not allow EIP unbinding.	EIP cannot be unbound.	Check whether the EIP status allows an unbinding operation.
400	DBS.200702	The node has been bound to a public IP address and cannot be bound again.	Node has already an EIP bound and cannot be bound with another EIP.	Check whether an EIP has been bound to the node.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.200705	The PublicIp does not exist.	Target EIP not found.	Check whether the EIP exists and whether the EIP and its ID match.
400	DBS.200816	Failed to create the database user.	Creating database user failed.	Check whether the database user name is valid and check the database status and instance status.
400	DBS.200817	Failed to obtain the database user list.	Obtaining database users failed.	Check the database status and instance status.
400	DBS.200818	Failed to delete the database user.	Deleting database user failed.	Check the database status and instance status.
400	DBS.200823	The database does not exist.	Database not found.	Check whether the database name is valid.
400	DBS.200824	The database account does not exist.	Database account not found.	Check whether the database user name is valid.
400	DBS.200826	The database name already exists.	Database name already exists.	Check whether the database name is valid.
400	DBS.200827	The database user already exists.	Database account name already exists.	Check whether the database user is valid.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.200828	Built-in database accounts cannot be edited.	This is an internal database account and cannot be operated by users.	Check whether the database user is valid.
500	DBS.200998	The system is busy. Try again later.	The system is busy. Try again later.	The system is busy. Try again later.
403	DBS.201000	The status of DB instance {0} does not allow the {1} operation.	Operation cannot be performed in current state of the instance.	Check whether the instance status or the ongoing operation on the instance conflicts with the request.
400	DBS.201006	Parameter error.	Parameter error.	Check whether transferred parameters or URLs are correct.
403	DBS.201014	Operation cannot be executed in current state of the DB instance.	Operation cannot be performed in current state of the instance.	Check whether the instance status or the ongoing operation on the instance conflicts with the request.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
403	DBS.201015	This operation cannot be performed because another operation is being performed on the DB instance or the DB instance is faulty. Try again later.	Operation cannot be performed because another operation is being performed.	Check whether the instance status or the ongoing operation on the instance conflicts with the request.
400	DBS.201020	Invalid DB engine.	Invalid DB API.	Check whether the database API is the API supported by the GeminiDB service.
403	DBS.201028	The DB instance does not exist.	Instance not found.	Check whether the instance belongs to the tenant and whether the instance exists.
400	DBS.201035	The database name must be different.	Database name already exists.	Check whether the same database name exists.
400	DBS.201038	The collection name must be different.	Collection name already exists.	Check whether the same collection name exists.
400	DBS.201101	Invalid backup period.	Invalid backup cycle.	Check whether the backup cycle meets the requirements.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.201106	Invalid retention period.	Invalid retention period.	Check whether the backup retention period is correct.
400	DBS.201201	The backup already exists.	Backup file already exists.	Check whether the backup name or ID already exists.
400	DBS.201202	Operation cannot be executed in current state of the DB instance.	Operation cannot be performed in current state of the instance.	Check whether the instance status or the ongoing operation on the instance conflicts with the request.
400	DBS.201204	Operation cannot be executed in current state of the DB instance.	Backup file not found.	Check whether the backup file exists and matches the instance.
400	DBS.201212	Backup ID is illegal.	Invalid backup ID.	Check whether the backup ID exists.
400	DBS.201214	The backup file does not exist.	Backup file not found.	Check whether the backup file exists and matches the instance.
400	DBS.201215	Time is illegal.	Invalid time.	Check whether the time meets the requirements.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.201319	Deleting backup file is not allowed because a restoration task is currently in progress. Please wait.	Original backup file cannot be deleted during restoration.	Check whether the backup is being used to restore instances.
400	DBS.201501	The DB instance does not exist.	Instance not found.	Check whether the instance belongs to the tenant, whether the instance name or ID is correct, and whether the instance exists.
400	DBS.201502	The DB instance does not exist.	Instance not found.	Check whether the instance belongs to the tenant, whether the instance name or ID is correct, and whether the instance exists.
400	DBS.280308	The AZ is inconsistent.	Inconsistent AZs.	Check whether the instance is in the same AZ as dedicated resources.
400	DBS.212001	The parameter group {0} does not exist.	Parameter template not found.	Check whether the parameter template exists.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.212003	This operation is not permitted.	Operation not allowed.	Check whether the instance status or the ongoing operation on the instance conflicts with the request.
400	DBS.212006	The node associated with this parameter group is not available.	Node associated with the parameter template is abnormal.	Check whether the node associated with the parameter template is normal.
400	DBS.212008	The database type does not exist.	Unsupported database version.	Check whether the database version supports this operation.
400	DBS.212013	The parameter group does not exist.	Parameter template not found.	Check whether the parameter template exists or belongs to the current tenant.
400	DBS.212017	The parameter is invalid.	Invalid parameter.	Check whether transferred parameters or URLs are correct and meet the requirements.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.212019	Invalid parameter.	Invalid parameter.	Check whether transferred parameters or URLs are correct and meet the requirements.
400	DBS.212028	Invalid description.	Invalid parameter template description.	Check whether the parameter template description is valid.
400	DBS.212030	The parameter group name already exists.	Parameter template name already exists.	Check whether the parameter template name exists.
400	DBS.212031	Invalid parameter group name.	Invalid parameter template name.	Check whether the parameter template name is valid.
400	DBS.212032	The operation cannot be performed because this parameter group is being applied to one or more DB instance nodes.	Operation is not allowed because the parameter template is applied to one or more instance nodes.	Check whether the parameter template has been applied to the instance.
400	DBS.212035	Failed to associate this parameter group with the DB instance because the DB instance is currently being operated.	Operation cannot be performed in current state of the instance.	Check whether the instance is ongoing other operations.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.212037	Parameters are incorrectly set.	Parameter error.	Check whether the parameter value is valid or within the valid range.
400	DBS.216016	DB instance does not exist.	Instance not found.	Check whether the instance exists.
400	DBS.216029	This operation is not allowed.	Operation not allowed.	Check whether the operation is valid or whether the current engine supports the operation.
400	DBS.216030	The queried node does not belong to the current instance.	Queried node does not belong to the current instance.	Check whether the input node belongs to the current instance.
400	DBS.216044	Invalid vpc cidr.	Invalid VPC CIDR block.	Check whether the format of the input VPC CIDR block is valid. Check whether the DR instance has overlapping VPC CIDR blocks in DR scenarios.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.216045	Invalid db engine version.	Invalid DB API version.	Check whether the DB engine version is available or whether the version supports the current operation.
400	DBS.216046	Invalid subnet cidr.	Invalid subnet CIDR block.	Check whether the format of the input VPC CIDR block is valid.
400	DBS.238007	This operation cannot be performed in the current IP address status.	Operation cannot be performed in the current IP address state.	Check whether the delivered IP address is in use.
400	DBS.239010	Reduce num or target invalid.	Invalid quantity of nodes to be deleted or invalid nodes.	Check whether the input node belongs to the current instance.
400	DBS.239011	Reduce num invalid.	Invalid number of nodes to be deleted.	Check whether the number of input nodes is valid.
400	DBS.239012	Reduce target invalid.	Invalid node to be deleted.	Check whether the input node belongs to the current instance.
400	DBS.239013	resize flavor invalid.	Inappropriate specifications.	Check whether the input specifications are valid.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.239014	Current disk capacity does not support scaling in the number of nodes.	Current storage space does not allow deleting nodes.	Check the storage space of the instance.
400	DBS.240001	node num inconsistent.	Inconsistent nodes.	Check whether the number of nodes is consistent.
400	DBS.240002	node num incorrect.	Invalid quantity of nodes.	Check whether the number of nodes is valid.
400	DBS.240003	Invalid billing mode.	Invalid billing mode.	Check whether the billing mode is valid.
400	DBS.240004	Invalid subscription type.	Invalid subscription type.	Check whether the subscription type is valid.
400	DBS.240005	Invalid validity period.	Invalid required duration.	Check whether the required duration is valid.
400	DBS.240006	Invalid renewal mode.	Invalid renewal mode.	Check whether the renewal mode is valid.
400	DBS.240007	Invalid payment mode.	Invalid payment mode.	Check whether the payment mode is valid.
400	DBS.240008	Submit order failed.	Order submission failed.	Contact customer service.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.240009	Deleting nodes cannot be executed in current state of the DB instance.	Deleting nodes is not allowed in current state of the instance.	Check whether the current instance has nodes that can be deleted.
400	DBS.240010	The selected nodes do not support shrinkage.	Selected nodes cannot be deleted.	Check whether the selected nodes can be deleted.
400	DBS.240011	{0}	Order submission failed.	Check whether an error message is returned.
400	DBS.240012	The maximum number of resources that can be changed has been reached.	The maximum number of resources that can be changed has been reached.	Check whether the number of nodes in the current instance exceeds the upper limit.
400	DBS.240013	The current disk capacity cannot be changed to the target flavor.	Specifications cannot be changed because the storage space is too small.	Check whether the storage space of the current instance exceeds the upper limit defined by the target instance specifications.
400	DBS.280001	Parameter error.	Parameter error.	Check whether transferred parameters or URLs are correct and meet the requirements.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
500	DBS.280005	Server error. Try again later.	Server error. Try again later.	Contact technical support.
400	DBS.280015	Permission denied.	Insufficient permissions.	Check whether the token expires and whether the instance matches the tenant.
400	DBS.280016	Resource not found.	Resource not found.	Check whether transferred parameters are correct and whether the instance exists.
403	DBS.280019	Account suspended.	Account in arrears.	Check the account balance.
403	DBS.280032	You do not have permission to perform this operation. Contact the administrator to obtain permission.	Permission denied.	Check whether the user group to which the current user belongs has the corresponding operation permission.
400	DBS.280042	Invalid request.	Invalid request.	Check whether the current instance status and the ongoing operation allow this operation or whether the request is valid.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.280040	Parameter is null.	Parameter missed.	Check whether there are parameters not input.
404	DBS.280045	This parameter group does not exist.	Parameter template not found.	Check whether the parameter template exists in the request.
403	DBS.280056	Token invalid.	Invalid token.	Check whether the instance belongs to the tenant and whether the token has been obtained again.
400	DBS.280063	You do not have following permission to perform this operation: {0}. Contact the administrator to obtain permission.	Current user does not have the permission to perform xxx operation. Contact the account administrator.	Check whether the user group to which the current user belongs has the corresponding operation permission.
500	DBS.280064	Check PDP permissions failed.	Fine-grained authentication failed.	Contact customer service.
400	DBS.280066	Invalid log type.	Invalid log type.	Check whether the log type meets the requirements.
400	DBS.280067	Invalid start time.	Invalid start time.	Check whether the start time meets the requirement.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.280068	Invalid end time.	Invalid end time.	Check whether the end time meets the requirement.
400	DBS.280110	The DB instance does not exist.	Instance not found.	Check whether the instance belongs to the tenant, whether the instance name or ID is correct, and whether the instance exists.
400	DBS.280122	Invalid DB engine.	Invalid storage API.	Check whether the storage engine matches the instance engine.
400	DBS.280123	Invalid node number.	Invalid node quantity.	Check whether the number of nodes meets the requirements.
400	DBS.280124	Invalid backup.	Invalid backup ID.	Check whether the backup ID is correct and meets the requirements.
400	DBS.280125	Invalid backup policy.	Invalid automated backup policy.	Check whether the automated backup policy meets the requirements.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.280127	Invalid backup description.	Invalid backup description.	Check whether the backup description is correct and meets the requirements.
400	DBS.280200	The password contains invalid characters.	Password contains invalid characters.	Check whether the password is correct and meets the requirements.
400	DBS.280214	Invalid retention period.	Invalid retention period.	Check whether the backup retention period is correct.
400	DBS.280215	Invalid backup cycle.	Invalid backup cycle.	Check whether the backup start time, end time, and backup cycle are correct and meet the requirements.
400	DBS.280216	Invalid backup start time.	Invalid backup start time.	Check whether the backup start time meets the requirements and whether the relationship between the backup start time and end time is rational.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.280234	Invalid DB instance name.	Invalid instance name.	Check whether the instance name is correct and whether the instance exists.
400	DBS.280235	Invalid database type.	Invalid DB API.	Check whether the DB engine name is correct.
400	DBS.280236	Invalid database version.	Invalid database version.	Check whether the database version is supported.
400	DBS.280239	Invalid specifications.	Invalid specifications.	Check whether specification code is correct and whether the specifications are available in the AZ.
400	DBS.280240	The specification does not exist.	Specification code not found.	Check whether the specification code is correct and the specifications are available.
400	DBS.280241	Invalid storage type.	Invalid storage type.	Check whether the storage type is correct and meets the requirements.
400	DBS.280242	The storage space is out of range.	Storage space is out of range.	Check whether the storage space is correct.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.280244	Invalid AZ.	Invalid AZ.	Check whether parameters of the AZ are correct, whether the AZ exists, and whether the AZ matches the specifications.
400	DBS.280247	Invalid VPC.	Invalid VPC.	Check whether the VPC ID is correct and whether the VPC exists.
400	DBS.280248	Invalid subnet.	Invalid subnet.	Check whether the subnet ID is correct and whether the subnet exists.
400	DBS.280249	Invalid security group.	Invalid security group.	Check whether the security group ID is correct and whether the security group exists.
400	DBS.280251	Invalid backup period.	Invalid backup cycle.	Check whether the backup cycle meets the requirements.
400	DBS.280266	Invalid storage space.	Invalid storage space.	Check whether the storage space is correct and meets the requirements.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.280267	Specifications not match.	Specifications do not match.	Check whether the specification information is correct and whether the specifications match the instance.
400	DBS.280269	Invalid Datastore Info.	Invalid database information.	Check whether datastore information is correct and meets the requirements.
400	DBS.280277	Invalid backup name.	Invalid backup name.	Check whether the backup name is correct and meets the requirements.
400	DBS.280280	Invalid DB instance number.	Invalid instance quantity.	Check whether the number of instances is correct and meets the requirements.
400	DBS.280284	Invalid IP address.	Invalid IP address.	Check whether the IP address is correct and meets the requirements.
400	DBS.280292	Invalid database username.	Invalid username.	Check whether the username is correct and meets the requirements.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.280311	Invalid storage space.	Invalid storage space.	Check whether the storage space is correct and meets the requirements.
400	DBS.280314	Invalid storage space.	Invalid storage type.	Check whether the storage type is correct and whether the instance supports the disk type.
400	DBS.280327	Invalid node type.	Invalid node type.	Check whether the node type is correct, whether the node type matches the instance, and whether the node type matches the group ID and node ID.
400	DBS.280341	Invalid DB instance type.	Invalid instance type.	Check whether the GeminiDB APIs support the instance.
400	DBS.280342	Invalid DB instance mode.	Invalid instance type.	Check whether the instance type is correct and matches the instance ID.
400	DBS.280347	Unsupported database type.	Unsupported instance type.	Check whether the instance type is correct and meets the requirements.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.280365	Invalid payment mode.	Invalid payment mode.	Check whether the payment mode is correct and meets the requirements.
400	DBS.280366	Invalid order ID.	Invalid order ID.	Check whether the yearly/ monthly order ID is transferred and meets the requirements.
400	DBS.280391	Action check states.	Invalid action.	Check whether the input action meets the requirements.
400	DBS.280404	Invalid DB instance ID.	Invalid instance ID.	Check whether the instance ID is correct and meets the requirements.
403	DBS.280406	The DB instance cannot be deleted.	Instance cannot be deleted.	Check whether the DB API and billing mode support deletion of instances.
400	DBS.280407	Invalid node ID.	Invalid node ID.	Check whether the node ID is correct and meets the requirements.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.280408	Invalid project id.	Invalid project ID.	Check whether the project ID is correct and meets requirements.
400	DBS.280414	Invalid group type.	Invalid group type.	Check whether the instance group type is correct and meets the requirements and whether it matches the instance and group ID.
400	DBS.280416	Invalid backup end time.	Invalid backup end time.	Check whether the backup end time is missing and whether the backup end time period and format meet the requirements.
400	DBS.280421	Invalid EIP.	Invalid EIP.	Check whether the EIP exists and meets the requirements.
403	DBS.280433	Invalid enterprise project ID.	Invalid enterprise project ID.	Check whether the enterprise project ID meets the requirements.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.280434	The specifications are unavailable.	Invalid resource specification code.	Check whether the resource specification code exists and meets the requirements.
400	DBS.280437	Not support enterprise project.	Enterprise project not supported.	The current user has not enabled the enterprise project service. Enable it or do not transfer related parameters.
400	DBS.280438	Invalid encryption key.	Invalid storage encryption key ID.	Check whether there is a disk encryption key ID available in the request and whether the current DB API supports disk encryption.
400	DBS.280439	Invalid limit.	Invalid query limit.	Check whether the value of the limit parameter is valid.
400	DBS.280440	Invalid offset.	Invalid offset.	Check whether the value of the offset parameter is valid.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.280441	Invalid key.	Invalid key.	Check whether the tag key is valid.
429	DBS.280443	The maximum number of connections has been reached.	Maximum connections reached.	Too frequent API requests. Try again later.
400	DBS.280444	Invalid value.	Invalid tag value.	Check whether the tag value is valid.
400	DBS.280445	The DB instance class is not available.	Unavailable instance specifications.	The current instance specifications are unavailable. Select another one.
400	DBS.280446	The database information does not exist.	Database information not found.	Check whether the datastore parameter exists.
400	DBS.280453	Invalid DSS storage pool ID.	Invalid DSS storage pool ID.	Check whether the storage pool ID is correct.
400	DBS.280465	Invalid password.	Invalid user password.	Check whether the user password meets the requirements.
400	DBS.280469	Invalid ECS group policy.	Invalid policy associated with the ECS group.	Check whether the policy associated with the ECS group is correct.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.280480	The target specification is same as current.	Target specifications are the same as current specifications.	Check whether the target specifications are the same as the current specifications.
400	DBS.280456	This operation is not supported by the current billing mode.	Current billing mode does not support this operation.	Check whether the billing mode of the current instance meets the API requirements.
400	DBS.290000	Parameter error.	Parameter error.	Check whether transferred parameters or URLs are correct and meet the requirements.
400	DBS.301024	Invalid backup restore information.	Invalid backup and restoration information.	Check whether the backup and restoration information is valid.
400	DBS.301040	The target security group is same as current.	Target security group is the same as current security group.	Check whether the target security group is the same as the current security group.
400	DBS.301071	Invalid session id.	Invalid session ID.	Check whether the session ID is correct and meets the requirements.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.301072	Invalid plan summary.	Invalid execution plan.	Check whether the execution plan is correct and meets the requirements.
400	DBS.301073	Invalid operation type.	Invalid operation type.	Check whether the operation type is correct and meets the requirements.
400	DBS.301074	Invalid namespace.	Invalid namespace.	Check whether the namespace is correct and meets the requirements.
400	DBS.301075	Invalid cost time.	Invalid execution time.	Check whether the execution time is correct and meets the requirements.
400	DBS.301076	Query session failed.	Querying session failed.	Contact customer service to check the instance status.
400	DBS.301077	Kill session failed.	Killing session failed.	Contact customer service to check the instance status.
400	DBS.03000001	The instance has not built a dual-active DR relationship.	Instance has no dual-active DR relationships.	Check whether the instance has a dual-active DR relationship.

HTTP Status Code	Error Code	Error Message	Description	Handling Measure
400	DBS.0300000 2	Invalid threshold.	Incorrect input threshold.	Check whether the input threshold meets the requirements.
400	DBS.0300000 3	Invalid step.	Inputting autoscaling increment failed.	Check whether the input increase step meets the requirements.
400	DBS.0300000 4	Invalid switch option.	Invalid parameter settings.	Check whether the input switch option is on or off .

8.4 Obtaining a Project ID

Scenarios

When calling APIs, you need to specify the project ID in some URLs. To do so, you need to obtain the project ID first.

You can obtain the required project ID with either of the following methods:

- [Obtaining the Project ID by Calling an API](#)
- [Obtaining a Project ID from the Console](#)

Obtaining the Project ID by Calling an API

You can obtain the project ID by calling the API used to [query project information based on the specified criteria](#).

The API used to obtain a project ID is **GET https://{Endpoint}/v3/projects/{Endpoint}**. **{Endpoint}** is the IAM endpoint and can be obtained from [Regions and Endpoints](#). For details about API authentication, see [Authentication](#).

The following is an example response. The value of **id** is the project ID.

```
{
  "projects": [
    {
      "domain_id": "65382450e8f64ac0870cd180d14e684b",
      "is_domain": false,
      "parent_id": "65382450e8f64ac0870cd180d14e684b",
      "name": "project_name",
      "description": ""
    }
  ]
}
```



```

    "links": {
      "next": null,
      "previous": null,
      "self": "https://www.example.com/v3/projects/a4a5d4098fb4474fa22cd05f897d6b99"
    },
    "id": "a4a5d4098fb4474fa22cd05f897d6b99",
    "enabled": true
  }
],
"links": {
  "next": null,
  "previous": null,
  "self": "https://www.example.com/v3/projects"
}
}

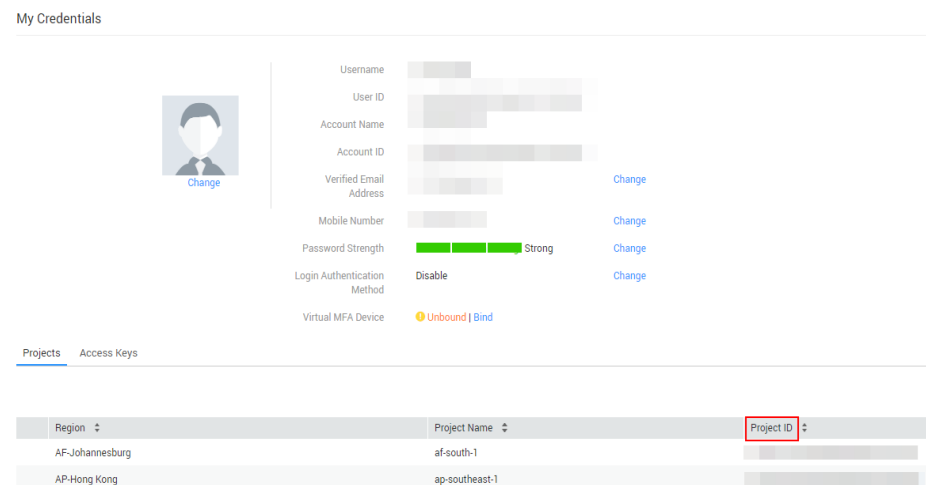
```

Obtaining a Project ID from the Console

1. Sign up and log in to the management console.
2. Move your pointer over the username and select **My Credentials** in the displayed drop-down list.

On the **Projects** tab page, view project IDs.

Figure 8-1 Viewing project IDs



8.5 Metrics

Function

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

Namespace

SYS.NoSQL

GeminiDB Redis Metrics

Table 8-3 Instance metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis816_cluster_disk_usage	Storage Usage	Storage usage of an instance Unit: %	0-100	GeminiDB Redis instances	1 minute
redis813_cluster_slow_query_count	Slow Queries per Instance	Number of slow queries on an instance Unit: count	≥ 0	GeminiDB Redis instances	1 minute
redis812_cluster_processed_command_count	Total Commands per Instance Total commands processed by an instance	Total commands processed by a node Unit: count	≥ 0	GeminiDB Redis instances	1 minute
redis811_cluster_max_connections_per_instance	Max. Connections per Instance	Maximum connections to an instance Unit: count	≥ 0	GeminiDB Redis instances	1 minute
redis808_cluster_new_client_connection	Connections Created Per Second	Instance connections created per second Unit: count	≥ 0	GeminiDB Redis instances	1 minute
redis807_cluster_all_connections	Total Connections	Total connections to the instance Unit: count	≥ 0	GeminiDB Redis instances	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis806_cluster_max_connection_usage	Connection Usage	Max. node connection usage of the instance Unit: %	0-100	GeminiDB Redis instances	1 minute
redis805_cluster_avg_hit_rate	Average Hit Rate of the Instance	Average hit rate of multiple nodes in the instance Unit: %	0-100	GeminiDB Redis instances	1 minute
redis804_cluster_all_p99_usec	p99 Latency	p99 latency of the instance Unit: μ s	≥ 0	GeminiDB Redis instances	1 minute
redis803_cluster_all_avg_usec	Average Latency	Average latency of the instance Unit: μ s	≥ 0	GeminiDB Redis instances	1 minute
redis802_cluster_max_response_argsc	Max. Elements Obtained in a Request	Max. elements obtained by the client in a request Unit: count	≥ 0	GeminiDB Redis instances	1 minute
redis801_cluster_max_response_bytes	Max. Bytes Obtained in a Request	Max. bytes obtained by the client in a request Unit: byte	≥ 0	GeminiDB Redis instances	1 minute
redis800_cluster_max_request_argsc	Max. Parameters Sent in a Request	Max. parameters sent in a request Unit: count	≥ 0	GeminiDB Redis instances	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis799_cluster_max_request_bytes	Max. Bytes Sent in a Request	Max. bytes sent in a request Unit: byte	≥ 0	GeminiDB Redis instances	1 minute
redis798_cluster_expire_key_counts	Keys with an Expiration Time Configured	Instance keys with an expiration time configured Unit: count	≥ 0	GeminiDB Redis instances	1 minute
redis689_ops_receive_total	Total Traffic Received by the Instance	Total traffic received by the instance, a reflection of the traffic volume on the application side. Unit: byte/s	≥ 0	GeminiDB Redis instances	1 minute
redis688_ops_send_total	Total Traffic Sent by the Instance	Total traffic sent by the instance, a reflection of the traffic volume on the application side. Unit: byte/s	≥ 0	GeminiDB Redis instances	1 minute
redis668_cluster_key_counts	Instance Keys	Total keys of a cluster Unit: count	≥ 0	GeminiDB Redis instances	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis667_cluster_qps	QPS of the Instance	The value of this metric is the QPS of the instance. Unit: count/s	≥ 0	GeminiDB Redis instances	1 minute

Table 8-4 Basic metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis817_slow_query_count	Slow Queries per Node	Slow queries on a node Unit: count	≥ 0	GeminiDB Redis instance nodes	1 minute
redis810_max_connections_count	Max. Connections per Node	Maximum connections to a node Unit: count	≥ 0	GeminiDB Redis instance nodes	1 minute
redis809_processed_commands_count	Total Commands per Node	Total commands processed by a node Unit: count	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis687_client_network_packet_return_p99	p99 Latency for Packets Returning to the Client (Send)	p99 latency for packets sent from the proxy to the client, a reflection of the network quality on the application side. Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis686_client_network_packet_return_max	Maximum Speed for Packets Returning to the Client (Send)	Maximum speed for packets sent by the proxy to the client, a reflection of the network quality on the application side. Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis685_client_network_packet_return_avg	Average Speed for Packets Returning to the Client (Send)	Average speed for packets sent by the proxy to the client, a reflection of the network quality on the application side. Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis684_client_packet_return_queue_p99	p99 Latency for Packets Returning to the Client (Queuing)	Queuing p99 latency for packets returned from the proxy to the client, a reflection of the network quality on the application side. Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis683_client_packet_return_queue_max	Maximum Speed for Packets Returning to the Client (Queuing)	Maximum queuing speed for packets returned from the proxy to the client, a reflection of the network quality on the application side. Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis682_client_packet_return_queue_avg	Average Speed for Packets Returning to the Client (Queuing)	Average queuing speed for packets returned from the proxy to the client, a reflection of the network quality on the application side. Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis678_max_response_argc	Maximum Elements Obtained in a Request	Maximum elements obtained by the client in a request Unit: count	≥ 0	GeminiDB Redis instance nodes	1 minute
redis677_max_response_bytes	Maximum Bytes Obtained in a Request	Maximum bytes obtained by the client in a request Unit: byte	≥ 0	GeminiDB Redis instance nodes	1 minute
redis676_max_request_argc	Maximum Parameters Sent in a Request	Maximum parameters sent by the client in a request Unit: count	≥ 0	GeminiDB Redis instance nodes	1 minute
redis675_max_request_bytes	Maximum Bytes Sent in a Request	Maximum bytes sent by the client in a request Unit: byte	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis674_max_pipeline_d	Maximum Commands Sent in a Pipeline	Maximum commands batch sent by the client in a pipeline Unit: count	≥ 0	GeminiDB Redis instance nodes	1 minute
redis673_wrong_auth	Failed Authentication Attempts	Failed authentication attempts per second on a node Unit: count	≥ 0	GeminiDB Redis instance nodes	1 minute
redis672_no_auth	Request Attempts Due to Authentication Failure	Failed request attempts due to authentication failure on a node per second Unit: count	≥ 0	GeminiDB Redis instance nodes	1 minute
redis671_new_client_connection	New Connections	Connections created on a node per second Unit: count	≥ 0	GeminiDB Redis instance nodes	1 minute
redis670_hit_rate	Hit Rate of a Key in Underlying Storage	Hit percentage of a key in underlying storage in a collection period. Formula: Hit keys / (Hit keys + Missed keys). Unit: %	0-100	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis669_connection_usage	Connection Usage	Connection usage in a collection period. Formula: Used connections/Total connections. Unit: %	0-100	GeminiDB Redis instance nodes	1 minute
redis319_all_qps	Proxy QPS	Proxy QPS on a node Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis318_all_p99	p99 Access Latency	p99 latency when a node executes all commands Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis317_all_max_usec	Maximum Access Latency	Maximum latency when a node executes all commands Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis316_all_avg_usec	Average Access Latency	Average latency when a node executes all commands Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis032_shard_qps	Shard QPS	Shard QPS on a node Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis021_proxy_send_client_bps	Traffic Send Speed	Outgoing traffic speed of a node Unit: byte/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis020_proxy_recv_client_bps	Traffic Receive Speed	Incoming traffic speed of a node Unit: byte/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis019_proxy_response_ps	Proxy Response Rate	Speed at which proxy responds to clients Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis018_proxy_request_ps	Request Receive Speed	Speed at which proxy receives requests from clients Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis017_proxy_accept	Client Connections	Client connections to a node Unit: count	≥ 0	GeminiDB Redis instance nodes	1 minute
nosql007_disk_used_size	Storage Space Usage	Used storage space of an instance Unit: GB	≥ 0	GeminiDB Redis instance nodes	1 minute
nosql006_disk_total_size	Total Storage Space	Total storage space of an instance Unit: GB	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
nosql005_disk_usage	Storage Usage	Storage usage of the current instance. Unit: %	0-100	GeminiDB Redis instance nodes	1 minute
nosql002_mem_usage	Memory Usage	Memory usage of the monitored system Unit: %	0-100	GeminiDB Redis instance nodes	1 minute
nosql001_cpu_usage	CPU Usage	CPU usage of the monitored system Unit: %	0-100	GeminiDB Redis instance nodes	1 minute

Table 8-5 String command metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis107_decr_qps	DECR QPS	QPS when a node executes the DECR command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis106_decr_p99	DECR p99 Latency	p99 latency when a node executes the DECR command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis104_decr_avg_usec	DECR Average Latency	Average latency when a node executes the DECR command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis103_incr_qps	INCR QPS	QPS when a node executes the INCR command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis102_incr_p99	INCR p99 Latency	p99 latency when a node executes the INCR command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis100_incr_avg_usec	INCR Average Latency	Average latency when a node executes the INCR command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis099_strlen_qps	STRLEN QPS	QPS when a node executes the STRLEN command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis098_strlen_p99	STRLEN p99 Latency	p99 latency when a node executes the STRLEN command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis096_strlen_avg_usec	STRLEN Average Latency	Average latency when a node executes the STRLEN command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis095_substr_qps	SUBSTR QPS	QPS when a node executes the SUBSTR command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis094_substr_p99	SUBSTR p99 Latency	p99 latency when a node executes the SUBSTR command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis092_substr_avg_usec	SUBSTR Average Latency	Average latency when a node executes the SUBSTR command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis091_setrange_qps	SETRANGE QPS	QPS when a node executes the SETRANGE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis090_setrange_p99	SETRANGE p99 Latency	p99 latency when a node executes the SETRANGE command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis088_setrange_avg_usec	SETRANGE Average Latency	Average latency when a node executes the SETRANGE command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis087_getrange_qps	GETRANGE QPS	QPS when a node executes the GETRANGE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis086_getrange_p99	GETRANGE p99 Latency	p99 latency when a node executes the GETRANGE command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis084_getrange_avg_usec	GETRANGE Average Latency	Average latency when a node executes the GETRANGE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis083_mset_qps	MSET QPS	QPS when a node executes the MSET command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis082_mset_p99	MSET p99 Latency	p99 latency when a node executes the MSET command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis080_mset_avg_usec	MSET Average Latency	Average latency when a node executes the MSET command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis079_mget_qps	MGET QPS	QPS when a node executes the MGET command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis078_mget_p99	MGET p99 Latency	p99 latency when a node executes the MGET command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis076_mget_avg_usec	MGET Average Latency	Average latency when a node executes the MGET command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis075_append_qps	APPEND QPS	QPS when a node executes the APPEND command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis074_append_p99	APPEND p99 Latency	p99 latency when a node executes the APPEND command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis072_append_avg_usec	APPEND Average Latency	Average latency when a node executes the APPEND command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis071_getset_qps	GETSET QPS	QPS when a node executes the GETSET command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis070_getset_p99	GETSET p99 Latency	p99 latency when a node executes the GETSET command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis068_getset_avg_usec	GETSET Average Latency	Average latency when a node executes the GETSET command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis067_get_qps	GET QPS	QPS when a node executes the GET command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis066_get_p99	GET p99 Latency	p99 latency when a node executes the GET command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis064_get_avg_usec	GET Average Latency	Average latency when a node executes the GET command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis063_set_qps	SET QPS	QPS when a node executes the SET command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis062_set_p99	SET p99 Latency	p99 latency when a node executes the SET command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis060_set_avg_usec	SET Average Latency	Average latency when a node executes the SET command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Table 8-6 Hashes command metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis159_hscan_qps	HSCAN QPS	QPS when a node executes the HSCAN command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis158_hscan_p99	HSCAN p99 Latency	p99 latency when a node executes the HSCAN command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis156_hscan_avg_usec	HSCAN Average Latency	Average latency when a node executes the HSCAN command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis155_hvals_qps	HVALS QPS	QPS when a node executes the HVALS command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis154_hvals_p99	HVALS p99 Latency	p99 latency when a node executes the HVALS command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis152_hvals_avg_usec	HVALS Average Latency	Average latency when a node executes the HVALS command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis151_hstrlen_qps	HSTRLEN QPS	QPS when a node executes the HSTRLEN command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis150_hstrlen_p99	HSTRLEN p99 Latency	p99 latency when a node executes the HSTRLEN command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis148_hstrlen_avg_usec	HSTRLEN Average Latency	Average latency when a node executes the HSTRLEN command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis147_hhlen_qps	HLEN QPS	QPS when a node executes the HLEN command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis146_hlen_p99	HLEN p99 Latency	p99 latency when a node executes the HLEN command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis144_hlen_avg_usec	HLEN Average Latency	Average latency when a node executes the HLEN command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis143_hkeys_qps	HKEYS QPS	QPS when a node executes the HKEYS command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis142_hkeys_p99	HKEYS p99 Latency	p99 latency when a node executes the HKEYS command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis140_hkeys_avg_usec	HKEYS Average Latency	Average latency when a node executes the HKEYS command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis139_hincrby_qps	HINCRBY QPS	QPS when a node executes the HINCRBY command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis138_hincrby_p99	HINCRBY p99 Latency	p99 latency when a node executes the HINCRBY command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis136_hincrby_avg_usec	HINCRBY Average Latency	Average latency when a node executes the HINCRBY command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis135_hexists_qps	HEXISTS QPS	QPS when a node executes the HEXISTS command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis134_hexists_p99	HEXISTS p99 Latency	p99 latency when a node executes the HEXISTS command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis132_hexists_avg_usec	HEXISTS Average Latency	Average latency when a node executes the HEXISTS command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis131_hgetall_qps	HGETALL QPS	QPS when a node executes the HGETALL command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis130_hgetall_p99	HGETALL p99 Latency	p99 latency when a node executes the HGETALL command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis128_hgetall_avg_usec	HGETALL Average Latency	Average latency when a node executes the HGETALL command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis127_hdel_qps	HDEL QPS	QPS when a node executes the HDEL command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis126_hdel_p99	HDEL p99 Latency	p99 latency when a node executes the HDEL command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis124_hdel_avg_usec	HDEL Average Latency	Average latency when a node executes the HDEL command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis123_hmget_qps	HMGET QPS	QPS when a node executes the HMGET command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis122_hmget_p99	HMGET p99 Latency	p99 latency when a node executes the HMGET command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis120_hmget_avg_usec	HMGET Average Latency	Average latency when a node executes the HMSET command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis119_h mset_qps	HMSET QPS	QPS when a node executes the HMSET command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis118_h mset_p99	HMSET p99 Latency	p99 latency when a node executes the HMSET command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis116_h mset_avg_usec	HMSET Average Latency	Average latency when a node executes the HMSET command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis115_h get_qps	HGET QPS	QPS when a node executes the HGET command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis114_h get_p99	HGET p99 Latency	p99 latency when a node executes the HGET command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis112_hget_avg_us ec	HGET Average Latency	Average latency when a node executes the HGET command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis111_hset_qps	HSET QPS	QPS when a node executes the HSET command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis110_hset_p99	HSET p99 Latency	p99 latency when a node executes the HSET command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis108_hset_avg_us ec	HSET Average Latency	Average latency when a node executes the HSET command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Table 8-7 List command metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis207_ltrim_qps	LTRIM QPS	QPS when a node executes the LTRIM command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis206_ltrim_p99	LTRIM p99 Latency	p99 latency when a node executes the LTRIM command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis204_ltrim_avg_us	LTRIM Average Latency	Average latency when a node executes the LTRIM command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis203_lset_qps	LSET QPS	QPS when a node executes the LSET command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis202_lset_p99	LSET p99 Latency	p99 latency when a node executes the LSET command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis200_lset_avg_usec	LSET Average Latency	Average latency when a node executes the LSET command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis199_lrem_qps	LREM QPS	QPS when a node executes the LREM command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis198_lrem_p99	LREM p99 Latency	p99 latency when a node executes the LREM command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis196_lrem_avg_usec	LREM Average Latency	Average latency when a node executes the LREM command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis195_lrange_qps	LRANGE QPS	QPS when a node executes the LRANGE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis194_lrange_p99	LRANGE p99 Latency	p99 latency when a node executes the LRANGE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis192_lrange_avg_usec	LRANGE Average Latency	Average latency when a node executes the LRANGE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis191_linsert_qps	LINSERT QPS	QPS when a node executes the LINSERT command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis190_linsert_p99	LINSERT p99 Latency	p99 latency when a node executes the LINSERT command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis188_linsert_avg_usec	LINSERT Average Latency	Average latency when a node executes the LINSERT command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis187_index_qps	LINDEX QPS	QPS when a node executes the LINDEX command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis186_index_p99	LINDEX p99 Latency	p99 latency when a node executes the LINDEX command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis184_index_avg_usec	LINDEX Average Latency	Average latency when a node executes the LINDEX command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis183_llen_qps	LLEN QPS	QPS when a node executes the LLEN command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis182_llen_p99	LLEN p99 Latency	p99 latency when a node executes the LLEN command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis180_llen_avg_usec	LLEN Average Latency	Average latency when a node executes the LLEN command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis179_rpoplpush_qps	RPOPLPUSH QPS	QPS when a node executes the RPOPLPUSH command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis178_rpoplpush_p99	RPOPLPUSH p99 Latency	p99 latency when a node executes the RPOPLPUSH command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis176_rpoplpush_avg_usec	RPOPLPUSH Average Latency	Average latency when a node executes the RPOPLPUSH command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis175_rpop_qps	RPOP QPS	QPS when a node executes the RPOP command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis174_rpop_p99	RPOP p99 Latency	p99 latency when a node executes the RPOP command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis172_rpop_avg_usec	RPOP Average Latency	Average latency when a node executes the RPOP command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis171_rpush_qps	R PUSH QPS	QPS when a node executes the R PUSH command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis170_rpush_p99	R PUSH p99 Latency	p99 latency when a node executes the R PUSH command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis168_rpush_avg_usec	R PUSH Average Latency	Average latency when a node executes the R PUSH command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis167_lpop_qps	LPOP QPS	QPS when a node executes the LPOP command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis166_lpop_p99	LPOP p99 Latency	p99 latency when a node executes the LPOP command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis164_lpop_avg_usec	LPOP Average Latency	Average latency when a node executes the LPOP command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis163_lpush_qps	L PUSH QPS	QPS when a node executes the L PUSH command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis162_lpush_p99	L PUSH p99 Latency	p99 latency when a node executes the L PUSH command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis160_lpush_avg_usec	LPU SH Average Latency	Average latency when a node executes the LPU SH command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis470_brpopush_qps	BRPOPLPU SH QPS	QPS when a node executes the BRPOPLPU SH command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis468_brpopush_p99	BRPOPLPU SH p99 Latency	p99 latency when a node executes the BRPOPLPU SH command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis466_brpopush_avg_usec	BRPOPLPU SH Average Latency	Average latency when a node executes the BRPOPLPU SH command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis465_brpop_qps	BRPOP QPS	QPS when a node executes the BRPOP command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis463_brpop_p99	BRPOP p99 Latency	p99 latency when a node executes the BRPOP command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis461_brpop_avg_us	BRPOP Average Latency	Average latency when a node executes the BRPOP command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis460_blpop_qps	BLPOP QPS	QPS when a node executes the BLPOP command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis458_blpop_p99	BLPOP p99 Latency	p99 latency when a node executes the BLPOP command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis456_blpop_avg_us	BLPOP Average Latency	Average latency when a node executes the BLPOP command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Table 8-8 Set command metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis247_srandmember_qps	SRANDMEMBER QPS	QPS when a node executes the SRANDMEMBER command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis246_srandmember_p99	SRANDMEMBER p99 Latency	p99 latency when a node executes the SRANDMEMBER command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis244_srandmember_avg_usec	SRANDMEMBER Average Latency	Average latency when a node executes the SRANDMEMBER command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis243_sdiff_qps	SDIFF QPS	QPS when a node executes the SDIFF command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis242_sdiff_p99	SDIFF p99 Latency	p99 latency when a node executes the SDIFF command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis240_sd diff_avg_use c	SDIFF Average Latency	Average latency when a node executes the SDIFF command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis239_si smember_q ps	SISMEMBE R QPS	QPS when a node executes the SISMEMBE R command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis238_si smember_p 99	SISMEMBE R p99 Latency	p99 latency when a node executes the SISMEMBE R command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis236_si smember_a vg_usec	SISMEMBE R Average Latency	Average latency when a node executes the SISMEMBE R command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis235_si nter_qps	SINTER QPS	QPS when a node executes the SINTER command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis234_sinter_p99	SINTER p99 Latency	p99 latency when a node executes the SINTER command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis232_sinter_avg_usec	SINTER Average Latency	Average latency when a node executes the SINTER command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis231_sunion_qps	SUNION QPS	QPS when a node executes the UNION command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis230_sunion_p99	SUNION p99 Latency	p99 latency when a node executes the UNION command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis228_sunion_avg_usec	SUNION Average Latency	Average latency when a node executes the UNION command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis227_srem_qps	SREM QPS	QPS when a node executes the SREM command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis226_srem_p99	SREM p99 Latency	p99 latency when a node executes the SREM command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis224_srem_avg_us	SREM Average Latency	Average latency when a node executes the SREM command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis223_smembers_qps	SMEMBERS QPS	QPS when a node executes the SMEMBERS command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis222_smembers_p99	SMEMBERS p99 Latency	p99 latency when a node executes the SMEMBERS command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis220_smembers_avg_usec	SMEMBERS Average Latency	Average latency when a node executes the SMEMBERS command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis219_scard_qps	SCARD QPS	QPS when a node executes the SCARD command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis218_scard_p99	SCARD p99 Latency	p99 latency when a node executes the SCARD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis216_scard_avg_usec	SCARD Average Latency	Average latency when a node executes the SCARD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis215_spop_qps	SPOP QPS	QPS when a node executes the SPOP command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis214_spop_p99	SPOP p99 Latency	p99 latency when a node executes the SPOP command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis212_spop_avg_usec	SPOP Average Latency	Average latency when a node executes the SPOP command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis211_sadd_qps	SADD QPS	QPS when a node executes the SADD command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis210_sadd_p99	SADD p99 Latency	p99 latency when a node executes the SADD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis208_sadd_avg_usec	SADD Average Latency	Average latency when a node executes the SADD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis485_sdiffstore_qps	SDIFFSTORE QPS	QPS when a node executes the SDIFFSTORE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis483_sdiffstore_p99	SDIFFSTORE p99 Latency	p99 latency when a node executes the SDIFFSTORE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis481_sdiffstore_avg_usec	SDIFFSTORE Average Latency	Average latency when a node executes the SDIFFSTORE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis480_sinterstore_qps	SINTERSTORE QPS	QPS when a node executes the SINTERSTORE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis478_sinterstore_p99	SINTERSTORE p99 Latency	p99 latency when a node executes the SINTERSTORE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis476_sinterstore_avg_usec	SINTERSTORE Average Latency	Average latency when a node executes the SINTERSTORE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis475_sunionstore_qps	SUNIONSTORE QPS	QPS when a node executes the SUNIONSTORE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis473_sunionstore_p99	SUNIONSTORE p99 Latency	p99 latency when a node executes the SUNIONSTORE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis471_sunionstore_avg_usec	SUNIONSTORE Average Latency	Average latency when a node executes the SUNIONSTORE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Table 8-9 Zset command metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis315_zremrangebylex_qps	ZREMRANGEBYLEX QPS	QPS when a node executes the ZREMRANGEBYLEX command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis314_zremrangebylex_p99	ZREMRANGEBYLEX p99 Latency	p99 latency when a node executes the ZREMRANGEBYLEX command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis312_zr emrangeby lex_avg_us ec	ZREMRAN GEBYLEX Average Latency	Average latency when a node executes the ZREMRAN GEBYLEX command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis311_zr emrangeby score_qps	ZREMRAN GEBYSCOR E QPS	QPS when a node executes the ZREMRAN GEBYSCOR E command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis310_zr emrangeby score_p99	ZREMRAN GEBYSCOR E p99 Latency	p99 latency when a node executes the ZREMRAN GEBYSCOR E command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis308_zr emrangeby score_avg_ usec	ZREMRAN GEBYSCOR E Average Latency	Average latency when a node executes the ZREMRAN GEBYSCOR E command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis307_zr emrangeby rank_qps	ZREMRAN GEBYRANK QPS	QPS when a node executes the ZREMRAN GEBYRANK command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis306_zr emrangeby rank_p99	ZREMRAN GEBYRANK p99 Latency	p99 latency when a node executes the ZREMRAN GEBYRANK command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis304_zr emrangeby rank_avg_u sec	ZREMRAN GEBYRANK Average Latency	Average latency when a node executes the ZREMRAN GEBYRANK command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis303_z popmin_qps	ZPOPMIN QPS	QPS when a node executes the ZPOPMIN command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis302_zpopmin_p99	ZPOPMIN p99 Latency	p99 latency when a node executes the ZPOPMIN command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis300_zpopmin_avg_usec	ZPOPMIN Average Latency	Average latency when a node executes the ZPOPMIN command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis299_zpopmax_qps	ZPOPMAX QPS	QPS when a node executes the ZPOPMAX command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis298_zpopmax_p99	ZPOPMAX p99 Latency	p99 latency when a node executes the ZPOPMAX command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis296_zpopmax_avg_usec	ZPOPMAX Average Latency	Average latency when a node executes the ZPOPMAX command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis295_zlexcount_qps	ZLEXCOUNT QPS	QPS when a node executes the ZLEXCOUNT command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis294_zlexcount_p99	ZLEXCOUNT p99 Latency	p99 latency when a node executes the ZLEXCOUNT command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis292_zlexcount_avg_usec	ZLEXCOUNT Average Latency	Average latency when a node executes the ZLEXCOUNT command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis291_zrevrank_qps	ZREVRANK QPS	QPS when a node executes the ZREVRANK command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis290_zrevrank_p99	ZREVRANK p99 Latency	p99 latency when a node executes the ZREVRANK command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis288_zrevrank_avg_usec	ZREVRANK Average Latency	Average latency when a node executes the ZREVRANK command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis287_zrank_qps	ZRANK QPS	QPS when a node executes the ZRANK command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis286_zrank_p99	ZRANK p99 Latency	p99 latency when a node executes the ZRANK command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis284_zrank_avg_usec	ZRANK Average Latency	Average latency when a node executes the ZRANK command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis283_zscore_qps	ZSCORE QPS	QPS when a node executes the ZSCORE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis282_zscore_p99	ZSCORE p99 Latency	p99 latency when a node executes the ZSCORE command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis280_zscore_avg_us	ZSCORE Average Latency	Average latency when a node executes the ZSCORE command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis279_zrem_qps	ZREM QPS	QPS when a node executes the ZREM command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis278_zrem_p99	ZREM p99 Latency	p99 latency when a node executes the ZREM command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis276_zrem_avg_us ec	ZREM Average Latency	Average latency when a node executes the ZREM command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis275_zcount_qps	ZCOUNT QPS	QPS when a node executes the ZCOUNT command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis274_zcount_p99	ZCOUNT p99 Latency	p99 latency when a node executes the ZCOUNT command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis272_zcount_avg_u sec	ZCOUNT Average Latency	Average latency when a node executes the ZCOUNT command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis271_zrange_qps	ZRANGE QPS	QPS when a node executes the ZRANGE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis270_zrange_p99	ZRANGE p99 Latency	p99 latency when a node executes the ZRANGE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis268_zrange_avg_usec	ZRANGE Average Latency	Average latency when a node executes the ZRANGE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis267_zrevrange_qps	ZREVRANGE QPS	QPS when a node executes the ZREVRANGE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis266_zrevrange_p99	ZREVRANGE p99 Latency	p99 latency when a node executes the ZREVRANGE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis264_zrevrange_avg_usec	ZREVRANGE Average Latency	Average latency when a node executes the ZREVRANGE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis263_zincrby_qps	ZINCRBY QPS	QPS when a node executes the ZINCRBY command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis262_zincrby_p99	ZINCRBY p99 Latency	p99 latency when a node executes the ZINCRBY command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis260_zincrby_avg_usec	ZINCRBY Average Latency	Average latency when a node executes the ZINCRBY command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis259_zscan_qps	ZSCAN QPS	QPS when a node executes the ZSCAN command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis258_zscan_p99	ZSCAN p99 Latency	p99 latency when a node executes the ZSCAN command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis256_zscan_avg_us	ZSCAN Average Latency	Average latency when a node executes the ZSCAN command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis255_zcard_qps	ZCARD QPS	QPS when a node executes the ZCARD command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis254_zcard_p99	ZCARD p99 Latency	p99 latency when a node executes the ZCARD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis252_zcard_avg_us	ZCARD Average Latency	Average latency when a node executes the ZCARD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis251_zadd_qps	ZADD QPS	QPS when a node executes the ZADD command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis250_zadd_p99	ZADD p99 Latency	p99 latency when a node executes the ZADD command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis248_zadd_avg_us	ZADD Average Latency	Average latency when a node executes the ZADD command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis495_zinterstore_qps	ZINTERSTORE QPS	QPS when a node executes the ZINTERSTORE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis493_zinterstore_p99	ZINTERSTORE p99 Latency	p99 latency when a node executes the ZINTERSTORE command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis491_zinterstore_avg_usec	ZINTERSTORE Average Latency	Average latency when a node executes the ZINTERSTORE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis490_zunionstore_qps	ZUNIONSTORE QPS	QPS when a node executes the ZUNIONSTORE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis488_zunionstore_p99	ZUNIONSTORE p99 Latency	p99 latency when a node executes the ZUNIONSTORE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis486_zunionstore_avg_usec	ZUNIONSTORE Average Latency	Average latency when a node executes the ZUNIONSTORE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Table 8-10 Bitmap command metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis440_bitfield_qps	BITFIELD QPS	QPS when a node executes the BITFIELD command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis438_bitfield_p99	BITFIELD p99 Latency	p99 latency when a node executes the BITFIELD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis436_bitfield_avg_usec	BITFIELD Average Latency	Average latency when a node executes the BITFIELD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis435_bitop_qps	BITOP QPS	QPS when a node executes the BITOP command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis433_bitop_p99	BITOP p99 Latency	p99 latency when a node executes the BITOP command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis431_bitop_avg_us	BITOP Average Latency	Average latency when a node executes the BITOP command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis430_bitpos_qps	BITPOS QPS	QPS when a node executes the BITPOS command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis428_bitpos_p99	BITPOS p99 Latency	p99 latency when a node executes the BITPOS command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis426_bitpos_avg_us	BITPOS Average Latency	Average latency when a node executes the BITPOS command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis425_bitcount_qps	BITCOUNT QPS	QPS when a node executes the BITCOUNT command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis423_bitcount_p99	BITCOUNT p99 Latency	p99 latency when a node executes the BITCOUNT command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis421_bitcount_avg_usec	BITCOUNT Average Latency	Average latency when a node executes the BITCOUNT command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis420_getbit_qps	GETBIT QPS	QPS when a node executes the GETBIT command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis418_getbit_p99	GETBIT p99 Latency	p99 latency when a node executes the GETBIT command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis416_getbit_avg_usec	GETBIT Average Latency	Average latency when a node executes the GETBIT command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis415_setbit_qps	SETBIT QPS	QPS when a node executes the SETBIT command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis413_setbit_p99	SETBIT p99 Latency	p99 latency when a node executes the SETBIT command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis411_setbit_avg_us	SETBIT Average Latency	Average latency when a node executes the SETBIT command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Table 8-11 Stream command metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis590_xreadgroup_qps	XREADGROUP QPS	QPS when a node executes the XREADGROUP command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis588_xreadgroup_p99	XREADGROUP p99 Latency	p99 latency when a node executes the XREADGROUP command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis586_xreadgroup_avg_usec	XREADGROUP Average Latency	Average latency when a node executes the XREADGROUP command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis585_xread_qps	XREAD QPS	QPS when a node executes the XREAD command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis583_xread_p99	XREAD p99 Latency	p99 latency when a node executes the XREAD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis581_xread_avg_usec	XREAD Average Latency	Average latency when a node executes the XREAD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis580_xinfo_qps	XINFO QPS	QPS when a node executes the XINFO command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis578_xinfo_p99	XINFO p99 Latency	p99 latency when a node executes the XINFO command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis576_xinfo_avg_usec	XINFO Average Latency	Average latency when a node executes the XINFO command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis575_xpending_qps	XPENDING QPS	QPS when a node executes the XPENDING command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis573_xpending_p99	XPENDING p99 Latency	p99 latency when a node executes the XPENDING command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis571_xpending_avg_usec	XPENDING Average Latency	Average latency when a node executes the XPENDING command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis570_xclaim_qps	XCLAIM QPS	QPS when a node executes the XCLAIM command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis568_xclaim_p99	XCLAIM p99 Latency	p99 latency when a node executes the XCLAIM command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis566_xclaim_avg_usec	XCLAIM Average Latency	Average latency when a node executes the XCLAIM command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis565_xrevrang_qps	XREVRANGE QPS	QPS when a node executes the XREVRANGE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis563_xr evrange_p99	XREVRANGE p99 Latency	p99 latency when a node executes the XREVRANGE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis561_xr evrange_avg_usec	XREVRANGE Average Latency	Average latency when a node executes the XREVRANGE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis560_xr ange_qps	XRANGE QPS	QPS when a node executes the XRANGE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis558_xr ange_p99	XRANGE p99 Latency	p99 latency when a node executes the XRANGE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis556_xrange_avg_usec	XRANGE Average Latency	Average latency when a node executes the XRANGE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis555_xlen_qps	XLEN QPS	QPS when a node executes the XLEN command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis553_xlen_p99	XLEN p99 Latency	p99 latency when a node executes the XLEN command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis551_xlen_avg_usec	XLEN Average Latency	Average latency when a node executes the XLEN command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis550_xtrim_qps	XTRIM QPS	QPS when a node executes the XTRIM command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis548_xtrim_p99	XTRIM p99 Latency	p99 latency when a node executes the XTRIM command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis546_xtrim_avg_us ec	XTRIM Average Latency	Average latency when a node executes the XTRIM command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis545_xdel_qps	XDEL QPS	QPS when a node executes the XDEL command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis543_xdel_p99	XDEL p99 Latency	p99 latency when a node executes the XDEL command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis541_xdel_avg_us ec	XDEL Average Latency	Average latency when a node executes the XDEL command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis540_xgroup_qps	XGROUP QPS	QPS when a node executes the XGROUP command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis538_xgroup_p99	XGROUP p99 Latency	p99 latency when a node executes the XGROUP command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis536_xgroup_avg_usec	XGROUP Average Latency	Average latency when a node executes the XGROUP command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis535_xack_qps	XACK QPS	QPS when a node executes the XACK command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis533_xack_p99	XACK p99 Latency	p99 latency when a node executes the XACK command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis531_xack_avg_us ec	XACK Average Latency	Average latency when a node executes the XACK command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis530_xadd_qps	XADD QPS	QPS when a node executes the XADD command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis528_xadd_p99	XADD p99 Latency	p99 latency when a node executes the XADD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis526_xadd_avg_us ec	XADD Average Latency	Average latency when a node executes the XADD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Table 8-12 Geo command metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis525_g eopos_qps	GEOPOS QPS	QPS when a node executes the GEOPOS command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis523_g eopos_p99	GEOPOS p99 Latency	p99 latency when a node executes the GEOPOS command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis521_g eopos_avg_ usec	GEOPOS Average Latency	Average latency when a node executes the GEOPOS command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis520_g eodist_qps	GEODIST QPS	QPS when a node executes the GEODIST command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis518_g eodist_p99	GEODIST p99 Latency	p99 latency when a node executes the GEODIST command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis516_g eodist_avg_ usec	GEODIST Average Latency	Average latency when a node executes the GEODIST command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis515_g eohash_qp s	GEOHASH QPS	QPS when a node executes the GEOHASH command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis513_g eohash_p9 9	GEOHASH p99 Latency	p99 latency when a node executes the GEOHASH command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis511_g eohash_av g_usec	GEOHASH Average Latency	Average latency when a node executes the GEOHASH command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis510_g eoradius_q ps	GEORADIU S QPS	QPS when a node executes the GEORADIU S command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis508_georadius_p99	GEORADIUS p99 Latency	p99 latency when a node executes the GEORADIUS command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis506_georadius_avg_usec	GEORADIUS Average Latency	Average latency when a node executes the GEORADIUS command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis505_georadiusby_member_qps	GEORADIUSBYMEMBER QPS	QPS when a node executes the GEORADIUSBYMEMBER command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis503_georadiusby_member_p99	GEORADIUSBYMEMBER p99 Latency	p99 latency when a node executes the GEORADIUSBYMEMBER command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis501_georadiusby_member_avg_usec	GEORADIUSBYMEMBER Average Latency	Average latency when a node executes the GEORADIUSBYMEMBER command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis500_geoadd_qps	GEOADD QPS	QPS when a node executes the GEOADD command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis498_geoadd_p99	GEOADD p99 Latency	p99 latency when a node executes the GEOADD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis496_geoadd_avg_usec	GEOADD Average Latency	Average latency when a node executes the GEOADD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Table 8-13 Hyperloglog command metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis455_pfmerge_qps	PFMERGE QPS	QPS when a node executes the PFMERGE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis453_pfmerge_p99	PFMERGE p99 Latency	p99 latency when a node executes the PFMERGE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis451_pfmerge_avg_usec	PFMERGE Average Latency	Average latency when a node executes the PFMERGE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis450_pfcount_qps	PFCOUNT QPS	QPS when a node executes the PFCOUNT command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis448_pfcount_p99	PFCOUNT p99 Latency	p99 latency when a node executes the PFCOUNT command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis446_pfcount_avg_usec	PFCOUNT Average Latency	Average latency when a node executes the PFCOUNT command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis445_pfadd_qps	PFADD QPS	QPS when a node executes the PFADD command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis443_pfadd_p99	PFADD p99 Latency	p99 latency when a node executes the PFADD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis441_pfadd_avg_usec	PFADD Average Latency	Average latency when a node executes the PFADD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Table 8-14 Pub/Sub command metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis620_p ubsub_qps	PUBSUB QPS	QPS when a node executes the PUBSUB command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis618_p ubsub_p99	PUBSUB p99 Latency	p99 latency when a node executes the PUBSUB command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis616_p ubsub_avg _usec	PUBSUB Average Latency	Average latency when a node executes the PUBSUB command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis615_p unsubscribe_qps	PUNSUBSC RIBE QPS	QPS when a node executes the PUNSUBSC RIBE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis613_punsubscribe_p99	PUNSUBSCRIBE p99 Latency	p99 latency when a node executes the PUNSUBSCRIBE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis611_punsubscribe_avg_usec	PUNSUBSCRIBE Average Latency	Average latency when a node executes the PUNSUBSCRIBE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis610_psubscribe_qps	PSUBSCRIBE QPS	QPS when a node executes the PSUBSCRIBE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis608_psubscribe_p99	PSUBSCRIBE p99 Latency	p99 latency when a node executes the PSUBSCRIBE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis606_ps unsubscribe_avg_usec	PSUBSCRIBE Average Latency	Average latency when a node executes the PSUBSCRIBE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis605_unsubscribe_qps	UNSUBSCRIBE QPS	QPS when a node executes the UNSUBSCRIBE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis603_unsubscribe_p99	UNSUBSCRIBE p99 Latency	p99 latency when a node executes the UNSUBSCRIBE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis601_unsubscribe_avg_usec	UNSUBSCRIBE Average Latency	Average latency when a node executes the UNSUBSCRIBE command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis600_subscribe_qps	SUBSCRIBE QPS	QPS when a node executes the SUBSCRIBE command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis598_subscribe_p99	SUBSCRIBE p99 Latency	p99 latency when a node executes the SUBSCRIBE command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis596_subscribe_avg_usec	SUBSCRIBE Average Latency	Average latency when a node executes the SUBSCRIBE command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis595_publish_qps	PUBLISH QPS	QPS when a node executes the PUBLISH command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis593_publish_p99	PUBLISH p99 Latency	p99 latency when a node executes the PUBLISH command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis591_publish_avg_usec	PUBLISH Average Latency	Average latency when a node executes the PUBLISH command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Table 8-15 Scripting command metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis635_script_qps	SCRIPT QPS	QPS when a node executes the SCRIPT command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis633_script_p99	SCRIPT p99 Latency	p99 latency when a node executes the SCRIPT command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis631_script_avg_usec	SCRIPT Average Latency	Average latency when a node executes the SCRIPT command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis630_evalsha_qps	EVALSHA QPS	QPS when a node executes the EVALSHA command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis628_evalsha_p99	EVALSHA p99 Latency	p99 latency when a node executes the EVALSHA command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis626_evalsha_avg_usec	EVALSHA Average Latency	Average latency when a node executes the EVALSHA command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis625_eval_qps	EVAL QPS	QPS when a node executes the EVAL command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis623_eval_p99	EVAL p99 Latency	p99 latency when a node executes the EVAL command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis621_eval_avg_usec	EVAL Average Latency	Average latency when a node executes the EVAL command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Table 8-16 Transactions command metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis660_discard_qps	DISCARD QPS	QPS when a node executes the DISCARD command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis658_discard_p99	DISCARD p99 Latency	p99 latency when a node executes the DISCARD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis656_discard_avg_usec	DISCARD Average Latency	Average latency when a node executes the DISCARD command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis655_exec_qps	EXEC QPS	QPS when a node executes the EXEC command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis653_exec_p99	EXEC p99 Latency	p99 latency when a node executes the EXEC command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis651_exec_avg_us	EXEC Average Latency	Average latency when a node executes the EXEC command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute
redis650_multi_qps	MULTI QPS	QPS when a node executes the MULTI command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis648_multi_p99	MULTI p99 Latency	p99 latency when a node executes the MULTI command Unit: μs	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis646_multi_avg_usec	MULTI Average Latency	Average latency when a node executes the MULTI command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis645_unwatch_qps	UNWATCH QPS	QPS when a node executes the UNWATCH command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis643_unwatch_p99	UNWATCH p99 Latency	p99 latency when a node executes the UNWATCH command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis641_unwatch_avg_usec	UNWATCH Average Latency	Average latency when a node executes the UNWATCH command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis640_watch_qps	WATCH QPS	QPS when a node executes the WATCH command Unit: count/s	≥ 0	GeminiDB Redis instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
redis638_watch_p99	WATCH p99 Latency	p99 latency when a node executes the WATCH command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute
redis636_watch_avg_usec	WATCH Average Latency	Average latency when a node executes the WATCH command Unit: μ s	≥ 0	GeminiDB Redis instance nodes	1 minute

GeminiDB Influx Metrics

Table 8-17 GeminiDB Influx metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
gemini001_cpu_usage	CPU Usage	CPU usage of the monitored system Unit: Percent	0-100	GeminiDB Influx instance node	1 minute
gemini002_memory_usage	Memory Usage	Memory usage of the monitored system Unit: Percent	0-100	GeminiDB Influx instance node	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
gemini003_bytes_out	Network Output Throughput	Outgoing traffic in bytes per second Unit: kbit/s	≥ 0	GeminiDB Influx instance nodes	1 minute
gemini004_bytes_in	Network Input Throughput	Incoming traffic in bytes per second Unit: kbit/s	≥ 0	GeminiDB Influx instance nodes	1 minute
nosql005_disk_usage	Storage Space Usage	Storage space usage of the monitored object. Unit: Percent	0-100	GeminiDB Influx instances	1 minute
nosql006_disk_total_size	Total Storage Space	Total storage space of the monitored object. Unit: GB	≥ 0	GeminiDB Influx instances	1 minute
nosql007_disk_used_size	Used Storage Space	Used storage space of the monitored object. Unit: GB	≥ 0	GeminiDB Influx instances	1 minute
influxdb001_series_num	Time Series	Total number of time series Unit: count	≥ 0	GeminiDB Influx instance nodes	1 minute
influxdb002_query_reqs_per_sec	Query Requests Per Second	Number of query requests per second Unit: count/s	≥ 0	GeminiDB Influx instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
influxdb003_write_req_ps	Write Requests Per Second	Number of write requests per second Unit: count/s	≥ 0	GeminiDB Influx instance nodes	1 minute
influxdb004_write_points_ps	Write Points	Number of write points per second Unit: count/s	≥ 0	GeminiDB Influx instance nodes	1 minute
influxdb005_write_concurrency	Concurrent Write Requests	Number of concurrent write requests Unit: count	≥ 0	GeminiDB Influx instance nodes	1 minute
influxdb006_query_concurrency	Concurrent Queries	Number of concurrent query requests Unit: count	≥ 0	GeminiDB Influx instance nodes	1 minute

GeminiDB Cassandra Metrics

Table 8-18 GeminiDB Cassandra metrics

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
nosql005_disk_usage	Storage Space Usage	Storage space usage of the monitored object. Unit: Percent	0-100	GeminiDB Cassandra instances	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
nosql006_disk_total_size	Total Storage Space	Total storage space of the monitored object. Unit: GB	≥ 0	GeminiDB Cassandra instances	1 minute
nosql007_disk_used_size	Used Storage Space	Used storage space of the monitored object. Unit: GB	≥ 0	GeminiDB Cassandra instances	1 minute
nosql009_dfv_write_delay	Storage Write Latency	Average delay of writing data to the storage layer in a specified period Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
nosql010_dfv_read_delay	Storage Read Latency	Average latency of reading data from the storage layer in a specified period Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra001_cpu_usage	CPU Usage	CPU usage of an instance Unit: Percent	0-100	GeminiDB Cassandra instance nodes	1 minute
cassandra002_mem_usage	Memory Usage	Memory usage of the instance Unit: Percent	0-100	GeminiDB Cassandra instance nodes	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cassandra003_bytes_out	Network Output Throughput	Outgoing traffic in bytes per second Unit: byte/s	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra004_bytes_in	Network Input Throughput	Incoming traffic in bytes per second Unit: byte/s	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra014_connections	Active Node Connections	Total number of connections attempting to connect to Cassandra instance nodes Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra015_read_latency	Average Read Latency	Average amount of time consumed by read requests Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra016_write_latency	Average Write Latency	Average amount of time consumed by write requests Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra037_pending_write	Suspended Write Tasks	Number of write tasks in waiting status Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cassandra038_pending_read	Suspended Read Tasks	Number of read tasks in waiting status Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra044_range_slice_latency	Scan Duration	Average amount of time consumed by scan operations Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra049_dropped_mutation	Dropped Writes	Average number of dropped writes Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra052_dropped_read	Dropped Reads	Average number of dropped reads Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra092_load_info	Data Volume on a Node	Data volume on a node Unit: byte	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra093_write_count_latency	Accumulated Write Requests	Total number of write requests initiated by a node Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra094_write_1min_rate	Average Write Rate in the Last Minute	Average write rate in the last minute Unit: count/s	≥ 0	GeminiDB Cassandra instance nodes	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cassandra095_write_p75_latency	p75 Write Latency	p75 write latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra096_write_p95_latency	p95 Write Latency	p95 write latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra097_write_p99_latency	p99 Write Latency	p99 write latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra098_read_count_latency	Accumulated Read Requests	Total number of read requests initiated by a node Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra099_read_1min_rate	Average Read Rate in the Last Minute	Average read rate in the last minute Unit: count/s	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra100_read_p75_latency	p75 Read Latency	p75 read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra101_read_p95_latency	p95 Read Latency	p95 read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra102_read_p99_latency	p99 Read Latency	p99 read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute

Metric ID	Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cassandra103_range_slice_count_latency	Accumulated Range Read Requests	Accumulated range read requests Unit: count	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra104_range_slice_1min_rate	Average Range Read Rate in the Last Minute	Average range read rate in the last minute Unit: count/s	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra105_range_slice_p75_latency	p75 Range Read Latency	p75 range read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra106_range_slice_p95_latency	p95 Range Read Latency	p95 range read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra107_range_slice_p99_latency	p99 Range Read Latency	p99 range read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra163_write_p999_latency	p999 Write Latency	p999 write latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute
cassandra164_read_p999_latency	p999 Read Latency	p999 read latency Unit: ms	≥ 0	GeminiDB Cassandra instance nodes	1 minute

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

GeminiDB Mongo Metrics

Table 8-19 GeminiDB Mongo metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
nosql001_cpu_usage	CPU Usage	CPU usage of the monitored system Unit: Percent	0-100	GeminiDB Mongo instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
nosql002_mem_usage	Memory Usage	Memory usage of the monitored system Unit: Percent	0-100	GeminiDB Mongo instance nodes	1 minute
nosql003_bytes_out	Network Output Throughput	Outgoing traffic in bytes per second Unit: bytes/s	≥ 0	GeminiDB Mongo instance nodes	1 minute
nosql004_bytes_in	Network Input Throughput	Incoming traffic in bytes per second Unit: bytes/s	≥ 0	GeminiDB Mongo instance nodes	1 minute
nosql005_disk_usage	Storage Space Usage	Storage space usage of the monitored object. Unit: Percent	0-100	GeminiDB Mongo instances	1 minute
nosql006_disk_total_size	Total Storage Space	Total storage space of the monitored object. Unit: GB	≥ 0	GeminiDB Mongo instances	1 minute
nosql007_disk_used_size	Used Storage Space	Used storage space of the monitored object. Unit: GB	≥ 0	GeminiDB Mongo instances	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
mongodb001_command_ps	COMMAND Statements per Second	Number of COMMAND statements executed per second Unit: count/s	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb002_delete_ps	DELETE Statements per Second	Number of DELETE statements executed per second Unit: count/s	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb003_insert_ps	INSERT Statements per Second	Number of INSERT statements executed per second Unit: count/s	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb004_query_ps	QUERY Statements per Second	Number of QUERY statements executed per second Unit: count/s	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb005_update_ps	UPDATE Statements per Second	Number of UPDATE statements executed per second Unit: count/s	≥ 0	GeminiDB Mongo instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
mongodb006_getmore_ps	GETMORE Statements per Second	Number of GETMORE statements executed per second Unit: count/s	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb007_connections	Current Active Connections	Total number of connections attempting to connect to DB instance nodes Unit: count	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb007_connections_usage	Percentage of Active Node Connections	Percentage of the number of connections that attempt to connect to the instance node to the total number of available connections Unit: Percent	0-100	GeminiDB Mongo instance nodes	1 minute
mongodb008_mem_resident	Resident Memory	Size of resident memory The unit is MB.	≥ 0	GeminiDB Mongo instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
mongodb009_mem_virtual	Virtual Memory	Size of the virtual memory The unit is MB.	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb010_regular_asserts_ps	Regular Asserts per Second	Number of regular asserts per second Unit: count/s	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb011_warning_asserts_ps	Warning Asserts per Second	Number of warning asserts per second Unit: count/s	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb012_msg_asserts_ps	Message Asserts per Second	Number of message asserts per second Unit: count/s	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb013_user_asserts_ps	User Asserts per Second	Number of user asserts per second Unit: count/s	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb014_queues_total	Operations Queued Waiting for a Lock	Number of operations queued waiting for the lock Unit: count	≥ 0	GeminiDB Mongo instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
mongodb015_queues_readers	Operations Queued Waiting for a Read Lock	Number of operations queued waiting for the read lock Unit: count	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb016_queues_writers	Operations Queued Waiting for a Write Lock	Number of operations queued waiting for the write lock Unit: count	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb017_page_faults	Page Faults	Number of page faults on the monitored nodes Unit: count	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb018_porfling_num	Slow Queries	Number of slow queries on the monitored nodes Unit: count	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb019_cursors_open	Maintained Cursors	Number of maintained cursors on the monitored nodes Unit: count	≥ 0	GeminiDB Mongo instance nodes	1 minute
mongodb020_cursors_timeout	Timeout Cursors	Number of timed out cursors on the monitored nodes Unit: count	≥ 0	GeminiDB Mongo instance nodes	1 minute

Dimensions

Key	Value
cassandra_cluster_id,cassandra_node_id	Cluster ID or node ID of the GeminiDB Cassandra instance
redis_cluster_id,redis_node_id	Cluster ID or node ID of the GeminiDB Redis instance
influxdb_cluster_id,influxdb_node_id	Cluster ID or node ID of the GeminiDB Influx instance
mongodb_cluster_id,mongodb_node_id	Cluster ID or node ID of the GeminiDB Mongo instance

8.6 Events Supported by Event Monitoring

Table 8-20 Events Supported by Event Monitoring for GeminiDB

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
NoSQL	Instance creation failure	NoSQL Create Instance Failed	Major	The instance quota or underlying resources are insufficient.	Release unnecessary instances and try again. You can also choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to adjust the quota.	Instances fail to be created.

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Specifications change failure	NoSQL Resize Instance Failed	Major	The underlying resources are insufficient.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console. Submit a service ticket to the O&M personnel to coordinate resources in the background and change the specifications again.	Services are interrupted.
	Node adding failure	NoSQL Add Nodes Failed	Major	The underlying resources are insufficient.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console. Submit a service ticket to O&M personnel to coordinate resources in the background, delete nodes that failed to be added, and add the nodes again.	None
	Node deletion failure	NoSQL Delete Nodes Failed	Major	Releasing underlying resources failed.	Delete the node again.	None

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Storage space scale-up failure	NoSQL ScaleUpStorageFailed	Major	The underlying resources are insufficient.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console. Submit a service ticket to O&M personnel to coordinate resources in the background and scale up storage again.	Services may be interrupted.
	Password resetting failure	NoSQL ResetPasswordFailed	Major	Resetting the password times out.	Reset the password again.	None
	Parameter template change failure	NoSQL UpdateInstanceParameterGroupFailed	Major	Changing a parameter template times out.	Change the parameter template again.	None
	Backup policy configuration failure	NoSQL SetBackupPolicyFailed	Major	The database connection is abnormal.	Configure the backup policy again.	None
	Manual backup creation failure	NoSQL CreateManualBackupFailed	Major	The backup files fail to be exported or uploaded.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	Data cannot be backed up.

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Automated backup creation failure	NoSQL CreateAutomatedBackupFailed	Major	The backup files fail to be exported or uploaded.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	Data cannot be backed up.
	Instance status abnormal	NoSQL FaultyDBInstance	Major	This event is a key alarm event and is reported when an instance is faulty due to a disaster or a server failure.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	The database service may be unavailable.
	Instance status recovery	NoSQL DBInstanceRecovered	Major	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

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Node status abnormal	NoSQL FaultyDBNode	Major	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 functional. Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	The database service may be unavailable.
	Node status recovery	NoSQL DBNodeRecovered	Major	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 switchover or failover	NoSQL PrimaryStandbySwitched	Major	This event is reported when a primary/secondary switchover or a failover is triggered.	No further action is required.	None

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Occurrence of hotspot partitioning keys	HotKey Occurs	Major	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.	<ol style="list-style-type: none"> 1. Choose a proper partition key. 2. Add service cache so that service applications read hotspot data from the cache first. 	The service request success rate is affected, and the cluster performance and stability deteriorates.
	BigKey occurrence	BigKey Occurs	Major	The primary key design is improper. There are too many records or too much data in a single partition, causing load imbalance on nodes.	<ol style="list-style-type: none"> 1. Choose a proper partition key. 2. Add a new partition key for hashing data. 	As more and more data is stored in the partition, cluster stability deteriorates.

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Insufficient storage space	NoSQL RiskyDataDiskUsage	Major	The storage space is insufficient.	Scale up storage space. For details, see section "Scaling Up Storage Space" in the user guide of GeminiDB.	The instance is set to read-only and data cannot be written to the instance.
	Data disk expanded and being writable	NoSQL DataDiskUsageRecovered	Major	The data disk has been expanded and becomes writable.	No further action is required.	None
	Index creation failure	NoSQL CreateIndexFailed	Major	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 created or is incomplete. Delete the index and create a new one.

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Write speed decrease	NoSQL Stalling Occurs	Major	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.	<ol style="list-style-type: none"> 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 success rate of service requests is affected.
	Data write stopped	NoSQL Stopping Occurs	Major	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.	<ol style="list-style-type: none"> 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 success rate of service requests is affected.
	Database restart failure	NoSQL Restart DBFailed	Major	The instance status is abnormal.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	The instance status may be abnormal.

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Restoration to new instance failure	NoSQL Restore ToNew Instance Failed	Major	The underlying resources are insufficient.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console. Submit a service ticket to O&M personnel to coordinate resources in the background and add nodes again.	Data cannot be restored to a new instance.
	Restoration to existing instance failure	NoSQL Restore ToExisting Instance Failed	Major	The backup file fails to be downloaded or restored.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	The current instance may be unavailable.
	Backup file deletion failure	NoSQL DeleteBackup Failed	Major	The backup files fail to be deleted from OBS.	Delete the backup files again.	None

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Failure to display slow query logs in plaintext	NoSQL SwitchSlowlog PlainTextFailed	Major	The DB API does not support this function.	Refer to the <i>GeminiDB User Guide</i> 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
	EIP binding failure	NoSQL BindEipFailed	Major	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 instance cannot be accessed from a public network.
	EIP unbinding failure	NoSQL UnbindEipFailed	Major	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 modification failure	NoSQL ModifyParameterFailed	Major	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

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Parameter template application failure	NoSQL ApplyParameterGroupFailed	Major	The instance status is abnormal. So, the parameter template cannot be applied.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	None
	Enabling or disabling SSL failure	NoSQL SwitchSSLFailed	Major	Enabling or disabling SSL times out.	Try again or submit a service ticket. Do not change the connection mode.	The SSL connection mode cannot be changed.
	Too much data in a single row	LargeRowOccurs	Major	If there is too much data in a single row, queries may time out, causing faults like OOM error.	<ol style="list-style-type: none"> 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. Check whether there are abnormal writes or coding, causing large rows. 	If there are too many records in a single row, cluster stability will deteriorate as the data volume increases.